# Why there is Only One Thing

## A Defence of Ontological Monism

*By*

David Michael Cornell

Submitted in accordance with the requirements for the
degree of Doctor of Philosophy

The University of Leeds

School of Philosophy, Religion and the History of Science

May 2013

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

*For my wife, Gemma, and our soon-to-be-born twins*

# Abstract

In this dissertation I will present and defend Ontological Monism: the thesis which states that there is only a single concrete object in existence – the world itself. According to this view, the world is mereologically simple; it has no proper parts. Because of this, monism may initially strike one as being absurd, for it may be thought that it is simply *obvious* that there is a plurality of concrete objects. There are trees, rocks, cats and larks; there are galaxies, planets, leptons and quarks. If monism denies this then it is blatantly false and, worse, patently absurd, or so one may think. It is true, I concede, that monism denies the existence of all these things. But it is false, I claim, that this renders it absurd. It is the purpose of this thesis to show that monism is actually a coherent, plausible and attractive metaphysical view. Indeed, it is the purpose of this thesis to convince the reader that monism is *true.*

The thesis will progress in a number of stages. I will begin in §1 by looking at the metaphysics of composite objects. Specifically, I will be revisiting Peter van Inwagen's Special Composition Question; the question which asks when two or more material objects compose a further material object. In §2, §3, and §4, I will consider various responses to this question, and conclude that the correct answer to the SCQ is 'never'. That is, I will argue for compositional nihilism, the view that there are no composite objects at all. In §5 I will present a number of arguments to show that monism represents the best form of nihilism. If you're a nihilist, you should also be a monist, or so I will claim. In §6 I will provide some independent support for the possibility of there being spatially extended, yet mereologically simple objects. In §7 I will explain, in detail, how the monist can explain the appearance of plurality given that there is, in fact, only singularity. I will introduce, explain, and defend a new type of property: irreducible structured distributional properties. Armed with this type of property, the monist can provide a satisfying explanation of why it seems as though there are many things, when there is in fact only one. In §8 I will address some objections that have been presented against monism, and show how they are best overcome.

# Acknowledgments

In the course of writing this thesis I have become indebted to a number of individuals and institutions. To begin with the institutions, I would like to thank the AHRC for generously funding the first three years of my study. I would also like to thank the Royal Institute of Philosophy for helping to fund my fourth year with the kind provision of a Jacobsen Bursary. Finally, I would like to thank the University of Leeds. The philosophy department at Leeds is a vibrant and stimulating place in which to conduct research, and there is no doubt that the quality of my work has benefited from being a part of it. The Brotherton Library also deserves a mention. It truly is a fantastic library and provides a wonderfully serene environment in which to work, providing, that is, one can find somewhere to sit. I have thoroughly enjoyed my time at Leeds and will always look back on it with great fondness.

Undoubtedly the best thing about Leeds, however, is the people who work and study there. I owe a great deal of these people a great deal of thanks. First of all, I would like to thank my supervisors, Robin Le Poidevin and Ross Cameron. Both have trawled tirelessly through my work over the past four years and have never failed to provide helpful, constructive and insightful criticism. Although theirs has been the job of highlighting the errors in my work (which is no insignificant task), both have been supportive and enthusiastic throughout, and have always helped me remain positive about my project. Without their help this thesis would have been immeasurably worse than it is; indeed, I doubt it would ever have been completed. I would also like to extend my thanks to all the other members of staff in the philosophy department at Leeds. It has been a pleasure and a privilege to work in the company of such fine minds. I just hope that some of their intellectual pedigree has rubbed off on me. I also want to thank all my fellow post-graduate students at Leeds. In particular, Sarah Adams, Danielle Adams, Jon Banks, Jordan Bartol, Becky Bowd, Thomas Brouwer, Mirja Holst, Wouter Kalf, Kerry McKenzie, Henry Merivale, Alexander Oldermeier, Will Perry, Robert Pezet, Levno Plato, Dave Race, and Carl Warom. These people, amongst others, make the post-graduate community at Leeds a truly thriving one, of which it has been a pleasure to have been a part. Finally, I would like to single out Mike Bench-Capon and Richard Caves, and thank them for regularly indulging me in lengthy discussions about all manner of metaphysical obscurities. Their philosophical insight has been invaluable.

Turning to the non-academic side of things, I would first like to thank my parents, Mary and Tim, for always being there for me. When writing a PhD, it can sometimes seem like one is fighting a never-ending, and losing, battle. At such times, some reassuring words from a familiar voice can be of immeasurable help. More often than not, those words of

reassurance came from them. Last, but most definitely not least, I would like to thank my wonderful wife, Gemma. Throughout the duration of my study, she has shown me nothing but love and support, which has been invaluable on helping me along that rocky road which is postgraduate study. She has also reminded me, and helped me, to try and keep in touch with reality; something which one can often lose sight of when immersed in the study of philosophy and the company of philosophers. For all of that, I will always be grateful.

# Abridged Table of Contents

## Part I – Material Composition



## Part II - Monism



## Part III – Objections and Replies

# Contents

# Part I

# Material Composition

# §1. Framing the Debate

## §1.1. The Special Composition Question

Suppose that you have five bits of wood: four lengths of, say, one metre each, and one flat, square board of, let's say, one metre-squared. Now suppose that you arrange these bits in such a way that the board is horizontal (parallel to the ground) and is being supported in each corner by one of the four lengths, each of which is standing vertically upright on the (level) ground. Now add the further supposition that, due to your excellent carpentry skills, the structure has a good degree of rigidity and stability. Now consider the following question:

1) Do your five pieces of wood compose a new object – let's call it a table?

Now suppose that you have a number of grains of rice, let's say it's precisely one thousand. You place these grains, keeping them as close together as possible, on top of the wooden structure you have just assembled. Now consider the following question:

2) Do your one thousand grains of rice compose a new object – let's call it a heap?

Standing back, you admire your morning's work. In front of you there are five pieces of wood, arranged in a certain way, and one thousand grains of rice, grouped together on top. Now consider the following, and final, question:

3) Do *all* these objects (the five bits of wood *and* the one thousand grains of rice) compose a new object – let's call it a heap-table?

These questions are perfectly straightforward to understand and, on the face of it at least, perfectly straightforward to answer. The answer to question one would appear to be an obvious 'yes'. A wooden structure with four legs and a flat, horizontal top, sturdy enough to

put things on *just is* a table. So *of course* our five bits of wood compose a table. Question two may not be quite so clear-cut, but one would normally, I assume, be inclined to say that the one thousand grains of rice *do* compose a heap of rice. We would certainly *call* it a heap, after all. So if it weren't a heap, why would we call it so? Common sense, then, tells us that the rice *does* compose a heap. So according to common sense, both questions one and two would be answered in the affirmative. The bits of wood compose a table; the grains of rice compose a heap. Question three, however, is quite different. The answer is just as obvious, it seems, yet it is quite the opposite. That is to say, according to common sense or intuition, we would surely want to say 'no'. There is no object composed of the bits of wood *and* the rice – there is no such thing as a heap-table! Instead, I will assume, we would normally be inclined to say there are two separate objects, a heap of rice and a table. It is simply that the former object is *on* the latter.

So far I have not said anything controversial, or so I will assume. In fact, all I have said so far seems quite obvious. Yet surprisingly (or perhaps *un*surprisingly, depending on your view of philosophers) the vast majority of philosophers who have turned their attention to this subject disagree with at least *some* of these statements. The reason for disagreeing is usually grounded in a complete disability to answer a crucial question: *why*? Why do the pieces of wood compose a new object, and the grains of rice compose a new object, yet the wood and the rice *taken together* do not compose a new object? Is there something special about the way these objects are arranged, or grouped, that makes it such that there is a table and a heap, but no heap-table? Is there some systematic and general rule which we can apply in order to determine when certain objects *do* compose a further object, and when they do not? If so, what is it? If not, then why do questions of material composition seem, on the face of it at least, so easy to answer? And indeed, just how do we arrive at these answers which seem so obviously apparent?

In his 1987 article and, at greater length, his 1990 book, Peter van Inwagen grapples with questions of exactly this sort.[1] His research makes it apparent that questions about material composition, such as those mentioned above, are much more problematic and difficult to answer than they may initially seem. In pursuit of a solution, he proposes what he calls the Special Composition Question (SCQ):

[1] van Inwagen (1987; 1990)

(SCQ): Under what conditions does composition occur?[2]

A satisfactory answer to this question should tell us why our five bits of wood compose a table, and our grains of rice compose a heap, yet why the rice and wood taken together do not compose a heap-table. Indeed a satisfactory answer to the SCQ should give us a *general and universal* rule, which would be applicable in *any* case where there are two or more material objects, to tell us whether those objects compose a further object or not.[3] What van Inwagen's work shows is that it is actually very difficult to provide an answer to the SCQ that accommodates our common sense intuitions about material composition. That is to say, it is very difficult to justify and defend an answer to the SCQ which allows that our pieces of wood compose a table, and that our grains of rice compose a heap, yet which at the same time denies that the wood and rice taken together compose a heap-table. Indeed van Inwagen finds it *so* difficult to defend such an answer, that he concludes that our intuitions must in fact be mistaken. Instead, he provides an answer to the SCQ that runs completely counter to common sense; for he claims that the *only* cases where composition occurs are those in which the object composed is a living organism. According to this answer, our five pieces of wood *do not* compose a table, and our grains of rice *do not* compose a heap. According to van Inwagen there are no tables or heaps – indeed there are no inanimate composite objects at all. Van Inwagen calls this somewhat startling conclusion "the denial".[4]

In the following four chapters (including this one), I will provide a thorough evaluation of the SCQ and its potential answers. First off, I will make a brief attempt to *motivate* the SCQ. That is, I will outline why I think it is an important question in the first place. I will then assess and evaluate a number of candidate answers that could be (and have been) given to the SCQ. Since van Inwagen formulated the SCQ just over twenty years ago, there has been a growing body of literature (which is now vast) dedicated to answering it. Because of this, not much of what I will say in the next few sections will be entirely new, but rather, will be going over ground that has already been covered. It is still necessary that this

[2] van Inwagen (1990), 39

[3] It has been suggested that there is *no* systematic and general answer to the SCQ (see Markosian 1998b; 2008). This approach will be considered, and dismissed, in §2.5.

[4] Van Inwagen (1990), 1

is done, however, for two main reasons. Firstly, since my chosen answer to the SCQ will play a big part in shaping the metaphysic to be fleshed out in this thesis, it is important that I show why I have chosen *that* particular answer and rejected others. I will do this by showing which arguments in this debate are strongest, and why, and explaining why I give more weight to some considerations than I do to others. Secondly, despite there being a wealth of literature on the SCQ, there is not (to my knowledge) a complete, thorough, *and up-to-date* review of all the significant extant positions, arguments and counter-arguments. A comprehensive review of this type would, I think, be of benefit to the philosophical landscape, and so in what follows, that is what I will aim to provide.

## §1.2. Why is the SCQ Important?

In the following few sections I will be dedicating a significant amount of time to the SCQ. I will be considering a variety of candidate answers, pushing them to their logical conclusions, and weighing up the costs and benefits of each in turn. The hope is to find an answer which I – and hopefully the reader too – will find satisfactory. Moreover, the answer on which I shall settle will play a significant role in determining not only the shape of this thesis, but also the shape of my metaphysical view of the world in general. It would be fair to say, then, that I consider the SCQ to be an important question, and one that is in need of an answer.

Not everyone shares this view. Firstly, for the philosophically uninclined – those who are often referred to by philosophers, as 'the folk'[5] – the SCQ may have never even been considered. Your standard member of the non-philosophical community is unlikely to have ever considered whether heap-tables exist or not. But even if the SCQ were brought to the attention of such a person, I suspect it would not be taken too seriously. Indeed I suspect that it may well be met with an incredulous stare. Secondly, even within the philosophical community there are those who dismiss the SCQ. For a variety of reasons, these philosophers claim the SCQ (as well as certain other ontological questions) to be either

[5] I can't help but feel that the term 'the folk' is distastefully patronizing, so I won't be employing it in what follows.

unimportant, unanswerable, or even incoherent.[6] I will not be tackling these deflationary views in this thesis, since it would take the discussion too far afield from where I want it to go. However, the mere presence of dissenting views suggests that a defence of the importance of the SCQ is required. In what follows, I will sketch out such a defence, and thus justify why I think the SCQ is worth grappling with. My comments will be quite loose at this stage, and I certainly don't expect them to pacify the concerns of the deflationist philosophers. However, whether one agrees with my remarks or not, they should make clear, or at least understandable, to philosophers and non-philosophers alike, my motivation for treating the SCQ seriously.

My motivation ultimately consists of the conjunction of two fundamental beliefs. The first is my conviction in a robust realism about ontology. I am very much committed to the view that the external world exists mind-independently. If rocks exist, then they would continue to exist if there were no minds, and they would have existed if there had never been any minds. I am in agreement with van Inwagen, when he says: "what there is, is never a matter of stipulation or convention".[7] The second belief is about the goal of metaphysics itself. I firmly believe that the ultimate goal of metaphysics is to discover *what the world is really like.* In a recent paper, Ted Sider says "the central task of metaphysics is illuminating the fundamental structure of reality".[8] Whilst the phrasing may sound a little grandiose, I think Sider is bang on the money. The world has an objective structure and character; it is the job of the metaphysician to discover what it is.

With that in mind, then, we could say that in an ideal world, the ultimate fruit of metaphysical inquiry would be a *perfect theory of the world.* A theory that explained completely and exhaustively the nature of reality; a theory such that if you understood it in its entirety, there would be nothing left for you to know, and nothing left for you to understand, about what the world is like. Such a theory is, of course, a thing of fantasy.[9] But

[6] Hirsch (2002; 2005; 2009) and Schaffer (2009), 358-359, for example, both attempt to deflate the importance of the SCQ; Hirsch does so by claiming it reduces to a merely verbal dispute rather than representing a genuine ontological one; Schaffer does so by claiming it to have a trivial and obvious answer. Chalmers (2009) claims that there may well be no determinate answer to the SCQ, whereas Thomasson (2009) suggests that the SCQ (and other related questions) may be defective, and thus entirely unanswerable. These are just but a few of the current crop of 'ontological deflationists'.

[7] Van Inwagen (1990), 7

[8] Sider (2009), 401

[9] I mean it is unobtainable from our humble human perspective. It may be possible that some God-like being could be omniscient in the way described.

this does not mean that we can't strive towards it. Furthermore, even if the perfect theory is unobtainable, we can still speculate about what sort of features it should have and what sort of phenomena it should explain. And it seems quite obvious to me that a crucial part of the perfect theory must be a list of what there is; a complete and exhaustive inventory of all the things that exist. After all, how could one profess to have a perfect knowledge of the world if one did not even know what was in it? I cannot argue for these claims other than to say simply that they form part of the bedrock of my basic philosophical beliefs. I merely hope that they do not sound unreasonable.

In light of this, then, the import of the SCQ should be apparent. A complete inventory of all the constituents of the world would include all composite material objects (should there be any such things at all). Thus if we were ever to compile such a list we would need to know exactly how to identify those composites. (Imagine trying to count all the sheep in a field, for example, if you did not know how to identify a sheep). And this is exactly what the SCQ purports to do; it aims to establish the necessary and sufficient conditions under which some things are to be considered a composite object. Answering the SCQ, then, is a pre-requisite for compiling an inventory of what there is. And *that* is why the SCQ is important. Although highly theoretical and abstract, it is a fundamental and substantive question which must be answered if we are to get any closer to our ultimate goal of understanding what the world is really like.

## §1.3. Technicalities and Terminology

The SCQ is inextricably bound up with the concepts of *mereology*. Mereology is the study of parthood, or the study of the relation(s) between parts and the wholes they compose, and indeed of the relation(s) between parts themselves. (The word derives from the Greek 'meros' (μερος) meaning 'part' or 'portion'). The fact that mereology is central to this debate should be fairly clear, for another way of formulating the SCQ would be to ask: under what conditions do two or more objects compose a whole of which they are the (only) parts? Or perhaps even: when should a material object be considered a part of another material object?

Mereology has long been of interest to philosophers, but in the last century in particular, it has been developed into a precise and technical field of metaphysical inquiry.[10] Before we begin the nitty-gritty of evaluating answers to the SCQ, then, it should be worthwhile to clarify and define some of the central mereological terms and concepts that will be employed in so doing. This should reduce the possibility of any ambiguity and/or misinterpretation in the subsequent discussion. For the reader who is already metaphysically inclined, these concepts should be fairly familiar, and for the reader already *au fait* with the SCQ, they will probably be more like second nature. It will do no harm, however, to go over them again here. I should note that what follows is by no means intended to be a complete and exhaustive overview of mereology *per se*. That would take an entire thesis in itself.[11] Rather, I will just outline some of the central mereological concepts that will be useful for the purpose of this thesis.

I take parthood to be a primitive relation, in that it is a relation not receptive to reductive analysis. I don't take this to be at all problematic, however, since I am assuming that it is very easy to get an *intuitive grip* on what it means to say that something is a part of something else. Furthermore, despite it being a primitive relation, one can still lay down some precise constraints on the way the term 'part' (or 'is a part of', etc.) is to be applied. Indeed that is just what I shall do.

Firstly, in this thesis I will be using the term 'part' in a much narrower sense than one may use it *per se.* For there are, in fact, many different ways in which the term 'part' can be employed in the English language. Consider, for example, the following sentences:

(1) The steering wheel is a *part* of the car
(2) The ship is a *part* of the fleet
(3) The first movement is a *part* of the symphony
(4) Learning to give and take is a *part* of growing up

[10] The formalising of mereology is largely due to the work Lesniewski (1916) and then later, Leonard and Goodman (1940). But interest in the part-whole relation stretches right back to ancient Greece. In particular, see Plato's *Parmenides,* and Aristotle's *Metaphysics,* Δ, 1023-24.
[11] For an excellent, and in depth study, into mereology, see Simons (1987)

All these sentences seem to express mereological relations of some sort or other, for they all refer to some thing as being a part of something else. But the term 'thing' here is clearly being used quite loosely. In (1), for instance, we have two material objects flanking either side of the 'is a part of' relation. Sentence (1), then, refers to a case of what I shall call the *material parthood relation.* The other sentences are different. In (2), whilst a ship certainly *is* a material object, it is not so clear that the fleet is too. One may, for instance, prefer to think of the fleet as being a *class* or *set* of objects, rather than being a material object in its own right. In (3), the objects related by 'is a part of' are almost certainly *not* material.[12] They are, if they exist at all, most likely to be considered *abstract*. And with regard to (4), whilst a mereological relation does *seem* to be being expressed, it would be with great reluctance that I would admit that the things it purports to relate are actually *things* at all. At the very most, it seems to express a relation between *concepts*.

So there are many and varied applications of the parthood relation, and many different kinds of thing which it can take as its *relata*. But I do not want to get drawn into a debate of how best they should be classified, for that is not really of present concern.[13] What I will say, however, is that in *this* thesis, I am concerned only with what I have called the *material parthood relation.* That is, I am only concerned with the mereology of material objects (indeed that is why I have called this section 'Material Composition' rather than just 'Mereological Composition'). It is important that this distinction is made, because as will become clear, there are differences (albeit subtle ones) between the material parthood relation and the parthood relation *per se*. From here on, then, when I refer to 'parthood', I am referring to 'material parthood', unless otherwise stated.

So now we are in a position to make the parthood relation a little more formal. Firstly, I will assume that the parthood relation is reflexive, transitive and anti-symmetric. So if we let 'P' denote the two-place parthood predicate, such that P$xy$ can be taken to stand for '$x$ is a part of $y$', we have:

*Reflexivity:* $\forall x(\mathrm{P}xx)$

[12] This view is not universally accepted, of course, as there are some who believe that musical works *are* in fact material objects. See, for example, Caplan & Matheson (2006). I take it, however, that this is very much a minority view, and a controversial one at that.

[13] See Winston, Chaffin & Herrmann (1987) for an attempt to classify a taxonomy of the varying types of part-whole relation.

*Transitivity:* $(Pxy \,\&\, Pyz) \rightarrow Pxz$

*Anti-Symmetry:* $(Pxy \,\&\, Pyx) \rightarrow x = y$

The reflexivity condition shows that my understanding of parthood is somewhat different from what you might call the *common sense* understanding of parthood, for it means that every object is a part of itself. This is probably at odds with the ordinary conception of parthood. After all, if you asked someone to make an inventory of all the parts of a car, for instance, it is unlikely that they would include the car itself on the list. The common sense understanding of parthood concords better with what is known technically as 'proper parthood', which I will define in due course.

A brief note about transitivity is in order. Although it is often assumed that parthood is transitive (van Inwagen assumes so, for instance, claiming transitivity to be "a non-negotiable feature of parthood"[14]), it is not universally accepted. And one may be right not to accept it *if* one is concerned with parthood *in its widest possible sense* (as opposed to the restricted sense of *material* parthood). An example from David Sanford shows this nicely. He says: "My spleen is a part of me, and I am a part of this Book Symposium, but my spleen is not a part of this Book Symposium".[15] Sanford is surely correct; I for one would not want to admit that Book Symposiums are the types of things that could have spleens as parts. However, it also seems clear that Book Symposiums are not material objects.[16] So whilst parthood, taken in its widest possible sense, may be non-transitive, it seems much more plausible to insist that it *is* transitive when restricted solely to the domain of material objects. And this is precisely what I shall do.[17]

[14] Van Inwagen (1990), 65

[15] Sanford (1993), 221

[16] And if one disagrees, here, and thinks that book symposiums *are* material objects, then I would have no hesitation in changing tack and admitting that spleens *could* be parts of them.

[17] There have been attempts to show that even the *material* parthood relation is non-transitive. Rescher (1955), for example, suggests that "a part (i.e. a biological sub-unit) of a cell is not said to be a part of the organ of which that cell is a part" (p.10). I don't find this example at all convincing. Whether or not it is 'said' to be a part of the organ is a matter of classification for biologists to dispute, but whether it actually *is* a part (I mean *physically* a part) does not even seem to be up for debate. Having said that, I do actually think that some more plausible examples could be given. Consider: 'the tuning peg is a part of the violin and the violin is a part of the orchestra, but the tuning peg is not a part of the orchestra'. Whilst this example may have *some* degree of plausibility, it is not enough for me to reject the transitivity of material parthood. As such, I am prepared to simply bite the bullet and admit that the tuning peg *is* a part of the orchestra.

So with (material) parthood as our primitive, we can now define some further relevant mereological relations *in terms* of that primitive:

*Proper Parthood (PP):* $PPxy =_{df} Pxy \ \& \sim Pyx$

In other words: *x* is a proper part of *y* just in case *x* is a part of *y* and *y* is not a part of *x.*

(Another way of defining proper parthood would be: $PPxy =_{df} Pxy \ \& \sim(x = y)$)

*Overlap (O):* $Oxy =_{df} \exists z(Pzx \ \& \ Pzy)$

In other words: *x* overlaps *y* just in case they share a part in common.

*Disjointedness (D):* $Dxy =_{df} \sim Oxy \ \& \ x \neq y$

In other words: *x* and *y* are disjoint just in case they are distinct and do not overlap.

The following definitions will help clarify some mereological concepts which are specific to, and indeed *central* to, the debate over material composition. I will follow van Inwagen by making use of "plurally referring expressions" such as 'the *x*s' to refer to a given plurality of objects:[18]

*Composition:* the *x*s compose $y =_{df}$ (i) the *xs* are all parts of *y,* (ii) no two of the *x*s overlap, and (iii) every part of *y* overlaps at least one of the *x*s.

*Fusion: y* is a fusion of the *x*s $=_{df}$ the *x*s compose *y*.

Finally, it is necessary to provide some precise definitions of the various terms I will be employing to describe various types of objects. I am using the term 'object' in its widest sense, such that it refers to any 'thing' or 'entity' you care to think of. 'Object', then, like 'parthood', is taken to be primitive:

[18] These are taken from van Inwagen (1990), 23-29. The use of plurally referring expressions is to increase clarity and neutrality. For instance, if we had some objects, and we wondered whether they composed a further object, we could simply ask "does this plurality of objects compose?" But the term 'plurality' is a noun, a singularly referring term, which would seem to suggest that there is some single object to which it refers. But that is the very thing that is up for debate! And the same goes for other possible terms such as 'set' or 'aggregate' and so on. By using plurally referring expressions, such as 'the *x*s', the terminology is much more neutral.

$x$ is an abstract object $=_{df}$ $x$ is an object which exists neither spatially nor temporally

$x$ is a material object $=_{df}$ $x$ is not an abstract object

$x$ is a material simple $=_{df}$ $x$ is a material object that has no proper parts

$x$ is a material composite $=_{df}$ $x$ is a material object and is not a material simple.

$x$ is an ordinary composite $=_{df}$ $x$ is a material composite that is *ordinarily* recognised by common sense, e.g. tables, chairs, trees, cars, people, etc.

$x$ is an exotic composite $=_{df}$ $x$ is a material composite and is not an ordinary composite, e.g. a heap-table, or the object composed of the Eiffel Tower and the Great Pyramid at Giza, etc.

With these definitions at hand, we are now in a position to formulate a precise version of the SCQ:

(SCQ): For any $x$s, where those $x$s are material objects, what are the jointly necessary and sufficient conditions under which there is some further material object, $y$, which those $x$s compose?

Furthermore, we can now specify exactly what form our desired answer needs to take. It needs to be as follows:

(ANSWER): for any $x$s (where those $x$s are material objects) there is a further material object, $y$, composed of those $x$s if and only if _____.

Importantly, for any purported answer to be deemed satisfactory, it must not make use of mereological terms on the right-hand side of the biconditional, for fear of circularity. To see why this is important, consider the futility of completing the right-hand side of the biconditional with 'there is a $y$ such that $y$ is a composite object of which the $x$s are the proper parts'. Such a response would be entirely trivial, and would illuminate no interesting principles about material composition.

With all the relevant concepts and terminology having been precisely and formally defined, we are now in a position to tackle the Special Composition Question. The following three chapters will be devoted to doing just that.

# §2. Compositional Restrictivism

There is a tripartite division among candidate answers to the SCQ. There are those answers which say composition never occurs (i.e. there are no composite objects), those answers which say composition always occurs (i.e. for any two or more material objects whatsoever, there is *always* a further object which they compose), and those answers which say composition only sometimes occurs (i.e. some material objects compose further objects, and others don't). This triumvirate is complete and exhaustive; no matter what answer you give to the SCQ it will fall under one of these heads.[19] I will use the following terminology to refer to these three types of position on composition[20]:

- ***Compositional Nihilism*** is the view that composition never occurs. According to this view there are no composite material objects; the only material objects in existence are material simples (i.e. objects without proper parts).[21]
- ***Compositional Universalism*** is the view that composition always occurs. According to this view, for any two (or more) material objects, there is always a further material object which they compose.[22]
- ***Compositional Restrictivism*** covers a variety of views, all of which say that composition sometimes, but not always, occurs. restrictivist views can be fleshed out in many different ways, depending on what conditions are stipulated as being necessary and sufficient for composition to occur.[23]

[19] That is, of course, so long as you take the question seriously in the first place. If, for some reason, you thought the question defective, for example, then you most likely won't settle on any of these three answers. But then to say the question is defective is not really to answer it at all.

[20] The term 'nihilism' was first used in this context by Peter Unger (1979a; 1979b). 'Universalism' was coined by van Inwagen (1990), 74. Both terms have now become accepted in the literature, so I seen no reason to use different terminology here. It is unclear exactly where the term 'restrictivism' was first used in this context, but I will stick with it since it too is widely used in the literature on this topic.

[21] Examples of nihilists are: Cameron (2010b); Dorr (2005); Horgan & Potrč (2000; 2008); Sider (forthcoming); Unger (1979a; 1979b; 1979c; 1980a; 1980b).

[22] Examples of universalists are: Armstrong (1997); Hudson (2001); Leonard & Goodman (1940); Lewis (1986a); Rea (1998); van Cleve (1986; 2008); to name but a few.

[23] Examples of restrictivists are: Markosian (1998b; 2008); Merricks (2001); van Inwagen (1990).

It is often claimed that common sense points to the fact that composition is restricted. This was hopefully highlighted by the earlier example about the table and the heap of rice, in that we are (pre-theoretically at least) inclined to accept the existence of ordinary composites, like tables and heaps, but not exotic composites, like heap-tables. It is because of this fact that I will begin my evaluation of composition with an investigation into restrictivist answers to the SCQ. For if our common sense intuitions suggest that composition is restricted, then it would be quite satisfying, I will assume, if we could formulate an answer to the SCQ that tallies up with those intuitions. More specifically, therefore, I will begin not by just looking at any old restrictivist answers, but those restrictivist answers which aim to preserve our common sense intuitions about which composite objects do, and do not, exist. I will call any answer to the SCQ that does preserve these intuitions, (a variant of) *Common Sense restrictivism*.

## §2.1. Common Sense Restrictivism (CSR)

It is widely agreed that common sense points towards a restrictivist stance on composition - and perhaps for good reason. For it appears that one does not need to engage in any serious philosophical reflection to realise that the great pyramid of Giza, for instance, is composed of limestone blocks and that the Eiffel Tower is composed of iron girders. Yet equally obvious is the fact that there is no object which these two great edifices together compose. So since it is plainly evident that there are both cases where objects do compose a further object, and cases where they do not, it may also be thought plainly evident that composition is restricted. This line of thinking commonly crops up in the literature on composition. More often than not, however, it is merely assumed, rather than explicitly argued for.[24] In the interests of maximum clarity, then, it may be worth explicitly formulating this line of thought into an argument for restricted composition. Let's call it the argument from obviousness.

[24] Trenton Merricks, for instance, claims that restrictivism results from the conjunction of "two common sense claims", those claims being that some objects clearly compose (i.e. my atoms compose me) and that some objects clearly do not compose (i.e. my atoms and your atoms do not compose any further object). See Merricks (2005), 615. Similar sentiments are echoed in Markosian (1998b; 2008), and many others. (Note that Merricks himself does not actually endorse common sense restrictivism).

### *§2.1.1 The Argument from Obviousness (AFO)*

1. It is obvious that composition occurs in some cases, e.g. it is obvious that the Eiffel Tower exists and is composed of smaller parts.
2. It is obvious that composition does *not* occur in some cases, e.g. it is obvious that there is no object composed of the Eiffel Tower and the Great Pyramid of Giza.
3. Therefore, it is obvious that composition is restricted.

AFO is deductively valid.[25] Despite this, however, I do not consider it particularly illuminating. There are two main reasons for this. The first reason is that even if the argument is sound (which is doubtful)[26], the truth of the conclusion does not entail the truth of the claim that composition *is* restricted. The conclusion tells us only what we find to be obvious; it says nothing whatsoever about what is actually the case. So if we were to draw any *objective* conclusion about composition from this argument, we would need a further premise stating that obviousness entails truth. Yet I am assuming that such a premise would be dubious to say the very least, on the basis that many things that have at one time appeared obvious have turned out to be false. After all, at one time people presumably thought it perfectly obvious that the earth was flat, and indeed, stationary. (Indeed if obviousness *did* entail truth, then the whole discipline of philosophy would be largely redundant, as there would be little or no call for reflection about anything; instead we would simply accept what immediately *seems* to be the case)!

This is not to say, of course, that there is no correlation whatsoever between what appears obvious and what is actually the case. I am not for a minute suggesting we should treat the obviously apparent and the patently absurd on a perfectly equal footing. If it appears obvious that *p*, then that may well be a *good indication* that it is actually the case that *p* (and likewise, if it appears absurd to suggest that *p*, this may well be an indication that it is

[25] Actually, this could be resisted, because of the part played in the argument by the word 'obvious'. Since 'obvious' is very much a psychological, or intentional, predicate, it might be thought to complicate the logical structure of the argument. I will ignore this technicality here.

[26] I say it is doubtful because the deeper one reflects on questions of composition, the less obvious it becomes as to which cases (if any) of composition are clear-cut, to the extent that both premises 1 and 2 may reasonably be rejected as false.

not the case that *p*). But this is very different from obviousness *entailing truth.* With this in mind, then, we may reasonably concede that AFO gives an *initial indication* that composition is restricted, but we should probably be careful not to read much more into it than that.

There is a second reason why AFO is not particularly illuminating, and that is that even if we accept that it gives an indication that composition is restricted, it doesn't tell us nearly *enough.* That is to say, it certainly doesn't constitute an answer to the SCQ. Yes, it indicates that composition is restricted, in that it sometimes occurs and sometimes doesn't, but it doesn't tell us *under what circumstances* composition occurs. The SCQ, recall, asks for the necessary and sufficient conditions under which some *x*s compose a further object. AFO does not provide us with these conditions; it does not constitute an answer to the SCQ.

All we have then, if common sense is anything to go by, is a hunch; a hunch that composition is restricted. But at this early stage of investigation, a hunch is not a bad thing to go on. Indeed it gives us a good place to start. The challenge now, then, is to formulate a full answer to the SCQ on the basis of this hunch. That is to say, we need to precisely specify the necessary and sufficient conditions under which composition occurs, such that they are satisfied by ordinary composites, yet not satisfied by exotic composites. As will soon become clear, this is no easy task.

### *§2.1.2 Simple Bonding Answers*

In his book, Peter van Inwagen followed a similar chain of reasoning to that given above, and likewise, he too began his attempt to answer the SCQ by looking at answers that attempt to preserve the intuitions of common sense.[27] Ultimately, however, he found all such answers to be entirely unacceptable, as have most philosophers who have subsequently taken on the same challenge.[28] I will not spend too much time considering these common sense answers, partly because van Inwagen's excellent treatment of them is still readily available to the reader who may be interested, and partly because it simply doesn't take much analysis to show that they are all hopeless. But they do need at least a cursory mention, however, in order to demonstrate that whenever we do pursue the common sense

[27] Van Inwagen (1990), 33-38 and 56-61

[28] The one exception I know of is Markosian (1998b). I will come to Markosian's view later.

view on composition, we will inevitably run into some serious metaphysical problems, and as such, when answering the SCQ, we may be forced to look elsewhere.

The first few common sense answers that van Inwagen considers all have one thing in common: they are all based on the strong intuition that in order for any material objects to compose a further material object they must, somehow or other, be bonded together, or at the very least be in contact with one another. For it is surely this intuition, or something very similar to it, that makes one baulk at the idea that there may be an object composed of the Eiffel Tower and the Great Pyramid at Giza. These two objects are separated by a distance of over two thousand miles – how on earth could they be considered the sole parts of a further material object? Spurred by this thought, van Inwagen's first formulation of an answer to the SCQ, which he calls 'contact', is as follows:

> CONTACT: for any *x*s (where the *x*s are material objects), there is a further material object, *y*, composed of those *x*s iff those *x*s are in contact.[29]

Before we can subject CONTACT to any scrutiny, a little clarification is in order, for one may wonder what *exactly* we mean when we say that the *x*s are in contact. Van Inwagen defines the notion of 'being in contact' as follows:

> The *x*s are 'in contact' if they do not spatially overlap and are 'clumped together'. That is, the *x*s are in contact if (1) no two of them overlap spatially, and (2) if *y* and *z* are among the *x*s, then *y* is in contact with *z*, or *y* is in contact with *w*, which is one of the *x*s, and *w* is in contact with *z* – and so on.[30]

All paradigm examples of ordinary composites *seem* to have inter-connected parts in this sense. Take a car, for example, the tyres are in contact with the wheels which are in contact with the hubs which are in contact with the axles which are in contact with the drive-shaft, and so on and so forth. For any two parts of the car, if they are not themselves in direct contact, they will be linked together via a chain of directly connected objects, all of which are themselves parts of the car. Objects with parts that are *not* in contact (or perhaps we should

[29] Van Inwagen (1990), 33. This is not van Inwagen's precise wording, but it expresses the same thesis.
[30] Van Inwagen (1990), 33

say *alleged* objects), by contrast, just don't seem to be the kinds of thing that the ordinary person on the street would recognise, as demonstrated by the supposition that there is an object composed of the Eiffel Tower and the Great Pyramid at Giza. *Scattered objects*, to use the philosophical jargon, just don't seem to tally up with the ontology of common sense.[31]

Despite this, however, CONTACT is a most unsatisfactory answer to the SCQ – particularly if it is an answer that is aiming to preserve the ontology of common sense. For it only takes a minimal amount of reflection to realise that CONTACT, if true, would have us countenance a whole host of exotic composites that are not recognised by common sense. Heap-tables, for instance, should be regarded as existing objects if CONTACT were true, for all their parts are in contact. Worse still, to use van Inwagen's popular example, if CONTACT were true, it would have the most undesirable consequence that whenever two people shook hands (or indeed came into contact of any sort), a new object would instantaneously come into existence, only to annihilate once the greeting ended and their hands parted.[32]

But not only would CONTACT have us countenance many exotic composites we may not want to, on the flip-side, it would also rule out the existence of certain ordinary composites that we are generally quite fond of. Consider some examples. A copy of the Oxford English Dictionary seems to be an example of a perfectly ordinary composite. Yet its parts are not in contact. It comes in a number of volumes, each disconnected (spatially) from the next.[33] Or take the landmass we call the United States of America. This too seems to be a perfectly ordinary composite. Yet it too is scattered, for Hawaii is separated from the mainland by a distance of over two thousand miles. If CONTACT were true, then, it would turn out that there are no existing copies of the Oxford English Dictionary, and even worse, that approximately 312 million Americans would have no country of residence. And these facts certainly don't sit well with our common sense intuitions – the very things that we are trying to preserve.

Perhaps even more damaging is the fact that modern physics tells us that at the sub-atomic level, the particles which are said to 'make up' atoms (i.e. leptons and quarks, etc.)

[31] The term 'scattered object' was, I believe, first introduced by Richard Cartwright (1975), who said: "let us say that a material object is *scattered* just in case the region of space it occupies is disconnected" (p.157).
[32] Van Inwagen (1990), 35
[33] This example is taken from Cartwright (1975), 157.

never actually come into direct contact at all. Indeed the vast majority of an average atom is empty space. So according to CONTACT, these sub-atomic particles *do not* compose atoms after all. And by extension, if, as we are often led to believe, all material things are made of atoms, then CONTACT would have us conclude there are no composite material objects at all. For if sub-atomic particles are not inter-connected, and ordinary composites are nothing but collections of sub-atomic particles, then it follows that the (alleged) parts of all ordinary composites are *not* in contact after all. Thus if CONTACT is true, there are no ordinary composites at all, which is a conclusion that could hardly be further removed from our common sense intuitions.

After rejecting CONTACT, van Inwagen went on to consider three other types of 'simple bonding answers' to the SCQ: FASTENING, COHESION, and FUSION.[34] All three answers are based on the principle that for any objects to compose a further object, they must be bonded together somehow or other (*fastened* together, bonded in such a way that the parts *cohere*, and *fused* together, respectively). Each proposed answer can be seen as a progression of the last in that it demands a stronger type of bond in order for composition to occur (FUSION, for instance, the most demanding of the three answers, states that for two objects to be fused together, they must be "melt[ed] into each other in a way that leaves no discernible boundary").[35] In light of the above comments, however, it shouldn't take long to realise that they are all entirely unsatisfactory as variants of CSR. For one thing, there are certain examples in which it just doesn't seem to matter how strongly bonded the objects in question may be. Take two human beings, for instance. You could bond them together as strongly as physically possible, but it still doesn't seem (according to common sense at least) that you would have brought a new object into existence.

But the real problem with these answers, and indeed with *any* type of bonding answer, is that they cannot account for scattered objects. If common sense recognises the existence of scattered objects (and it would appear, in light of the above examples, that it does), then bonding answers to the SCQ are all doomed to failure. For scattered objects, *by definition*, have parts that are not in contact with one another (let alone bonded together). With bonding answers, then, we have reached a dead-end.

[34] Van Inwagen (1990), 56-60
[35] Van Inwagen (1990), 59

### *§2.1.3 Some Other Simple Answers: Quickly Considered, Quickly Dismissed*

Perhaps we could find an answer to the SCQ not in terms of physical bonding, but in terms of *functionality?* That is to say, for any material objects, those objects compose a further material object *iff* they collectively perform some discernible function. The thought behind this answer is that ordinary composite objects *do* seem to perform functions: chairs function to support us and stop us falling to the ground; clothes function to keep us warm and dry; trees function to reproduce and maintain the continuation of their species; and so on and so forth. Furthermore, exotic composites seem *not* to have any discernible function. Take our example of the object composed of the Eiffel Tower and the Great Pyramid at Giza. If there were such an object, *what would it do; what would it be for?* It would appear to have no discernible function or purpose whatsoever, and so perhaps, on this view, that is why there is no such object. Let us call such an answer, *functionality:*

> FUNCTIONALITY: For any *x*s (where the *x*s are material objects) there is a further material object, *y*, that those *x*s compose iff the *x*s perform, collectively, a discernible function.

Functionality has some intuitive pull about it. Yet unfortunately, there are three reasons why it simply won't do.[36] Firstly, the notion of a 'discernible function' is particularly vague. For FUNCTIONALITY to constitute a satisfactory and precise answer to the SCQ, one would need to give a precise definition of what it is to perform a discernible function, and that may prove difficult. Secondly, just as with the simple bonding answers, there are numerous counter-examples available. That is, for many of the ordinary composites that are recognised by common sense, it is very difficult say what their function actually is. Take our heap of rice, for instance. We may want to allow that the grains of rice *do* compose a heap, but what function is the heap performing? It seems hard to say. The same goes for countless other ordinary composites too: works of art; mountains; the planet Mars; to name just but a few examples that come to mind.

[36] It is quite possible, of course, that there are a great deal more than three reasons why this answer fails. I will consider only three.

Finally, and most fatally, there is a further problem with FUNCTIONALITY. For it strikes me that the function various objects perform is purely a matter of stipulation or convention. Objects (or at least, the vast majority of objects) don't have functions *intrinsically,* but rather, we *impose* functions upon them. I can sit on a rock, and thus state that it functions as a seat, but it only functions as such *because* I have chosen to sit on it and impose this purpose upon it. Performing this function is not part of the rock's intrinsic nature. And indeed it is hard to see how performing *any* function could be part of its intrinsic nature. And the same surely goes for all inanimate objects. But this presents a serious problem for FUNCTIONALITY. For if material composition is determined by the collective functionality of material objects, and the functionality of objects is a matter of convention, then it follows that composition *too* must be a matter of convention (or at the very least, that compositional facts must *supervene* on conventional facts). But that contradicts one of the fundamental assumptions made at the outset: that what there is is *never* a matter of convention – what there is *just is.* This contradiction is enough, in my view, to render FUNCTIONALITY completely untenable.

Another feature of ordinary composites seems to be their *unity.* That is to say that the parts of the (alleged) composite objects that are usually recognised by common sense seem to form, in some sense or other, a unified whole. Furthermore, scattered objects *too* appear to be able to be unified in this sense, thus saving this type of answer from the counter-examples that ruined the simple bonding answers. We recognise the Oxford English Dictionary to be a single object, for instance, despite its volumes being spatially separated, because the volumes taken together, form a *unified whole.* Likewise, when I take the lid off my pen, I still recognise the two spatially separated parts to compose the one and the same pen, because the parts are unified in some sense or other. So perhaps here we have a potential answer to the SCQ:

> UNITY: for any *x*s (where the *x*s are material objects) there is a further material object, *y*, that those *x*s compose *iff* the *x*s display, collectively, a sufficient degree of unity.

As it stands, UNITY is far from being a satisfactory answer to the SCQ. For we do not yet have a clear and precise understanding of what it is for some *x*s to display a sufficient degree of unity. In fact, it seems very difficult to define 'unity' in this sense in a non-question-begging manner. What is it about ordinary composites that makes them appear *unified* that exotic composites lack? It can't be to do with the spatial relations that hold between parts. Because no matter how far I separate the volumes of the OED, it doesn't affect the perceived unity of the whole. Perhaps it may be to do with the actual *stuff* that objects are made of? That is to say, perhaps there needs to be a certain homogeneity of parts to effect a unity of the whole? But this can't be right. Some ordinary composites are made of a complete mish-mash of substances and materials, yet we still recognise them as wholes. Think of the vast array of metals, plastics, fabrics and liquids that make up an average car, for instance, these can hardly be said to be homogenous. Conversely, being made of the same stuff is often not enough for us to recognise a unified whole. For example, one could stack some bricks against the wall of a house, and even if the house is made of the very same batch of bricks, we don't consider the loose bricks a part of the house. So material homogeneity is clearly not the source of unity we recognise in ordinary composites.

It strikes me that what we are actually referring to when we speak of the 'unity of wholes', is a type of *conceptual unity.* That is, the reason we take a copy of the OED to be a unified whole (along with tables, chairs, people, and all other ordinary composites) is that we *conceive* of it as a unified whole. But if this is so, it immediately invites the *Euthyphro-style* question: do we conceive of an object as unified because it *is* unified, or is it unified because we conceive it so? If it is the latter, then it seems that we will end up facing the same problem that faced FUNCTIONALITY. That is, if being a unified whole is purely a matter of being *conceived* to be a unified whole, then it looks like being a unified whole comes down to mere convention. But that would again imply that what composite objects there are is a matter of mere convention, which is unacceptable. If we take the former option, however, we face a different problem. For if certain objects (i.e. ordinary composites) are genuinely (i.e. objectively) unified wholes, and it is this genuine, objective unity that makes us conceive of them as unified wholes, then we now face the question of what grounds their unity? In other words, *why are unified wholes unified*? But this question seems just as difficult as the SCQ! Indeed it looks as though by answering the SCQ with UNITY, we have simply pushed

the question back a level rather than answer it in any informative fashion. For we now need to fill in the right-hand side of the following biconditional:

UNITY: the *x*s form a unified whole iff _____.

The candidate explanations that could be used to complete this biconditional look to be just the same as those which could complete the original SCQ biconditional. Is it, for instance, the way in which the *x*s are bonded together that suffices for their unity? Or is it the fact that they collectively perform a discernible function? Thus it seems that we are back to where we started. It may well be the case that ordinary objects display some kind of unity, but the mere recognition of that fact will not help us answer the SCQ unless we can explain *why* they display that unity. And explaining *that* looks to be no easier than the initial challenge of answering the SCQ. UNITY, then, gets us nowhere.

## §2.2. The First Charge against Restrictivism: Arbitrariness

From the above discussion, I hope to have shown that CSR, in any form, is a very difficult position to defend. The reason for this is that there is seemingly no principled way in which to distinguish ordinary composites from exotic composites. We have seen that it can't be the way in which objects are *bonded* that determines whether they compose; nor can it be to do with the function that objects collectively perform, or whether they exhibit a particular degree of unity. Of course, these considerations do not exhaust the possible CSR answers to the SCQ, but I for one cannot think of any other plausible candidates, and neither am I aware of any others in the existing literature.[37] Moreover, it seems *likely*, albeit not certain, that any other purported CSR answers would be subject to similar difficulties to those considered above. In particular, it is hard to image a restriction on composition that preserves common sense intuitions, and that is *not* subject to at least one significant counter-example. The upshot of all this is that restrictivism is in big trouble. For if there is no *principled* way in which to distinguish those collections of objects that do compose from

[37] Excluding Markosian (1998b), of course, to whom we shall come in §2.5

those which don't, then it looks like any such distinction that one does hold on to must be entirely arbitrary. If you can't explain, for instance, why some limestone blocks compose a pyramid and some iron girders compose a tower, yet at the same time why the blocks and the girders taken together compose nothing, then your assertion that this is in fact the case looks entirely unjustified.

I am assuming that unjustified arbitrariness is a big problem for CSR. And that is because I am assuming that to draw a line *arbitrarily* between cases in which composition does occur and cases in which it does not is, philosophically speaking, intolerable. The underlying thought here is that if we are to assert that some objects do compose a further object and other objects don't, then we need a *reason* for so asserting. If we don't have such a reason, then our assertion reveals itself as being entirely arbitrary and completely lacking in justification. And in philosophy at least, arbitrary and unjustified assertions are surely to be avoided. Terry Horgan has clearly articulated this point, elucidating what he calls *the principle of the non-arbitrariness of composition (NAOC):* "Even though explanation must presumably bottom out somewhere, it is just not credible – or even intelligible – that it should bottom out with specific compositional facts which themselves are utterly unexplainable and which do not conform to any systematic and general principles".[38]

I think Horgan probably pushes it a little far to suggest that it is not even *intelligible* that compositional facts should be unexplained, for I think I could at least *conceive* of such a scenario. But he is surely right to suggest it is incredible. However, it should be noted that no *argument* is provided in support of NAOC; it is merely asserted. In fact, *arguing* for the truth of NAOC proves particularly difficult.[39] However, it is surely a most plausible principle. Indeed it may be fair to say it is overwhelmingly plausible when compared to its negation. In light of this, then, we can reasonably say of anyone wishing to defend CSR, that the burden of proof lies with them to provide a satisfactory reason for restricting composition where they do. And since it does not look as though any such reason is forthcoming, the prospects for CSR look bleak.

---

[38] Horgan (1993), 695

[39] In §2.5 I consider this in more detail, and provide an argument (*of sorts*) that supports NAOC. Not wanting to be misleading, I should point out that it is not a *direct* argument for NAOC (I'm not sure if such an argument is actually possible), but rather, an argument against the only extant restrictivist position that violates NAOC (Ned Markosian's).

## §2.3. Non-Common Sense Variants of Restrictivism

There are other variants of restrictivism in the extant literature that are *not* in line with the ontology of common sense. Most notable among these theories are those of van Inwagen himself, and Trenton Merricks.[40] Van Inwagen is quick to recognise the consequences of the arguments I have considered above, and he has no hesitation in following his philosophical convictions, no matter how unintuitive they may be, into denying the existence of the vast majority of ordinary composites. This is his famous "denial" mentioned in the introduction of his book. He does not, however, deny the existence of *all* composite objects; he allows that material objects can compose *if* their collective activities *constitute a life.* This, then, is van Inwagen's answer to the SCQ:

> (VI): for any *x*s (where those *x*s are material objects) there is a further material object, *y*, which those *x*s compose *iff* the activity of the *x*s constitutes a life (or if there is only one of the *x*s).[41]

The implications of this answer are that the only material objects in existence are material simples and living organisms. Trenton Merricks has adopted a similar position to this, although his reasons for adopting it are quite different from those of van Inwagen.[42] Thus although both views are restrictivist, they are both very much at odds with common sense. According to both van Inwagen and Merricks there are no rocks or planets, tables or chairs; there are, in fact, no inanimate composite objects whatsoever.

Despite the significance that these two views have in the literature on composition, and despite the impact that they have had on the overall debate, I will not be considering them in any more detail here. The reason for this will become clear in the next section, where I will set out an argument that undermines virtually all species of restrictivism, Merricks's and van Inwagen's included.[43] Because I think that this argument is strong enough to reject the restrictivist views of van Inwagen and Merricks, there seems little point in weighing up

[40] Van Inwagen (1990); Merricks (2001)
[41] Van Inwagen (1990), 82.
[42] Merricks (2001)
[43] The only restrictivist theory which may not succumb to the argument is Markosian's. I will consider his view later.

the various merits and demerits *of* those views. Instead, I shall press right on with the argument.

## §2.4. The Second Charge against Restrictivism: Vagueness

We have seen already that restrictivist answers to the SCQ are problematic. This is particularly so when trying to formulate an answer that preserves our common sense intuitions about what exists. Indeed it seems that there is *no* way of formulating such an answer without endorsing some kind of intolerably arbitrary line-drawing. Even van Inwagen's proposed answer, whilst striking a startling contrast with the ontology of common sense, cannot totally evade accusations of arbitrariness. However, whilst charges of arbitrariness may severely diminish the credibility of restrictivist answers to the SCQ, there is a stronger species of argument which warrants, in my view at least, their outright rejection: *arguments from vagueness.* In what follows I will show how this species of argument works, how it can take many different guises, and why it is so damaging to the restrictivist cause.

### *§2.4.1 Simple Sorites Arguments*

Arguments from vagueness directed against restrictivism usually begin by appeal to the ancient puzzle of the sorites.[44] Sorites paradoxes, or sorites series, are a form of chain-argument which draw utterly absurd conclusions from the most innocuous and reasonable of premises. The classic example of a sorites argument is based on the concept of a heap (in fact the word 'sorites' derives from the Greek '*soros*' which means 'heap'), and could be formulated as follows:

1. 1 grain of rice does not make a heap of rice.
2. 2 grains of rice do not make a heap of rice.
3. 3 grains of rice do not make a heap of rice.

[44] And it is indeed and *ancient* puzzle, believed to be originally formulated by Eubulides of Miletus, a contemporary of Aristotle.

4. In general, if *x* grains of rice do not make a heap, then *x + 1* grains do not make a heap.
5. Therefore, (by extension) 1 million grains of rice do not make a heap.

The argument can also be formulated in reverse, with perhaps even more startling conclusions:

1. 1 million grains of rice make a heap.
2. 999,999 grains of rice make a heap.
3. 999,998 grains of rice make a heap.
4. In general, if *x* grains of rice make a heap, then *x – 1* grains also make a heap.
5. Therefore, (by extension) zero grains of rice make a heap.

The sorites series are so puzzling because in each case, the individual premises strike us as perfectly reasonable, yet the conclusions they collectively entail strike us as patently absurd. For it seems perfectly reasonable – indeed, *obvious* – to think that the addition a single grain of rice would not be enough to turn a non-heap into a heap, or conversely that the subtraction of a single grain would not be enough to turn a heap into a non-heap. In other words, it seems implausible to suggest that there is an abrupt cut-off point between heaphood and non-heaphood. Yet if there is no abrupt cut-off point, then it seems we must say that either every member of the series either is a heap or is not a heap, which brings us back to our absurd conclusion.

These sorites-style arguments can be extended such that they have consequences not just for the existence of heaps, but for the existence of *any* macroscopic material objects whatsoever. Take any ordinary composite you want, and you will find that it is susceptible to the challenge of the sorites. An average cat, for instance, according to modern science at least, is composed of approximately $3 \times 10^{26}$ atoms.[45] That is a simply colossal number. But the astronomical number of microscopic parts that this hypothetical feline is purported to have, and indeed the truly miniscule size of each one, actually increases her susceptibility to the sorites argument. For one could imagine that if we had some incredibly high-precision tool, we could begin the laborious task of removing atoms from this unfortunate moggy, one by one. (The fact that no such tool actually exists is of no consequence, for it is the *possibility*

[45] This is, of course, *very* approximate. But it is, if anything, an underestimate.

of such a tool on which the force of the argument rests, and there is no obvious reason to think that such a tool is an impossibility). We would eventually reach a point where we had only a single atom left, and of course, a single atom is clearly not a cat. The problem for the restrictivist, then, is clear. If one wants to insist that composition is restricted, then one would need to draw a line somewhere in our feline-sorites series, such that the atoms on one side of the line compose a cat and those on the other side do not. But such an idea is surely preposterous! The very idea that the removal of a *single* atom -an object so incredibly small that to even conceive of it in isolation is almost impossible – could make the difference between there being a cat and there not being a cat is quite ludicrous. And the same goes for any material objects supposedly made up of many atoms: chairs, tables, trees, rocks, cats, and even people.

Peter Unger is perhaps the best known proponent of using sorites arguments in support of some startling conclusions. He did not use them to argue against restrictivism *per se,* but rather, on the basis of sorites arguments he claimed that "there are no ordinary things", and even more controversially, that "I do not exist".[46] But the fact that sorites arguments spell trouble specifically for the restrictivist is manifest. In addition, it should be noted that the more *extreme* answers to the SCQ (i.e. universalism and nihilism) are not obviously affected by these arguments, since neither of them is required to posit a cut-off point between composition and non-composition. That is because in one case (that of the nihilist) composition is said never to occur, and in the other (that of the Universalist) composition is said always to occur. The sorites arguments, then, represent a formidable foe for any friend of restrictivism.

## *§2.4.2 Introducing Ontic Vagueness*

One may be wondering at this point, what all of this has to do with vagueness. For as yet, there has been no mention of such a phenomenon. The strength of these *Ungerian* sorites arguments, as formulated above, lies purely in the fact that they force the restrictivist into the rather embarrassing (in fact, worse than embarrassing; untenable) position of having to specify exact cut-off points – *down to the very last atom* – between cases of composition and

[46] See Unger (1979a; 1979b; 1979c; 1980a; 1980b). Wheeler (1979) also presented similar arguments in support of similar conclusions.

cases of non-composition. But this conclusion rests on a significant assumption: that whether some *x*s compose or not *is always a determinate matter.* Perhaps, then, the restrictivist could reject this assumption, and maintain that (i) composition *is* restricted, (ii) there are *no* exact cut-off points between cases of composition and cases of non-composition, and (iii) it is sometimes simply *indeterminate* whether some *x*s compose a further object or not.

If we return to our previous example about the cat, then the thought here is that whilst the original 3 x $10^{26}$ atoms definitely *do* compose a cat, and whilst a single atom definitely *does not* compose a cat, there will be a phase somewhere during the atom removal process where it is genuinely indeterminate as to whether the remaining atoms compose a cat or not; a phase in which it is a *vague* matter as to whether there is a cat present or not. It must be made absolutely clear here that we are not talking about *linguistic* vagueness. That is to say, we are not suggesting that it is the extension of our sortal term 'cat' which is vague, and as such that there will be a phase during the atom removal process during which it will be unclear whether our sortal term 'cat' does or does not apply to the present collection of atoms. (Such a phase would inevitably arise, but it would be of little consequence to this debate). Instead, we are talking about something much more controversial: genuine vagueness *in the world.* What I mean by that is that even if we were able to tighten up our definition of the term 'cat' such that it had a perfectly exact and precise extension - down to the very last atom – we *still* wouldn't be know whether it applies to the present collection of atoms. The reason for that would be, that on this view, there is genuinely *no fact of the matter* as to whether the atoms compose a cat or not. Let us call this species of vagueness, if indeed it is even possible, *ontic vagueness.*

Ontic vagueness is problematic to say the least. For a start, it is quite difficult to actually conceive of what it would mean for the world to be vague. For some, the claim that there can genuinely be no fact of the matter as to whether some purported object exists or not, just makes no sense. For it has rightly been said that "the claim that *there is some thing* such that it is indeterminate whether *that thing* exists, is hard, if not impossible, to make sense of".[47] The thought here is that either we *are* quantifying over something or we are *not.*

[47] Barnes (2010b), 960. It should be noted that Barnes herself does not take ontic vagueness to be unintelligible – far from it. Barnes is one of the leading defenders of ontic vagueness. See Barnes (2005; 2009; 2010a) and Barnes & Williams (2009; 2011a; 2011b)

If we are, then it must determinately exist, and if we are not, then it does not. For this reason, the very idea of ontic vagueness is sometimes dismissed as an incoherent impossibility.[48]

But it is not just an instinctive distaste for ontic vagueness that makes it controversial. Some have presented arguments in an attempt to show that it just cannot occur. Ted Sider has provided just such an argument.[49] In a nutshell, Sider's claim is that *if* composition can be vague, then it would be also vague as to how many objects there were in the world. So suppose, for sake of argument, you had to count all the material objects in existence. That would include all composite objects, if indeed there were any such things. But suppose now we had some *x*s such that it was indeterminate as to whether they composed a further object. All the *x*s would be counted; they would go down on the list. But it would be indeterminate as to whether there was an *extra* object: the fusion of the *x*s in question. As Sider points out: "that would mean that some *numerical sentence* – a sentence asserting that there are exactly *n* concrete objects, for some finite *n* – would be indeterminate... [but] numerical sentences can never be indeterminate in truth value".[50] The reason for thinking that numerical sentences cannot be vague is that they contain only logical apparatus, and logical apparatus is supposed to be perfectly precise.

With all this in mind, then, sorites-style arguments put the restrictivist on the horns of very nasty dilemma. Either she accepts that there are exact cut-off points in sorites series, such that a single atom can make the difference between composition and non-composition, *or* she must accept that there is an abundance of ontic vagueness. Neither of these options looks good. Neither are necessarily impossible, but both are extremely controversial, so to endorse either one would represent a significant cost to the original restrictivist theory.

### *§2.4.3 The Final Charge against Restrictivism*

The preceding comments on vagueness are quite brief. And it would be far too quick if I were to simply conclude here that the 'arguments from vagueness' signify the death knell for any species of restrictivism. For it should be noted that in recent years much work has

[48] See Dummett (1975), 260, and Lewis (1986a), 212-13 for explicit views of this type. A similar view can even be found in Russell (1923)
[49] Sider (2001), 127
[50] Sider (2001), 127

gone into the defence of ontic vagueness. So much so, that it is now a fairly widely held view that ontic vagueness (or to use the more popular jargon 'metaphysical indeterminacy') is not only possible, but indeed most likely.[51] So one may think, then, that the restrictivist can once again breathe a sigh of relief, for it may seem that entailing ontic vagueness is not such a significant theoretical cost after all.

Unfortunately, however, I think this can be shown not to be the case. For even if one allows that ontic vagueness is a genuine possibility (and I for one would not want to rule it out), by positing ontic vagueness the restrictivist *does not,* I claim, overcome the destructive force of the sorites arguments. To sum up briefly, the reason is this. Consider once more our sorites series concerning the average cat. Indeed let's specify, for ease of reference, a particular cat: my own cat, Brian. The original sorites argument accused the restrictivist of having to demarcate an exact cut-off point in the atom-removal process, whereby there is a point at which the present atoms *do* compose Brian, yet if just a single atom were removed, those remaining would fail to compose Brian. It was supposed wildly implausible that such a cut-off point exists, and as such, restrictivism was supposed to have fallen by the wayside. But now, in light of the above comments, one can imagine the restrictivist defiantly responding: 'but there is no exact cut-off point! Rather, there is a phase of *indeterminacy;* a phase in which there is simply no fact of the matter as to whether the remaining atoms compose Brian or not. *Of course* you start off with a cat; and *of course* when you get down to a single atom you definitely have no cat, but that doesn't mean that there must be a precise cut-off point'.

I do not accept this response. I think it can be shown that even if one endorses some kind of ontic vagueness, and posits a phase in the sorites series in which it is indeterminate whether composition occurs or not, one still has to posit exact cut-off points which would be arbitrary and objectionable to the extreme. The difference between the restrictivist who does endorse ontic vagueness and the restrictivist who doesn't is actually minimal. For both have to posit exact cut-off points in a given sorites series, it is just that the latter is drawing the line between determinate composition and determinate non-composition, and the former is drawing the line between determinate composition and indeterminate composition. And

[51] In addition to the work of Barnes and Williams cited above, see Akiba (2004); Parsons & Woodruff (1995); Rosen & Smith (2004) for varying defences of ontic vagueness.

there seems no reason to suppose that one of these line-drawings is any more palatable than the other.

This type of objection, appealing to what is known as higher-order vagueness, has been noted elsewhere.[52] Perhaps the most lucid exposition of the problem is provided by Mark Heller, who imagines himself in the rather unusual situation of being atop Mt Everest, accompanied by God, playing 'the vagueness game'.[53] Below I have taken the liberty of paraphrasing Heller's example, altering the context somewhat, such that it is specific to the debate over material composition. So let me begin, like Heller, by setting a somewhat improbable scene. Imagine that you (the reader) are playing a game with God; the vagueness game. Also involved in the game (although not strictly participating in it) is Brian the cat. How the game works is this: God begins to remove atoms, one by one, from the unsuspecting Brian (God can do this – he is God, after all). After the removal of each atom, the game stops, and you ask Him: 'do the remaining atoms compose Brian?' God, playing the game in good spirit, will answer each question to the very best of his omniscient ability.

So the game begins with God's answering 'yes' to your question, for all the atoms present *do* compose Brian. And presumably, the game will continue for quite some time with God making this same assertion each time, for it will take the removal of a large number of atoms, one would assume, before they no longer compose a cat. But what should be clear, however, is that *at some point,* His assertion will *change*. Now the crucial point is that it does not matter what his assertion changes *to*, rather, it matters only that it change *at all.* For instance, it could be the case that after the removal of a particular atom, let's say it's the 1,348,946,627th, God answers 'no' to your question. He may even go on to assert: 'the remaining atoms *do not* compose Brian'. Such an assertion would imply that, contrary to our earlier misgivings, there *is* in fact a sharp cut-off point – down to the very last atom – between cat-hood and non-cat-hood. But alternatively, suppose that there was no such cut-off point, and instead there was a phase of indeterminacy between cat-hood and non-cat-hood. Even in that case there would be an exact point, after the removal of a specific atom, where God's assertion would change. It may change to, 'it is now indeterminate whether the

[52] See Horgan (1994), 173-4; Heller (1996); Hudson (2001). All three of these examples actually use their arguments to claim that ontic vagueness itself is impossible. I do not wish to argue for such a strong claim (I will remain neutral on that for the time being), but merely to claim that restrictivism is false.
[53] Heller (1996)

remaining atoms compose Brian', or perhaps, 'it is neither true nor false that the remaining atoms compose Brian', or perhaps, 'it is true to degree *x* that the remaining atoms compose Brian', or even, 'it is true according to only some precisifications of 'cat-hood' that the remaining atoms compose Brian'. Importantly, however, *whatever* God's changed assertion may be, it still has the consequence that the difference of a single atom can make a significant difference to the status of Brian's cat-hood. Now remember that the number of atoms that supposedly make up an average cat such as Brian is $3 \times 10^{26}$. Just try to consider, *if you can*, how unimaginably small each atom must be for that titanic number of them to be able to fit in the area of space occupied by an average cat. And now seriously try to entertain the idea that the removal of a single atom could make any difference whatsoever to the truth-value of the assertion 'these atoms compose a cat'. I defy anyone to maintain that such a situation could obtain. To suppose that a single, nugatory atom could make such a significant difference on the macroscopic scale would be to suppose nothing short of what Unger appropriately calls "a miracle of metaphysical illusion".[54] It is for this reason that I claim that, even if ontic vagueness is a genuine metaphysical possibility, material composition still cannot be restricted.

## §2.5. One Last Try: Brutal Composition

Throughout all the extant literature on this subject there has been only one (as far as I know) serious attempt to flesh out an answer to the SCQ that preserves our common sense intuitions about material composition. That attempt was made by Ned Markosian in his 1998 paper entitled 'Brutal Composition'.[55] Markosian labels his position 'Brutal', because it states that there is no *true and interesting* answer to the SCQ. Rather, whenever composition does or does not occur is, on this view, simply a brute fact. As will soon become clear, I do not think Brutal Composition is at all plausible, mainly because there are few good reasons, if any, to endorse it, whilst there are a number of good reasons to reject it. However, since it is the only genuine formulation of a common sense restrictivist position, it deserves to be

[54] Unger (1979b), 246
[55] Markosian (1998b). See also Markosian (2008).

taken seriously at the very least. This is particularly true in light of the reception that Markosian's thesis has had from the philosophical community. It is routinely dismissed as false, yet its falsity seems always to be assumed rather than argued for. Horgan and Potrč, for instance, say: "if one bunch of physical simples compose a genuine physical object, but another bunch of simples do not compose any genuine object, then there must be some reason *why*; it couldn't be that these two facts are themselves at the explanatory bedrock of being".[56] But alas no argument is given to support this conclusion, it is merely asserted. (I agree with Horgan & Potrč that their claim is immensely plausible, but we have already seen that when concerned with material composition, plausible claims can often lead up blind alleys). Indeed it seems that the vast majority of those who have contributed to the composition debate agree with Horgan & Potrč. Hud Hudson seems to be the only philosopher to have objected to Markosian's thesis by way of reasoned argument, and even his treatment is fairly brief.[57] James van Cleve, too, at least takes Brutal Composition seriously, as he admits it to be "respectable and not easily refutable".[58] Eventually, however, van Cleve's dismissal of Brutal Composition consists only of him finding it "pretty hard to swallow".[59]

I too think Brutal Composition is a respectable thesis, yet I also think it is refutable (although perhaps not *easily* refutable). But claiming it too be hard to swallow, or simply asserting it as false, clearly doesn't count as a refutation. In what follows, then, I will provide a refutation of Brutal Composition, *and argue for it.* I will set out exactly what Markosian's position is, why he endorses it, and ultimately, why it should be rejected.

## *§2.5.1 What is Brutal Composition?*

Markosian formulates Brutal Composition as follows:

(BC): There is no true, non-trivial, and finitely long answer to the SCQ.[60]

[56] Horgan & Potrč (2008), 19
[57] Hudson (2001), p.22-24
[58] Van Cleve (2008), 333
[59] Van Cleve (2008), 333
[60] Markosian (1998b), 213

Instead, whenever composition does (or doesn't) occur is just a brute fact, where a brute fact is defined as follows:

> *F* is a *brute fact* $=_{df}$ *F* is a fact, and it is not the case that *F* obtains in virtue of any other fact or facts.

BC is based on a number of assumptions, some of which are more reasonable than others. Firstly, Markosian assumes that any satisfactory answer to the SCQ should take the following form (or at least, should be able to be put in the following form):

> Necessarily, for any *x*s, there is an object composed of those *x*s iff _____.[61]

This seems a perfectly reasonable assumption to make, since any answer that was of this form would give us exactly what we want: the necessary and sufficient conditions under which composition occurs.

Secondly, Markosian assumes that any satisfactory answer to the SCQ must be non-trivial. This too seems perfectly acceptable, indeed it seems most advisable, since a trivial answer to the SCQ would not tell us anything of any interest (that is, after all, why we would call it trivial). For example, one could give the following answer to the SCQ:

> (Triv): Necessarily, for any *x*s, there is an object composed of those *x*s iff the *x*s compose.

Triv is clearly true, and since it is of the required biconditional form, it does constitute an answer (of sorts) to the SCQ. It is clearly trivial however, as it tells us nothing of any interest, and illuminates no principles pertaining to composition.[62] So Markosian is right, in my view, to assume that any satisfactory answer to the SCQ will be non-trivial.

---

[61] Markosian (1998b), 214

[62] The *reason* it provides no illumination is that it contravenes van Inwagen's stipulation that the right-hand side of the biconditional should not contain any mereological terms. Essentially, then, Triv tries to explain 'composition' in terms of 'composition'.

Markosian's final assumption, however, is not so obviously true. Firstly, he rightly points out that even if compositional facts were brute we could, if we had the relevant knowledge, fill in the right flank of the biconditional by simply listing each and every individual case of composition. However, he then assumes that such a list would be infinitely long: "Such a sentence would have to be long indeed. In fact, even if the number of cases of composition in our world is finite, it seems clear that any sentence that expressed a necessary truth about composition in such an enumerative fashion would have to be infinitely long".[63] It is not at all clear to me why this should be the case. After all, suppose for sake of argument that compositional facts were brute. In that case it would be perfectly possible that there was only a single case of composition, for instance, the set of simples {S1, S2, ... Sn} (call it set A) composing object O. In that case, it seems we could provide an instance of the above biconditional schema with not only a finite right-hand side, but a positively short right-hand side:

> Necessarily, for any *x*s, there is an object composed of those *x*s iff each of the *x*s is identical to a distinct member of set A, and there are the same number of *xs* as there are members of Set A.

With this in mind, then, it is somewhat strange that Markosian assumes that any such enumeration must be infinitely long. However, this is probably of little consequence, since I am in agreement with Markosian that any such enumeration that *was* infinitely long should *not* be considered an informative answer to the SCQ. Whether or not a body of brute compositional facts could or could not be enumerated in a finite list may be up for debate, but either way, Markosian's position is relatively unaffected. For if they *couldn't* be listed finitely then BC would be correct, and if they *could* be finitely listed then although BC would be incorrect, Markosian's position would be even stronger, since the finite list *would* constitute an informative, non-trivial answer to the SCQ: the very thing we have been looking for since the start. So despite the fact that I disagree with Markosian over this particular point, it does nothing to undermine his central claim that compositional facts are brute.

[63] Markosian (1998b), 5

## *§2.5.2 Why Should One Believe in Brutal Composition?*

Markosian gives only one argument to support BC: the argument from elimination. Essentially, Markosian's line of thought is that all other candidate answers to the SCQ are implausible to the point of being untenable, resulting in the conclusion that composition must be brute. He formulates the argument as follows:

(1) [Compositional] nihilism is false

(2) [Compositional] universalism is false

(3) There is no true moderate[64] answer to the SCQ

(4) If (1) – (3), then BC is true

(5) Therefore, BC is true[65]

The argument from elimination is deductively valid. (Provided, that is, that one agrees with Markosian that an infinitely long disjunction enumerating each and every case of composition *does not* constitute an informative answer to the SCQ. As stated earlier, I see no reason to disagree over this point, so I am happy to concede that the argument is valid). Whether it is convincing or not, however, is a very different matter. This will depend on how plausible one takes the first three premises to be.

I will not take issue with premise (3) here. It should hopefully have been demonstrated in the preceding sections that restrictivist answers of all stripes (excluding BC of course) are subject to so many difficulties that they should be rejected as false. Markosian points to similar difficulties in his defence of (3), placing a particular emphasis on the argument from vagueness. Since I agree that these difficulties represent fatal flaws in restrictivist theories of composition, I will also agree that (3) is true. (Indeed it is a strength of BC that it is unaffected by the vagueness and sorites arguments considered in the last section. It remains unaffected because the BCer endorses exactly what the sorites arguments suggest: that there is a sharp cut-off point between composition and non-composition. Sharp cut-off points have so far been the thorn in the side of restrictivist s, because there is seemingly no non-arbitrary way of determining where they should lie. But the proponent of

[64] According to Markosian, a moderate answer is any restrictivist answer that can be expressed in a finitely long sentence, or sentences.
[65] Markosian (1998b), 27

BC is unruffled by such a worry. As Markosian notes, "she can just shrug and say 'there is no reason. It is a brute fact'").[66]

Markosian's defence of (1) and (2) are much less convincing, however. Indeed he does not provide any *arguments* to show the falsity of nihilism or universalism at all. Rather, they are rejected on purely intuitive grounds. For example, here is Markosian's rejection of nihilism:

"according to my intuitions there are far more composite objects in the world than nihilism allows. This seems to me to be a fatal objection to nihilism, and I conclude, on the basis of this objection, that nihilism is not the correct answer to the SCQ".[67]

Universalism, too, is dispensed with in a similarly confident, and cursory, manner:

"there is what seems to me to be a fatal objection to universalism: universalism entails that there are far more composite objects than common sense intuitions allow. [...] On the basis of this objection I reject universalism".[68]

What should have been made clear in the preceding sections is that I place little importance on common sense intuitions when trying to answer the SCQ. Or perhaps to be a little more cautious, I should say that it seems that the more one reflects on composition, the more it seems that our intuitions about it are mistaken, and so perhaps they should not be given all that much credence. To be sure, our intuitions may give us a good place to start, but it only takes a little serious reflection to realise that they bring with them some serious metaphysical problems. For this reason, then, one will not be surprised to discover that I am unconvinced by Markosian's rejection of both nihilism and universalism. Markosian, however, is unlikely to be concerned by this. What it boils down to is a base disagreement between us over how much of a part common sense should play in our metaphysical theorising. Whilst I am of the view it should play only a cameo role at most, Markosian would no doubt take the Moorean line that our common sense intuitions about what things exist are so strong as to overrule any philosophical arguments that may contradict them. I don't want to get drawn in to a discussion of the merits and demerits of Mooreanism here, so I will say only that I recognise it a serious philosophical stance to adopt, albeit one I

[66] Markosian (1998b), 37. It will be argued later that the acceptance of sharp cut-off points *does* present a problem for the BCer, in spite of Markosian's claims to the contrary.
[67] Markosian (1998b), 14
[68] Markosian (1998b), 22-23

generally disagree with, and as such, on this point at least, I am willing to agree to disagree with Markosian. However, herein lies a problem: if it is common sense that is driving Markosian's approval of BC, then upon closer inspection, his position begins to look somewhat confused. For there are a number of reasons to think that BC is, in fact, in great conflict with common sense. I will list these reasons below.

### *§2.5.3 Three Reasons to Disbelieve Brutal Composition*

**REASON 1:** Think back to when I set out the vagueness game, involving you, God, and Brian the cat. Now suppose that BC is true and, as Markosian suggests, that whatever the brute compositional facts are, they preserve our common sense intuitions about composition. In that case, it would seem that at the very beginning of the vagueness game, when *all* Brian's atoms are untouched, it must be a fact that those atoms compose – they compose Brian. Moreover, near the very end of the game, when there are, say, only two atoms remaining, then if common sense is anything to go by at least, it must be a fact that those two atoms *do not* compose a cat. The conjunction of these two facts entails that there must be a sharp cut-off point, somewhere during the game, at which the removal of a single atom makes it the case that Brian ceases to be. But Markosian would be quite happy to accept this. Recall, he thinks that the defender of BC can just shrug her shoulders about sharp cut-off points, and say "it's a brute fact".

But God, of course, would know the precise location of this cut-off point, due to his omniscience. So let's suppose he removed all the atoms he could from Brian until he reached that cut-off point, i.e. the stage at which, if just one more atom were removed, the remaining atoms would no longer compose anything. Now I presume that the collection of atoms in front of you at this stage would look remarkably like a cat. After all, they do compose a cat, and the compositional facts are, *ex hypothesi*, supposed to preserve our common sense intuitions about what composite object there are. But now suppose that God removed one more atom. I will presume that there would be absolutely no perceptible difference whatsoever between the collection of atoms in front of you now, and the previous collection that included the single atom that God has just removed. If he had removed the atom while your back was turned, for instance, you would have had no way of recognising that

anything had changed once you returned your gaze upon Brian. But despite this entirely negligible, completely imperceptible difference, you are now being told, by God, that there is no longer a cat in front of you; you are being told that Brian has ceased to be. But this does not sit well with our intuitions! Indeed, such a claim would fly in the very face of common sense. Any normal observer would surely resist this claim. "Of course Brian has not ceased to be", one can reasonably envisage them proclaim, "he is still right there in front of me"! So the fact that BC entails that there are sharp cut-off points between cases of composition and cases of non-composition is the first reason to suspect that it does not do what it says on the tin; it does not preserve our common sense intuitions.

**REASON 2:** Common sense may well point towards the fact that composition is restricted, but, as James van Cleve has rightly noted, it surely also points towards the fact that there is a *reason* why it is restricted.[69] According to common sense, for instance, there is a collection of iron girders on the *Champs du Mars* in Paris that *do* compose; they compose the Eiffel Tower. But there is surely a reason why they compose, or at least, that is what our common sense suggests, i.e. it is because the girders are fixed together in a particular way, or that they form a unified, purpose-built whole, or whatever. And it is for a similar reason that common sense suggests that a collection of limestone blocks in Giza compose a great pyramid. Likewise, it is the conspicuous *lack* of such a reason that the iron girders in Paris and the limestone blocks in Giza do *not* compose an object when taken together, or so our common sense would have it. So it would surely run completely counter to common sense to suggest that we are in fact mistaken about this, and there is in fact, no explanation whatsoever as to why these compositional facts obtain. But that is, of course, what the defender of BC has to maintain.

**REASON 3:** The final reason why BC clashes with intuition is that it would, if true, leave us completely in the dark as to what the compositional facts actually are. How would one go about establishing what composite objects exist if there was no explanatory principle underlying composition in general? Markosian, I think, would suggest that the answer is in fact easy. He presents his thesis of brutal composition as one which is "consistent with

[69] Van Cleve (2008), 333

standard, pre-philosophical intuitions about the universe's composite objects".[70] So I imagine he would say that we have a very simple way of establishing what composite objects exist: if intuition says a certain composite object exists, then we should accept that it *really does* exist. What could be simpler than that? But, unfortunately, this just won't work. As I have previously made clear, our intuitions about composite objects are nowhere near fine-tuned enough to detect the fine divisions between composites and non-composites that BC entails there must be. Consider Brian, in the vagueness game, right on the cut-off point of composition. Call Brian on the right side of the cut-off point (i.e. an existent composite object) 'Real Brian', and call the atoms on the wrong side of the cut-off point (i.e. one atom short of composing a cat) 'Faux Brian'. Common sense and/or intuition would have absolutely no way of distinguishing Real Brian from Faux Brian. Either *both* would seem like a composite object (a cat) or *neither* would. It is beyond any stretch of the imagination to suppose that we could distinguish, intuitively, between the two. Thus in either case, our intuitions about composition would be mistaken. Either we would take Faux Brian to be a cat (which he is not), or we would take Real Brian to *not* be a cat (which he is). But, of course, if our intuitions can be mistaken in this case, then what reason do we have to suppose that they are right in other cases? What reason do we have to suppose that they are right in *any* cases? I would suggest that there is no such reason, other than pure optimism on the part of the Brutal Composition Theorist.

It is for these three reasons that the motivation for BC disappears. We are told that BC is the only candidate answer to the SCQ that preserves our pre-theoretical intuitions about composition. Moreover, it is *because* of this very fact that Markosian suggests that we should believe that BC is true. But as I have just shown, BC does *not* preserve our intuitions about composition. And what is more, it actually *violates* what is a strong and central intuition about composition that the other candidate answers to the SCQ successfully preserve, namely, that whatever the compositional facts are, there must be a *reason* as to why they obtain. So in light of this, then, there is simply no motivation for believing BC in the first place.

[70] Markosian (1998b), 211

## *§2.5.4 A Brief Digression about Intuition*

Before concluding, I just want to say one more thing about the intuitions that are the driving force behind Markosian's thesis. Markosian's starting point is that it is just *obviously apparent* that ordinary composites, like tables, chairs, people, stars, etc., exist, and that exotic composites, like heap-tables, do not exist. However, I simply disagree that this intuition is as strong as he claims, and I believe that this can be shown with just a minimal amount of reflection. Sure, if you ask the average person on the street, "do tables exist?" they will unthinkingly answer "yes" (if they grace your question with an answer at all!), and likewise, if you ask "is there an object composed of your dining table and the moon?" they will, more than likely, say "no". But it would normally take only a few more simple questions to make their conviction in these answers wane, or so I claim. The best way I can demonstrate this is by recounting a conversation I have had a number of times with first-year philosophy undergraduates whom, incidentally, constitute perhaps the best representation of common sense intuitions. I say this because they are willing to take philosophical questions seriously (usually), yet at the same time their philosophical views are fairly undeveloped. The conversation tends to go a little like this:

Me: *Do crowds of people exist?*

Them: *Yes*

Me: *Okay, well suppose there is a crowd of 100 people in front of you, and suppose for sake of argument that people are mereologically simple. How many objects are in front of you, 100 or 101?*

Them: *What do you mean?*

Me: *I repeat: how many objects are in front of you? 100 (i.e. just the people), or 101 (i.e. the people and the crowd)?*

Them: *Aah, I see. Well now you come to mention it, I'm not so sure.*

Me: *Now look at that table in front of you. Now suppose, for sake of argument, it is made of 100 mereological simples. Now tell me how many objects are in front of you: 100 (i.e. just the simples) or 101 (i.e. the simples and the table)?*

The conversation usually goes one of three ways at this point. Either, they will plump for there being 101 objects, in order to preserve the table, or they will tentatively opt for there

being only 100 objects, unwittingly starting along the road to nihilism. Or sometimes they maintain that whilst there are definitely only 100 objects, the table still exists because the table *just is* the simples taken together, thus unwittingly endorsing that controversial thesis that composition is identity. It strikes me that at this early stage of investigation, each of these three answers has something going for it, and what's more, each of the three answers is perfectly understandable for them to settle on. But what this surely shows is that our intuitions about what exists are *not* as strong as they may initially seem, and certainly not as strong as Markosian takes them to be. Whilst one's initial reaction may be a complete conviction in the claim that tables exist, under only gentle questioning that conviction will often dissipate significantly. The intuition that tables exist is only strong when taken *completely uncritically*, and it seems that this is the way Markosian takes it when forming his views about composition.

As a final point on this, it seems that the conclusions of nihilism and universalism (in fact, of nihilism in particular) only seem so counter-intuitive because they too are taken completely uncritically. For instance, nihilists are regularly reminded (often in a derogatory tone) that their stance on composition entails that there are no tables and chairs. This is quite true. But the conclusion is only *seriously* counter-intuitive when taken uncritically. For the claim 'there are no tables and chairs', when taken uncritically, would have us believe that we have no surfaces to put things on and nowhere to sit down. But of course it doesn't mean that! If it did then we would be right to reject it as absurd. As we well know, the nihilist can sit down just as comfortably as anyone else; she merely maintains that what she is sitting on is not a single composite object (a chair), but rather a plurality of little objects (i.e. simples arranged chair-wise). When this is understood, the conflict between nihilism and common sense starts to lose its bite and look more of a pseudo-conflict. The conflict only looks damaging when our intuitions are taken completely uncritically. In light of this, then, I would suggest that it may be somewhat misguided to build a philosophical theory of material composition upon intuitive foundations that appear to have been subjected to little or no reflective scrutiny.

## §2.6. Concluding Remarks

I hope to have shown that restrictivism, in any of its guises, is false. The vast majority of restrictivist theories involve theoretically intolerable arbitrariness, since they fail to provide a justified and principled reason to restrict composition in the way they do. But even those restrictivist theories which claim to have such a principled reason are subject to the destructive force of the sorites. If composition is restricted, then there must be sharp cut-off points between cases of composition and cases of non-composition, and this is regardless of whether or not one endorses the possibility of ontic vagueness. Such cut-off points are implausible and objectionable to the extreme, and thus any theories which entail that there are such things should be rejected. Finally, I have shown that Brutal Composition, the one version of restrictivism that is markedly different from the rest, should also be rejected. It too endorses the existence of implausible and objectionable sharp cut-off points. Moreover, despite Markosian's claims to the contrary, there is no good reason to believe in BC in the first place, since the claim that it preserves our common sense intuitions about composition is entirely false. It is for these reasons that I conclude that material composition is *not* restricted.

# §3. Universalism

If restrictivist theories of composition are untenable, and in light of comments made in the preceding chapter I will assume that they are, then one is left only with what have been called the "extreme answers" to the SCQ.[71] That is, one must say either that composition *always* occurs (i.e. endorse universalism) or say that composition *never* occurs (i.e. endorse nihilism). Of these two options it is the former that has proved by far the most popular among contemporary philosophers; indeed it would probably fair to say it is the default view.[72] Universalism, recall, is defined as follows:

> **Universalism:** for any two or more *x*s, where the *x*s are material objects, there is always a further material object which those *x*s compose.

I should point out here that universalism is quite distinct from, yet often conflated with, the somewhat stronger thesis of Unrestricted Mereological Composition (UMC). Universalism is a thesis concerned only with *material* objects, whereas UMC states that for any objects whatsoever (material or not) there is a further object that they compose. Thus UMC entails the existence of many strange and unusual objects which universalism does not. Armstrong, for instance, a famous advocate of UMC, recognises that his view commits him to objects such as that composed of "the number 42 and the Murrumbidgee River".[73] It may be thought (indeed, until recently I myself thought) that UMC *entails* universalism, since the set of all material composites is just a subset of the set of all composites recognised by UMC, but this is not actually the case.[74] This is because although UMC says that for any objects there is always a further object they compose, it does not - or *need* not - insist that the object which they compose is a *material* object. The defender of UMC can remain neutral on the *nature* of composite fusions. This is, however, a technicality which is probably safe to ignore. For it

71 Van Inwagen (1990), 72
72 Defenders of universalism include, Armstrong (1989, 92; 1997, 13), Heller (1990, 49-51), Hudson (2000; 2001, 107-112), Leonard & Goodman (1940), Lesniewski (1916), Lewis (1986a, 211-213; 1991, 7), McGrath (1998), Rea (1998), Sider (2001, 121-32), and many others.
73 Armstrong (1997), 13
74 This point was brought to my attention by reading Effingham (2007), 57

would be a most unusual thesis which states that the fusion of two material objects was itself not a material object – I certainly know of no one who defends such a thought. So whilst some of the philosophers I will refer to as defenders of universalism are actually defenders of UMC, one should not worry that I am misrepresenting their positions by doing so. For despite the fact that UMC does not entail universalism, these particular philosophers *do* endorse universalism. As a final point on this, it should be noted that universalism does not entail UMC, thus while any observations I make about universalism *may* carry over to UMC, it should not be assumed that they will.

To the reader who is unfamiliar with the debate on composition it may come as something of a surprise to learn that the view of the informed majority is that material composition is unrestricted. This is because the truth of universalism implies the existence of a plethora of exotic composites. There would be a veritable abundance of heap-tables, for instance. And it would turn out that there is, after all, an object composed of the Eiffel Tower and the Great Pyramid at Giza. Universalism is entirely indiscriminate; it matters not how disparate or incontiguous any two objects may be, according to universalism they will compose. Despite this, however, universalism remains very much the majority view.

In what follows I will build a case against universalism. The case will consist of two parts. In the first part, I shall consider some of the key arguments that are often given in support of universalism, and expose them to be flawed and unconvincing. In the second part I shall present a number of arguments against universalism. I don't think that any of these should be considered as knock-down arguments when taken on their own, but when taken together, they constitute a battery of reasons to reject universalism. Coupled with the fact that there is little motivation to believe universalism in the first place, I will conclude that it is false.

# §3.1. Arguments *for* Universalism

## *§3.1.1 Initial (Naive) Advantages*

Universalism has some clear advantages over restrictivism. Firstly, unlike most restrictivist theories, universalism is completely unaffected by sorites arguments and arguments from vagueness. The reason for this is that since composition occurs in *all* cases, there is no need to posit any sharp cut-off points, or phases of indeterminacy, between cases of composition and cases of non-composition. According to universalism, there are no cases of non-composition! Secondly, and also unlike most restrictivist theories, universalism cannot be accused of unjustified arbitrariness, for it is perfectly consistent across the board. Universalism is uniform and indiscriminate: material objects always compose! Thirdly, and perhaps following on from the last point, universalism is often touted as being simple and elegant.[75] Since simplicity and elegance are widely regarded as theoretical virtues, this is supposed to count in universalism's favour.

I am perfectly willing to concede all these points in favour of universalism. Indeed it seems quite clear to me that universalism is a much better answer to the SCQ than restrictivism (in any of its many guises). However, by now we have already eliminated restrictivism, so touting advantages a theory has over restrictivism is not going be of great consequence at this stage. As stated above, with restrictivism out of the picture, we are left with a straight battle between universalism and nihilism. So for universalism to win the day, its proponents need to provide advantages it wields over nihilism, not over restrictivism. Yet, the three initial advantages mentioned above can also be forwarded by the nihilist. For nihilism is also perfectly unaffected by sorites or vagueness arguments. It can never be vague as to whether composition occurs, for composition never occurs. And neither is there any need to posit sharp cut-off points between cases of composition and cases of non-composition, for according to nihilism, there are *no* cases of composition. Neither can nihilism be accused of unjustified arbitrariness, since it too is consistent across the board: composition never occurs. And because of this nihilism is just as simple and elegant a theory

75 For example, see Markosian (2008), 345

as universalism.[76] So although the universalist certainly can lay claim to the three advantages set out above, since the nihilist can do the same, these factors will provide no help when it comes to deciding between the two.

### *§3.1.2 The Argument from Elimination*

There is an argument for universalism which seems to hold considerable sway with a great many philosophers, yet it is an argument which is rarely explicitly stated (let alone defended!) but usually just implied, or even assumed. It is an argument from elimination. The argument consists of the conjunction of two claims: firstly, that composition is not restricted (usually based on the types of consideration covered in the last chapter), and secondly, that composition clearly occurs in some cases (e.g. I exist, and I am composed of parts). The conjunction of these two claims is taken to entail the truth of universalism. Let's call this the *Argument from Elimination (*AFE*),* and formulate it as follows:

1. Composition is not restricted
2. Therefore, composition must either always occur or never occur
3. Composition definitely occurs in some cases
4. Therefore, composition must occur in all cases
5. Therefore, composition is unrestricted

David Lewis has endorsed this type of argument in favour of universalism. He says: "no restrictions on composition can serve the intuitions that motivate it. So restriction would be gratuitous. Composition is unrestricted".[77] Notice how quickly Lewis jumps from the premise that composition is not restricted to the conclusion that it must be *un*restricted. *It is surely far too quick!* Lewis jumps straight from premise (1) in AFE to the conclusion (5). This is, of course, an invalid inference. The extra premises (and in particular, premise 3) are *needed* in order for one to legitimately infer that composition is unrestricted. Lewis, I would suggest, simply assumes that these extra premises are obviously true – indeed this must

76 In fact it could be argued that nihilism is *more* simple and elegant than universalism, for nihilism only posits one type of material object (simples), whereas universalism posits two (simples and composites), but not much will hang on this point.
77 Lewis (1986a), 213

have been *so* obvious to him that it wasn't worth mentioning. (The other alternative, of course, is that Lewis *did* in fact make an invalid inference, but I am assuming that this is unlikely).

Ted Sider, too, has presented an argument similar to AFE.[78] I am referring here to his well-known 'Argument from Vagueness'; the very same argument that was considered in the last chapter. Sider's argument is directly influenced by Lewis's comments, yet it is formulated in more detail. The argument essentially states that composition can't be restricted (because of the concerns about vagueness and sharp cut-off points that were considered in the last chapter), therefore composition must be unrestricted. Unlike Lewis, however, Sider does recognise that his argument needs an extra premise in order to be valid. More precisely, Sider recognises that if nihilism were true, or even just possible, his argument would not go through. Sider's solution is to simply reject nihilism on independent grounds, and thus conclude that composition is unrestricted.

I will later argue that Sider's rejection of nihilism is unfounded, and thus by extension, that his endorsement of universalism is too. But that is for later. Now it just remains to be clarified that *without* a valid reason to reject nihilism, the argument from elimination does not support universalism. AFE, really, is just a reiteration that composition is not restricted, and that as a result, either universalism *or* nihilism must be true. Thus in an attempt to answer the SCQ it gets us no further than we were already. To decide between universalism and nihilism, we will need new, *independent*, arguments.

### *§3.1.3 The Argument from CAI*

It has been suggested that composite objects may be *identical* to their parts taken together. The rough idea here is that if you've got the parts, you've got the whole. The whole, although existent in its own right, is not anything additional to the parts. To use the often used parlance, the whole is *nothing over and above its parts*. This thesis is known as *Composition as Identity* (CAI).[79] It has also been claimed that CAI entails universalism, thus if

78 Sider (2001), 120-132

79 Defenders of this rough view include Lewis (1991), Baxter (1988a; 1988b), Sider (2007), Armstrong (1997)

the former is true, then so is the latter.[80] As such, any arguments in favour of CAI can be seen as being indirect arguments for universalism. In what follows I aim to show that CAI does not, in fact, entail universalism, and is perfectly compatible with other stances on material composition. More importantly, however, I think it can also be shown that there is no good reason to believe in CAI in the first place (in fact, I think there is good reason to reject it outright). But before I do this, it would be worthwhile to get a little clearer of what the thesis of CAI actually claims, *in detail.*

First and foremost, let me clarify what CAI is *not*. CAI is not the thesis that composite objects are identical to the sum of their parts, or the aggregate of their parts. For this would be to claim only that composites, if they exist, are self-identical, which would be a fairly trivial claim. Rather, the defender of CAI claims that for any composite object, *O*, composed of parts, *the x*s, *O just is* the *x*s. CAI, then, stretches the standard conception of identity. Identity, ordinarily conceived, is a one-one relation, i.e. a relation that holds between one single object and another single object (e.g. Phosphorus = Hesperus, or Elton John = Reginald Dwight). But according to CAI, identity can also be a many-one relation, i.e. a relation linking a plurality of objects (some parts) to a single object (a whole). Because of this, CAI is quite hard to express. One is tempted to say something like, 'according to CAI, for any composite, *O*, composed of parts, the *x*s, then the *x*s are identical to *O*'. But this could quite easily be misinterpreted; for taken at face value, it seems to suggest that *each* of the *x*s is identical to *O*, and that is obviously not what we are trying to say. In fact, Ross Cameron has suggested that discussing CAI *forces* one to be either ungrammatical or misleading.[81] One risks being misleading if one expresses it in the way just done, but in order to express it more precisely, one may need to sacrifice one's grammar and say something like, 'the *x*s is identical to *O*'. Not being completely satisfied with either of these terminological choices, I will, in what follows, use the slightly more cumbersome, yet less grammatically objectionable phrase, 'the *x*s taken together are identical to *O*'. For brevity, I will denote this symbolically by: 'the *x*s = *O*'.

CAI certainly has an appeal of sorts. For when we think of any given composite object, it seems that, materially speaking, the totality of its parts constitute the *entire* whole.

80 E.g. Merricks (2005); Sider (2007)
81 Cameron, (2012), 531n

There is no extra material in the whole - no extra stuff - that is not accounted for by all the parts taken together. (If you recall the example in the last chapter about the crowd of 100 people, it would seem that even for one who believes in the existence of crowds, one may hesitate to say that the crowd is an *extra* object, such that there are 101 objects present instead of just 100. The intuitive thought is: if you've got the people, you've got the crowd). This is the driving intuition behind CAI, and has perhaps been best expressed by Donald Baxter.[82] Consider some of his examples:

> Someone with a six-pack of orange juice may reflect on how many items he has when entering a 'six items or less' line in a grocery store. He may think he has one item, or six, but he would be astonished if the cashier said 'Go to the next line please, you have seven items'. We do not ordinarily think of a six-pack as seven items, six parts plus one whole.[83]

To reinforce the point, another example:

> Suppose a man owned some land which he divided into six parcels. Overcome with enthusiasm for the non-identity view [of composition] he might try to perpetrate the following scam. He sells off the six parcels while retaining ownership of the whole. That way he gets some cash while hanging on to his land. Suppose the six buyers of the parcels argue that they jointly own the whole and that the original owner now owns nothing. Their argument seems right. But it suggests that the whole was not a seventh thing.[84]

There is little doubt that these examples are powerful. After all, one certainly *would* be astonished if refused entry to the 'six items or less' queue, and one would rightly be aggrieved if a land-owner continued to lay claim to a plot of land you had just bought from him in the circumstances set out above. But there is equally little doubt that these examples do not constitute arguments to show that composition is identity. For they could, I think, just as forcefully be used to support nihilism. For the examples seem to demonstrate that if you have the parts, you have *everything there is to have*. If you have the six bottles of juice, you have *all* the juice. There is no six-pack over and above the bottles because there is no six-pack at all – there is just the six bottles! Of course, the examples don't show *conclusively*

82 Baxter (1988a)
83 Baxter (1988a), 579
84 Baxter (1988a), 579

either way. What they do, however, is highlight the fact that, perhaps, the relationship between parts and whole is not quite as straightforward as we may ordinarily assume.

Many have thought that the truth of CAI entails the truth of universalism.[85] The general idea is that *if* composition is identity, then the fusion of any objects just *is* those objects taken together. So for any objects whatsoever, you automatically get their fusion, because their fusion *just is* those objects. Trenton Merricks forwards just such a proposal, making the seemingly plausible claim that "it seems nonsensical to deny the existence of something that would, if it existed, be (identical with) things whose existence one already affirms".[86] But that is precisely what someone would be doing if they endorsed CAI but did not endorse universalism, or so the argument goes. Therefore, we are led to conclude that if CAI is true, universalism must be true also. Consider a six-pack of orange juice, for example. First, suppose that you accept, unremittingly, the existence of the six individual bottles of juice. Now according to CAI, the six-pack (the whole) *just is* the six bottles taken together, nothing more, and nothing less. So given the fact that you accept the existence of the six bottles, you *already* accept the existence of the six-pack. And the same goes for any collection of objects you can think of. It's as simple as that; CAI entails universalism.

Whilst this argument may appear to be quite powerful, there are reasons to believe that its power is only superficial. In a recent paper, Ross Cameron argues that the argument is not valid. Cameron's central point is that CAI is a thesis about *the nature* of composition (i.e. it tells us *what composition is* - identity), but it doesn't necessarily tell us *when* composition does and does not occur. For CAI tells us that when there *is* a composite object, that object is identical to its parts taken together. Furthermore, it tells us that when some objects are, taken together, identical to some single object, then they compose that object. Crucially, however, it doesn't tell us *when* some objects are identical to a single object and when they are not. As Cameron says: "[CAI] does not tell us whether, given some *x*s, they in fact compose; it only settles the biconditional: they compose iff there is some one to which they are identical".[87] In order for CAI to entail universalism, one must already assume that given any *x*s whatsoever, there is a single object to which those *x*s are identical – in other

85 E.g. Merricks (2005), 629; Sider (2007)
86 Merricks (2005), 629. Merricks does not endorse universalism, however. Whilst he does claim that CAI entails universalism, he does not believe that CAI is true.
87 Cameron (2012), 534

words, that there is a single object which those *x*s compose. But that is just to beg the question in favour of universalism!

Cameron's argument is a strong one. For it exposes the argument from CAI to be indubitably invalid. The argument from CAI attempts to argue straight from *p* to q:

*p*: There are some *x*s

*q:* there is a *y,* such that the *x*s = *y.*

In order for it to be valid, it of course needs an extra premise of $p \rightarrow q$. But one cannot assume that! For that would be to assume the truth of universalism and thus beg the question in an objectionable way.

What we can deduce from all this is that while CAI is clearly *compatible* with universalism, it certainly doesn't entail it; for it is compatible with many other stances on composition too. The restrictivist, for example, could say 'of course composition is identity! It just so happens that composition only sometimes occurs, but when it does, the composite object is identical to its parts'. There is no contradiction in such a claim. Even the nihilist could endorse CAI, but simply add the caveat that composition never actually occurs.

I am of the belief that Cameron's argument conclusively defeats the argument from CAI, and this is because I simply cannot see any way of showing that universalism follows from CAI *without* begging the question. But for those who remain unconvinced, there is another line of attack that threatens to undermine the argument from CAI. For it can be claimed – indeed it often is claimed – that CAI is simply not true, or worse, is incoherent. And of course, if CAI is not true, or is incoherent, then it is quite irrelevant as to whether it would entail universalism or not if it *were* true. Firstly, it may be the case that one simply cannot make sense of CAI. That is to say, given a classical conception of identity, it may be just incoherent to suggest that a single thing can be identical to many things taken together.[88]

It is often supposed that CAI fails because it violates Leibniz's Law. For according to Leibniz's law, or more precisely, according to the principle of the indiscernibility of identicals, if $x = y$, then any property of $x$ must also be a property of $y$, and likewise, any property of $y$ must also be a property of $x$. So the thought is that CAI clearly violates this

88 Van Inwagen (1994), for instance, claims not to be able to make sense of CAI.

principle because it supposes that many things (the parts) can be identical to one thing (the whole), but the parts have the property of being many, whereas the whole does not (it has the property of being one). Because of this, CAI is taken to be false. This kind of argument has been presented by a number of different philosophers.[89]

I don't accept this argument, for I think it misunderstands what the thesis of CAI actually states. I think that CAI, when properly understood, does *not* violate the indiscernibility of identicals.[90] To illustrate, consider our six-pack of orange juice again. What CAI states is that the parts are *identical* to the whole; the six bottles *are* the six-pack, and the six-pack *is* the six bottles. But this is surely just to state that the six bottles have the property of being one six-pack! (And of course, the converse holds: the six-pack has the property of being six bottles). But if this were the case then the indiscernibility of identicals would *not* have been contravened, since both the whole and the parts share all the same properties (i.e. the properties of being one six-pack and being six bottles). So to suggest that CAI is in conflict with Leibniz's Law is a mistake, and rests on a misconception of the thesis of CAI. Unfortunately for the defender of CAI, however, this does not make their view any more palatable. For whilst Leibniz's Law remains unscathed, CAI, when viewed this way, involves the claim that composite objects (and their parts) have incompatible properties. A six-pack of orange juice, for instance, simultaneously has the property of being one thing and the property of being six things! If anything, I would suggest that this consequence of CAI is even *more* troubling than the consequence that it contravenes Leibniz's Law. For it is surely impossible for an object to simultaneously have the property of being one thing and the property of being many things! But that, of course, would entail the impossibility of CAI.

But many defenders of CAI accept this conclusion. To overcome the problem, they simply insist that whilst composition is not *strictly* identity, it is something *very similar to* identity.[91] This notion of 'near-identity' or 'almost-identity' has been endorsed by both David Lewis and D. M. Armstrong.[92] But it is Ted Sider who has taken the most effort to explain what it actually is.[93] Sider claims the parthood relation to be uniquely intimate, and

89 Lewis (1991), 87, and McKay (2006), 38, have both presented arguments of this type.
90 I thank Ross Cameron for making this clear to me.
91 See Sider (2007), and Lewis (1991). Baxter is the only philosopher I know of who maintains that CAI uses the strict classical notion of identity
92 Lewis (1991), 84-5; Armstrong (1978), 36-80.
93 Sider (2007)

then goes on to say: "The intimacy of this connection must be explained. The best explanation is a conception of parthood that renders the connection between parts and wholes as intimate and identity-like as possible".[94] These words make it clear that Sider does not take composition to literally *be* identity, but only very similar to it. He aims to highlight this similarity by making a number of claims, which I shall paraphrase below:

1. Just as everything is identical to something, so too any *x*s compose something (this is just universalism).
2. Just as nothing can be identical to two distinct things, so too no *x*s can compose distinct things.
3. Just as identity does not hold relative to place, location, or sortal, so too composition does not hold relative to place, location, or sortal (i.e. composition is absolute).
4. Just as a single identity relation applies to all objects regardless of ontological category, so too a single composition relation applies to all objects regardless of ontological category.
5. Just as identity is never vague, composition is never vague.[95]

These claims are meant to strengthen the case for the claim that composition is 'nearly-identity'. As Sider says, "these theses draw composition as close to identity as possible, without going so far as to identify composition with identity".[96] Despite Sider's insistence, I remain unconvinced. First off, many of the five theses are independently controversial. 1, for instance, is simply an assertion that universalism is true, which is certainly not an uncontroversial claim. 2 could also be disputed, for some philosophers believe that statues, for example, are distinct objects from the lumps of clay/metal/etc that they are made of, but they would not deny that both objects are composed of the same *x*s.[97] 4, too, is dubious to say the least. For as was pointed out in chapter 1, it is quite possible that the relation of

94 Sider (2007), 25
95 These five claims appear in Sider (2007, 27), and are clearly heavily influenced by four claims made by Lewis (1991), 85.
96 Sider (2007), 27
97 See, for instance, Wiggins (1968a)

parthood may behave differently when relating abstract objects to how it behaves when relating material objects. For there is good reason to believe (indeed I have already assumed that this is the case) that *material* parthood is a transitive relation, whereas parthood *per se* is not. (Recall the example that Sanford's Spleen is a part of him, and he is a part of the book Symposium, but his spleen is not a part of the book Symposium?) And this suggests that there isn't a single relation of parthood that holds among all objects, regardless of ontological category, but rather, that there is one relation that holds among material objects and another, different, relation that holds when abstract objects are brought into play. So unless you already happen to believe the five controversial theses that Sider puts forward, it is unlikely, I would suggest, you would be persuaded that they show anything at all substantive about composition and the parthood relation.

But there is, I think, a more pressing problem with the claim that composition is *nearly* identity. The problem is that it is not entirely clear what such a claim is supposed to achieve. The initial advantage of suggesting that composition is identity was that its proponent could claim that composite objects are *nothing over and above* their parts. In slogan fashion: if you've got the parts, you've got the whole. But then it was seen that composition couldn't be identity, *strictly speaking*, because such a claim is incoherent (or maybe even *impossible*). So the next move was to suggest that composition, whilst not the *same* as identity, is *virtually* the same as identity. This is the step with which I wish to take issue. For it strikes me that the notion of identity is a particularly clear and precise one. Not many things in philosophy are black and white, but if anything is, surely identity is a prime candidate. Either *a* and *b* are identical, or they are not. There is no middle ground here, or so it seems to me. And surely the same goes for composition; either composition *is* identity or it is *not.* And if it is not, then composite objects are *not* identical to their parts. But if composite objects are not identical to their parts, then what reason is there to think that they are nothing over and above their parts? What justifies such a claim? I do not know the answer, and I am yet to encounter anyone who does. And this is surely a big problem for the composition-as-nearly-identity theorist.

At this stage, I envisage that the defender of composition-as-nearly-identity would accuse me of being obtuse, or at least uncharitable, in my interpretation of their position. This may well be so, but so far I am yet to encounter any good reason *why* I should be more

charitable, or, in response to the obtuseness charge, it has yet to be explained to me what more there is to understand. Simply put, identity is a precise notion. It does not come in degrees. *Similarity* comes in degrees for sure, and perhaps identity is a limiting case of similarity, but *identity* is just not the sort of thing to vary by degree. So it is of little importance how similar composition may be to identity; if it's not identity it's not identity, and that is that. So consider again David Lewis's claim that "composition is not identity, but it is analogous to it in many important respects", and ask yourself what such a claim achieves. My response would be that it achieves very little. For composition may well be analogous to identity in many respects, but it is not analogous in the most crucial respect of all: *it is not identity*.

To sum up, then, the argument from CAI is entirely unconvincing. First and foremost, CAI in its strictest sense (i.e. when it employs a classical notion of identity) is incoherent, for it would involve objects simultaneously instantiating incompatible properties. The more popular versions of CAI which claim that composition is nearly identity are much more plausible, inasmuch as they have the potential to be true. But they produce results that conflict with the very thought that motivated CAI in the first place: that composites are nothing over and above their parts. But most important of all is the fact that CAI does *not,* contrary to what some have thought, entail universalism. So even if we *can* make sense of it – even if we have reason to believe it may be true – we would still have no extra motivation to believe in universalism, because CAI is quite compatible with other theories of composition, nihilism included.

## §3.2. Arguments *against* Universalism

In what follows I will present four arguments against universalism. I don't think that any of these arguments, when taken individually, is enough to refute universalism outright. Rather, each argument is designed to show that universalism entails some consequence or other that is objectionable or unpalatable. When taken together, however, a number of arguments that may be individually quite weak, have the potential to form a conjunctive argument of much greater strength. I think that is the case here.

## *§3.2.1 The Gratuitousness of Universalism*

It is likely to be noted as a drawback of universalism – indeed it often *is* noted – that it posits too many objects. The ontology of the universalist will, for obvious reasons, be vast. To illustrate, consider once more our afore-mentioned heap-table. According to universalism, not only is there a heap-table composed of the table and the 1000 grains of rice, but there will also be one composed of the table and 999 of the grains of rice, and another composed of the table and 998 of the grains of rice, and so on and so forth. And what about our object composed of the Eiffel Tower and the Great Pyramid at Giza? Well that too will exist. And so will an object composed of the Pyramid and all of the Eiffel Tower bar a single bolt, and so on. And we haven't even begun to consider the countless objects made up of the heap-table (and its constituent parts) *and* the pyramid, *and* the tower. Now I may be labouring the point somewhat here, but the import should be very clear: if universalism is true, there will be a simply astronomical amount of material objects in existence. And this may well put people off - remember Markosian's "fatal objection" to universalism, that it "entails that there are far more composite objects that common sense intuitions allow"?[98] Markosian may be a little premature in suggesting this represents a *fatal* objection, but one thing is for sure: if parsimony is to be considered a theoretical virtue, and it often is amongst philosophers, then all the worse for universalism. Universalism can make no claims of parsimony – quite the opposite; it is gratuitously extravagant, or so the objector may say.

There are two ways (I can think of) that the universalist could respond to this charge. The first strategy I will mention is the one that most universalists do in fact employ, but it is also, in my view, by far the weaker of the two. This strategy is to claim that although Universalism does posit a huge number of objects, most of those objects are, to use the popular terminology, *ontologically innocent.* We have already encountered this kind of idea in §3.1.3 when considering the thesis that composition is identity. Needless to say, then, I am somewhat sceptical of this strategy. The idea is that if some objects compose a composite object, let's say, some atoms compose a table, then the table does not contain any extra matter over and above the atoms that compose it. Therefore, although the table exists, it in some sense *comes for free.* This kind of idea is commonly endorsed. D. M Armstrong, for

98 Markosian (1998a), 233

instance, tells us that "mereological wholes are not ontologically additional to their parts".[99] Similarly, Varzi has said: "the whole and the parts *encompass the same amount of reality* and should not, therefore, be listed separately in an inventory of the world".[100] David Lewis, too, echoes these sentiments by saying "it would be double counting to list the cats and then list their fusion".[101] Espoused by the likes of Lewis and Armstrong, then, the claim that composites are ontologically innocent is certainly backed by philosophical authority. But that does not mean, of course, it is correct.

In fact, it is not immediately clear to me exactly how one is meant to take such a claim. My concern is that if an object is ontologically innocent, then I am unclear as to what sense it is supposed to exist in at all? If a table, for instance, is ontologically innocent, yet one of the atoms (mereological atoms, I mean) that make it up is not (what the opposite of 'innocent' here is not obvious; to call an atom ontologically *guilty* does not seem quite right), then are we to say that these two entities exist *in one and the same sense*? For if they do, then what is actually meant by the term 'ontologically innocent'? But if they do not, then what exactly does the difference consist in? These are difficult questions and their answers not obvious. And it suggests to me that these slogan-like catchphrases such as 'ontologically innocent' or 'no ontological addition' are perhaps not informative to the degree which one really requires. I think one should require a better explanation of *what it is* to be ontologically innocent before one should accept that there is such a status at all.

Particularly revealing, I think, is this mention of 'double counting'. The thought seems to be that if you were making a list of everything there was, then you would only list the mereological simples; you would not list their fusions, because to list the fusions would be to count things twice. This is explicitly expressed in the quotes from Lewis and Varzi above. But it seems to deny the very thing that universalists want to affirm: *that composite objects exist.* Call me old-fashioned if you will, but I believe that *if* something exists, then it should be included on a list of what there is. Because that is just what such a list *is*: *a list of what exists!* To say that there are certain things, yet those things shouldn't appear on a list of

99 Armstrong (1997), 12
100 Varzi (2000), 285
101 Lewis (1991), 81

what there is just seems bizarre, if not plain contradictory.[102] (Imagine saying, 'that apple is red, but it's not making my list of all red things' for example). Now it may be claimed that this is too simplistic a way of looking at it, but in response I could only say that I cannot see any *other* way of looking at it. If something exists then it makes up part of your ontology; and if something makes up part of your ontology, then I am not sure how it can be considered ontologically innocent. (Or if it was still claimed to be ontologically innocent, then I would need an explanation of what that means). Perhaps, however, being ontologically innocent means *not* making up part of your ontology, but then it also must mean *non-existent.* What else could it mean? This is not clear.

Furthermore, if the universalist were not to include composite objects in her inventory of what there is, for fear of double-counting, then it seems that her inventory would be indistinguishable from that of the nihilist, in that it would include only mereological simples. But this cannot be right! The whole dispute between universalists and nihilists is precisely about what does and doesn't exist. If it turns out that their ontological inventories are indistinguishable then it is not clear that there is a disagreement here at all![103] This is surely not the result the universalist envisaged when claiming that composites are ontologically innocent. But to avoid it, much more explanation is required. To insist that composites exist, yet are ontologically innocent, is the metaphysical equivalent of having one's cake and eating it. Armstrong has said that the innocence of composite objects is like an "ontological free lunch".[104] And I think he sums it up perfectly, for as we all know, there ain't no such thing as a free lunch.

The more promising strategy for the universalist, I would suggest, is to simply bite the bullet. *Admit* that universalism is not very parsimonious in terms of the number of entities it posits, but then deny that parsimony in that regard is particularly important. It is a fairly common thought in metaphysics that ontological parsimony should only come into

102 Unless you adopt a Meinongian-type position whereby you concede that there are things that do not exist. This is a particularly controversial position, however, and I am pretty sure it is not what Armstrong *et al* had in mind when they said that composite objects are ontologically innocent.

103 Perhaps an even worse scenario for the universalist would be if the world was gunky, i.e. made of matter whose parts all have further parts. If this were the case, then it seems that the universalist's ontological inventory would be entirely empty, for if the world were gunky then *every* object would be composite and, of course, it's double counting to count the composites.

104 Armstrong (1997), 12. To be a little more precise, the free lunch concerns any cases of *supervenience*. If *a* supervenes on *b,* then *a* comes for free. This applies to composite objects, however, because Armstrong believes they supervene on their parts.

play when one is considering *types* of entity, not merely *tokens.*[105] So it may well be thought that once universalism allows that there are composite objects *at all,* it should not be overly concerned about the *number* of them. This seems a fairly reasonable thought. To give an example, consider Rutherford's postulation that there were sub-atomic particles. Imagine how absurd it would have been if his postulation had been rejected on the grounds that if it were true, then there would be far too many sub-atomic particles! It is merely a consequence of his theory that since there are a lot of atoms, there will of course be a lot of sub-atomic particles. And so it could be said for universalism. It is merely a consequence of the theory that since there are a lot of material objects, there will of course be a lot of fusions of those objects. This is the route that I think the universalist should take. It should be conceded that the ontology of the universalist will be vast, and this ontological profligacy should be taken on the chin. This represents a cost of the theory, but how *much* of a cost is up for debate. I think it is going too far to insist, as Markosian does, that it represents a *fatal* cost, but it has to be recognised as a cost of some sort nonetheless.

### *§3.2.2 The Counter-Intuitiveness of Universalism*

Let's suppose, for sake of argument, that our hypothetical objector is satisfied that universalism is not gratuitously extravagant in its ontological posits, either because many of the posits are seen as 'ontologically innocent', or because she is not much bothered by quantitative parsimony, or perhaps for some other reason I have not considered. She still may be concerned, however, that many of the objects countenanced by the universalist – ontologically innocent or not – are just not the type of things that exist. Common sense, she may say, just doesn't have room for heap-tables or any other gerrymandered objects that the universalist is committed to. Universalism is simply too counter-intuitive to be true. One could certainly imagine an objection of this sort being levelled at the universalist, and one would hope, for the sake of universalism, that a response would be forthcoming

Universalists, of course, are well prepared for this sort of objection, and they do usually come armed with a response. A popular strategy is to suggest that in ordinary thought and talk we tacitly restrict the domain our quantifiers such that they range only

105 I am somewhat sceptical of the value of ontological parsimony *per se.* I will say more about this in chapter 5.

over *ordinary* objects. For we often use quantifiers in a restricted sense in ordinary discourse, even if we don't realise we are doing it. To illustrate this, consider an example: if a mugger steals your wallet, you may tell the police that he stole all your money. But you would not literally mean *all* your money – for presumably the mugger did not also empty your bank accounts as well as gather all the loose change from the back of your sofa. What you would have meant, of course, is that he stole all the money you had with you at the time. So you would have been tacitly restricting the domain of your quantifiers such that they ranged only over the contents of your wallet, or perhaps, over whatever you had on your person. Indeed when you think about it, it seems that in ordinary talk and thought we restrict the domain of our quantifiers all the time (note that I don't really mean *all* the time). So perhaps this is what happens when we talk about composite objects? We tacitly restrict the domain of our quantifiers such that they exclude any *exotic* composites. So *this* is why there seems to be a conflict between universalism and common sense – because common sense employs restricted quantification. When faced with a table on which there are just two apples, for example, an ordinary speaker may say there are just two things on the table, whereas the universalist will say there are three (the two apples and their fusion).[106] But this conflict is merely *apparent*, because the ordinary speaker is using restricted quantifiers. As such, it should not count against universalism as a theory, or so the response goes. This line of thought can be found in the writings of many philosophers sympathetic to the universalist's cause, and can perhaps be summed up best by the following passage from David Lewis:

> Restrict quantifiers not composition. [...] We have no name for the mereological sum of the right half of my left shoe plus the moon plus the sum of all her Majesty's ear-rings, except for the long and clumsy name I just gave it; we have no predicates under which such entities fall, except for technical terms like 'physical object' (in a special sense known to philosophers) or blanket terms like 'entity' and maybe 'thing'; we seldom admit it to our domains of restricted quantification. It is very sensible to ignore such a thing in our ordinary thought and language. But ignoring it won't make it go away.[107]

106 Technically speaking, the universalist will say there are a lot more than three things on the table, like the fusion of the top half of one apple and the bottom half of the other, and so on (presuming, of course, that the apples are not mereologically simple). But this is of little importance, what matters is that the two speakers disagreed at all about how many things there are on the table.
107 Lewis (1986a), 213.

Now despite the fact that this kind of strategy is quite popular amongst universalists, there are reasons why one should be suspicious of it. In fact there are reasons to believe that it is simply not true. Most obvious is the fact that it is *prima facie* implausible; it is very difficult to believe that we actually *do* restrict our domain of quantification in such a way. Undoubtedly we do regularly place restrictions on our quantifiers, but it is not at all obvious that we do so in order to exclude exotic material composites. To illustrate this, consider once more the above example about the mugger. One could imagine a scenario where, on reporting the incident to the authorities, a particularly meticulous officer arched an eyebrow and remarked: "you really mean he stole *all* your money sir; every last penny you owned?" Now whilst you may well be exasperated by such a response, you would at least understand what the policeman meant; you would simply have to reiterate more precisely, that you meant only that the mugger stole all the money that was in your wallet. But now suppose that on telling the officer that there were exactly two items in the wallet – two twenty pound notes, say – he were to respond: "Only *two* items you say? But what about the object which those two notes compose? And what about the fusion of the left half of the first note and the right half of the second? Were you forgetting these items Sir?" Such a question would not exasperate but completely befuddle! It is wildly implausible to suppose that one would casually respond "Oh, I didn't realise you were counting *those* types of object too!"

What these observations show is that whilst we certainly do restrict our quantifiers in certain circumstances, it usually only takes minimal reflection (or perhaps for someone - like a fussy police officer - to point it out to us) for us to realise, *and to accept*, that we are doing so. But there is no controversy there – it is just something that we do. In contrast, it is most controversial to suggest that we regularly restrict our quantifiers to exclude exotic composites. For if you tried to point out to someone that they were doing *that*, it is unlikely that they would even understand what you were talking about, let alone accept that what you said was true. Moreover, once you *had* explained what you meant, it is still very unlikely that they would accept what you have said. Much more likely is that they would insist that the exotic composites you were attempting to refer to simply didn't exist. Seen in

this light, it simply stretches the limits of credibility to suggest that, in ordinary thought and talk, we restrict our quantifiers so as to exclude exotic composites.[108]

It is for this reason that I do not consider the restricted quantifier strategy that is so often employed by universalists to be at all convincing. I think that the counter-intuitiveness of universalism is yet another pill that universalists have to swallow. However, I have to confess that I don't think it is an overly bitter pill. The counter-intuitiveness of universalism is certainly a cost, but it is not a truly exorbitant one. This is for two reasons. Firstly, as I have now mentioned numerous times, there are many reasons to believe that our intuitions about composition are mistaken, as all attempts to reinforce our intuitions with theory seem to end in catastrophic failure. And, of course, if our intuitions about composition are mistaken, then it doesn't seem to be much of a problem if a theory conflicts with them. Secondly, one should remember that at this stage in the dialectic, having eliminated restrictivism as a possible candidate answer to the SCQ, we are left only with a choice between universalism and nihilism. But nihilism has consequences that conflict with our intuitions as well. For universalism may countenance many exotic composites, but nihilism countenances virtually nothing at all. For the nihilist, there are no exotic composites, but there are no ordinary composites either. There are no trees or rocks or cats or dogs, or even people.[109] So it would be unfairly biased to accuse the universalist of counter-intuitiveness whilst not recognising that same accusation could reasonably be levelled at the nihilist. Nonetheless, it still has to be conceded that universalism *does* entail some counter-intuitive consequences, and this constitutes yet another cost of the theory.

## *§3.2.3 The Argument from Primitive Cardinality*

A recent argument has been presented against universalism, by Juan Comesana, which claims that universalism places unacceptable restrictions on the *number* of material objects

108 Korman (2007) investigates this strategy at much greater length than I do here, and should be the first place to look for anyone interested.

109 Unless, of course, the nihilist takes trees and rocks and so on to be mereological simples, but that would be a rare and exotic version of nihilism indeed.

that a world could contain.[110] Comesana claims that universalism is incompatible with the thesis he calls *Primitive Cardinality* (PC):

(PC): For any *n* there could have been exactly *n* material things.[111]

According to PC, then, there is a possible world consisting of just one material thing, a possible world consisting of just two material things, and so on and so forth. More generally, according to PC, the set whose members are the numbers of material objects at a world, across *all* possible worlds, in ascending order, would exactly mirror the set of natural numbers. I.e. it would be {0, 1, 2, 3, 4, 5, 6, ...}.

According to universalism, however, PC is false. To illustrate this, consider a world consisting of exactly two material simples, call it $w_1$. According to universalism, these two simples will compose a further object, call it *O.* Therefore, at $w_1$ there are *three* material things, the two simples and *O.* Now suppose that one of the simples was annihilated. Now at $w_1$ there would be only *one* thing. Because by annihilating one of the simples, you also annihilate *O,* and you are left with just a single simple, which does not have a fusion. Now suppose that another simple was added back to $w_1$. The number of material things the world then contained would immediately go back up to three, because as soon as you have two simples, you also get their fusion. So no matter what you do, you can't have a world with only two things. There is no intermediate world between worlds of three things and worlds of just one thing. Two-thing worlds are, according to universalism, impossible!

Furthermore, universalism does not only rule out the possibility of *two*-thing worlds, but it also rules out the possibility of four-thing worlds, five-thing worlds, six-thing worlds, eight-thing worlds, and countless more. The reason for this is that with the addition of each individual simple, there will also be the automatic addition of numerous fusions composed of the previously existing simples and the newly added simple. More precisely, for any particular number of simples, *n*, the total number of material things (i.e. simples *and* fusions) will be $2^n$-1. So, universalism is incompatible with PC. Instead, the universalist must adhere to a different principle; let's call it Universalist Cardinality (UC):

110 Comesana (2008)
111 Comesana (2008), 32

(UC): For any *n*, there could only be worlds where the number of material things is $2^n$-1.

According to UC, then, the set whose members are the numbers of material objects at a world, across *all* possible worlds, in ascending order, would be {0, 1, 3, 7, 15, 31, ...}.

In light of this, then, we can explicitly formulate Comesana's argument as follows:

1. Universalism is incompatible with PC
2. PC is true
3. Therefore, universalism is false

How forceful one finds this argument will depend entirely on the strength of one's conviction in PC. Comesana claims that intuition supports the truth of PC. He claims that we have "particular pre-theoretical judgments that there could have been exactly two things, and exactly three things, and...", whereas universalism is supported purely by abstract and theoretical principles.[112] Moreover, he claims that it is "standard methodological procedure" in many areas of philosophy to give precedence to pre-theoretical judgments over general theoretical principles, when they conflict. Because of this, he claims that this constitutes *prima facie* evidence in favour of PC.

I think there is a lot to be said for Comesana's argument. For it does seem perfectly reasonable to suppose that there could be a world consisting of just two material objects. I mean, why *couldn't* there be such a world? When taken in isolation, there is nothing logically, epistemically, or metaphysically, objectionable about such a supposition at all. Because of this, I think that once again the universalist will have to simply take it on the chin that her theory entails a somewhat counter-intuitive conclusion. However, I do think that the universalist can soften the blow of the argument somewhat. Firstly, she can point out that whilst universalism is incompatible with PC, it is perfectly compatible with what we could call 'Primitive Simple Cardinality':

(PSC): For any *n*, there could have been exactly *n* simples.

112 Comesana (2008), 2

PSC is perfectly acceptable by universalist lights. Moreover, the universalist might be able to tell some kind of story to the effect that it is actually PSC which our pre-theoretical judgments are driving at, not PC. How plausible such a story could be made to be, however, I am not sure.

Finally, I think that the universalist would be free to remind us that when reflecting on matters of composition and the metaphysics of ordinary objects, it is an unavoidable fact that some of our intuitions will be violated. I have already shown in the previous chapter that there is no plausible way of answering the SCQ in a way that preserves our intuitions about what composite objects do and don't exist. So the universalist can simply say that this argument serves merely to add another entry onto the long list of violated intuitions that accumulates when one concerns oneself with answering this particular metaphysical question. So, the argument from PC does not strike a *fatal* blow against universalism, but it strikes a blow nonetheless. It highlights yet another counter-intuitive conclusion that the universalist must stomach. Moreover, it should be noted that the nihilist is unaffected by the argument, since nihilism is perfectly compatible with PC. So in the ongoing cost-benefit analysis of answers to the SCQ, this strikes another black mark against universalism.

### *§3.2.4 The Identity Argument*

Perhaps the most well-known argument against universalism is that presented by Peter van Inwagen.[113] I won't be presenting van Inwagen's argument here, since it rests on some assumptions which I find questionable.[114] However, what I will be presenting is an argument that has been greatly inspired by van Inwagen's, and one that is similar to it in many respects. I have taken some of the key premises in van Inwagen's argument and reformulated them into an argument which I feel is more neutral and, hopefully, more compelling. The thrust of the argument is that the universalist owes us answers to some particularly tricky questions concerning the identity of ordinary composite objects. If my

113 Van Inwagen (1990), 75

114 Specifically, van Inwagen makes some assumptions that simply beg the question which is up for debate. He assumes, for instance, that he exists and is a composite object. Given that we are currently in the process of trying to answer the SCQ (the question which asks when objects compose), I think it is objectionably biased to assume the existence of any composite objects, regardless of how convinced you may be that they exist.

argument goes through, then the universalist may not be able to countenance that there are any ordinary objects at all. It will be seen that whilst there are ways to get around these conclusions, none of them come without significant theoretical costs.

The argument rests on two crucial theses:

1) If universalism is true, then whether some objects compose has nothing whatsoever to do with the spatial or causal relations that may or may not hold between those objects.
2) If universalism is true, then some particular set of objects, the *x*s, can never compose more than a single object, either simultaneously or successively.

I believe that one *has* to accept these theses if one endorses universalism. I will say a little bit more about *why* I believe this in due course, but first, let me present the argument.

Let us begin by assuming that universalism is true. Now consider a typical ordinary composite object, let's say, a tree – call it 'Sycamore'. According to universalism, Sycamore is a composite object, and more precisely, Sycamore is the fusion of a particular set of simples that are, currently, arranged tree-wise.[115] Let us call this particular set of simples '*S*', and call the fusion of all its members '*F*'. At this point, then, call it *t*, I am assuming that our universalist would uncontroversially endorse the statement 'Sycamore exists at *t*'. Moreover, I am assuming that Sycamore *just is* the fusion of simples in *S*, i.e. Sycamore = *F*.[116] But now suppose, that in an act of divine rage, a thunderbolt was sent to earth and poor Sycamore was completely vaporised. That is to say, she was completely destroyed such that all her constituent parts (i.e. all the simples that are the members of *S*) were separated far and wide throughout the surrounding area. Let us call this point *t'*. Now this is where things begin to get a little peculiar. For at *t'*, I am assuming that Sycamore no longer exists, i.e. the statement 'Sycamore exists at *t*'' is false. And I am assuming that this is uncontroversial since the simples that once composed her are no longer arranged tree-wise, and are no longer collectively engaged in any tree-like activity. (In other words, there is no tree standing where

115 This is not necessarily true, since universalism is compatible with gunk; it could be the case that Sycamore was gunky, and therefore not composed of simples. I will continue as though Sycamore were composed of simples, however, purely for sake of simplicity. The possibility that Sycamore could be gunky does not affect the argument.

116 I mean, what else could Sycamore be? Sycamore just is a fusion of simples that are arranged tree-wise. Likewise, *F* is a fusion of simples that are arranged tree-wise. It seems quite straightforward: Sycamore = *F*.

Sycamore once stood). But *F*, however, still exists, since all the simples that are members of *S* still exist, and so long as they exist, they automatically compose a fusion (i.e. *F*). But earlier we stated, quite in accordance with universalism, that Sycamore = *F*. And this seems to introduce a worrying contradiction. For if Sycamore = *F*, then it is clear that if one of the two does not exist, then neither can the other. But here we are at *t'* in that very situation: *F* exists, yet Sycamore is no more. So the conclusion seems to be that we must have been wrong in initially stating that the simples arranged tree-wise composed a tree. For they still compose *F* at time t', but at t', *F* is definitely not a tree. So those simples can *never* have composed a tree, because according to (2) the same set of simples can *never* compose two (or more) distinct objects, either simultaneously or successively.

So what is the upshot of this argument? Well, considering that it is possible that *any* tree could suffer the same fate as Sycamore, it seems that we have to conclude that if universalism is true, then there are no trees at all. But, of course, it is not only trees that the argument affects. It could be set up with *any* type of ordinary object you care to mention. But seeing as though it is possible for *any* ordinary object to be vaporised and destroyed, it seems that we have no option but to conclude that *if* universalism is true, there are no ordinary objects at all! There are no trees, no cats or dogs, no planets or stars. Indeed, there are no composite objects at all that are capable of being destroyed, and that would even include human beings like you or me. There *are*, of course, mereological fusions like *F*. There will be a whole host of indestructible mereological fusions, composed of all manner of disparate material objects. But these cannot be the ordinary objects of common sense! Because fusions like *F* can survive vaporisation; you can separate their parts to all corners of the universe and they will be entirely unaffected. No *ordinary* object could withstand such a fate. So universalism, it seems, is a most peculiar thesis indeed. If true, then it would have us believe that the world is populated by a vast number of indestructible composite fusions, yet at the same time, is quite empty of all the composite objects which we would normally suppose there to be. But why would one ever believe such a theory? Surely one of the main reasons that defenders of universalism are so keen to preserve composition is that in so doing they can preserve all the ordinary composite objects of common sense. But if this cannot actually be done, as this argument purports to demonstrate, then why bother preserving composition at all?

I will consider some possible responses to this argument in the next section. But first, if one is to be convinced by it at all, I need to say something in support of the two theses that were assumed at the outset. Let's consider (1) first. If universalism is true, then for any material objects whatsoever there is a fusion of those objects. If you have the objects then you *automatically* have their fusion. This is merely a statement of what universalism is. But already, then, it suggests that *whether* objects compose or not has nothing to do with *where* those objects are, or how they are interacting with one another. Moreover, if the spatial arrangement of objects *did* have something to do with whether they composed or not, then that would imply that there are some spatial arrangements of those objects which would *not* be sufficient for them to compose. But we know that is not true. As far as the universalist is concerned, the objects merely have to exist in order to compose. Spatial and/or causal relations never come into it. (1), therefore, is true.

(2) is not quite so obvious, but nonetheless, it is something the universalist must accept. Since universalism is indiscriminate as to which objects compose (they all compose), it is also indiscriminate as to *when* they compose, and *what* they compose. In this respect, universalism is very similar to Set Theory. For any objects whatsoever, there will always be a set which has just those objects as members. As long as those objects exist, the set will also exist. And moreover, the very identity of that set is determined purely by the identity of its members. In other words, it is impossible to have two distinct sets which have exactly the same members. The same is the case for universalism.[117] For any objects whatsoever, there will always be a fusion which has those members as parts. As longs as those objects exist, the fusion will also exist. And moreover, the very identity of that fusion is determined by the identity of its parts. In other words, it is impossible to have two distinct fusions that have exactly the same parts. (2), therefore, is true.

Given the truth of these assumptions, the conclusion of the argument clearly follows. For if ordinary composite objects are taken to be fusions of other material objects, then since those fusions would be able to survive the complete scattering of their parts, the ordinary objects with which they are identified should also be able to survive such scatterings.

117 Indeed it is common for proponents of universalism to argue *for* their thesis on the grounds that it can underpin, or explain, set theory. For example, see Armstrong (1997; 2004) or Lewis (1991)

Ordinary objects clearly cannot survive their parts being scattered far and wide, so the conclusion is clear, ordinary composite objects cannot be mere mereological fusions.

At this point, one may have noted that the argument just presented has a number of similarities to the well-known problem of material constitution.[118] Indeed, one may just think that I am merely rehashing that problem, by exposing the fact that an object has different persistence conditions from the stuff it is made of (e.g. a statue cannot survive being squashed, whereas the lump of clay it is made from can). And if the argument *is* just a restatement of the problem of the statue and the lump, then why can't the universalist just endorse one of the many responses that have been given to that case?

The first thing I would say to this response is that even if the argument *were* the very same problem as the statue and the lump, then it is still a problem nonetheless. Just because the problem has been noted before doesn't make it any less problematic. However, whilst I can see that the argument does have its similarities to the problem of the statue and the lump, it is *not,* in fact, the same problem. Actually, I think it is a problem that cuts much deeper than the statue and the lump. For according to this argument, the universalist could not even allow that there are lumps of clay, let alone any statues that they constitute. For if the fusion, *F,* of some simples arranged lump-wise is taken to be a lump of clay, then the same problem arises. For the fusion would survive even if its atoms were scattered to the furthest corners of the universe, but the lump of clay would not. So contrary to our initial suggestion, *F* is not actually a lump of clay, but something else entirely. And if fusions of simples arranged lump-wise are not lumps of clay, then what *are* lumps of clay? And, indeed, what on earth are these mysterious fusions?

So the argument is in fact much more severe than the problem of the statue and the lump. It shows that if universalism is true, it cannot explain what ordinary composite objects are. They cannot be the ubiquitous mereological fusions that the universalist posits, since they have vastly different properties. So what *are* ordinary composite objects, if they are not mereological fusions, and moreover, what are these mysterious, indestructible fusions that allegedly populate our world?

118 For those not familiar with this problem, I will be going over it in more detail in the next chapter.

## *§3.2.5 A Response to the Identity Argument: Constitution*

There are a various ways in which one could respond to the identity argument. Whilst some of these responses are more successful than others, I think that *all* of them come at some sort of cost, as I shall presently show.

Firstly, one could endorse what has been called *the constitution view.*[119] This view would employ a technical relation of *constitution.* Constitution is taken to be a particularly intimate relation that holds between spatially coincident objects. Crucially, however, although this relation is taken to be intimate, it is not taken to be *identity.* This view is often adopted in response to the problem of the statue and the lump, such that the statue and the lump are taken to be distinct (albeit spatially coincident) objects, and the latter constitutes the former.[120] In response to the identity argument, then, the response would be that ordinary objects are *constituted* by mereological fusions, but are not *identical* to them.

This view would overcome the identity argument. The reason for this is that it would render a crucial premise in the argument false. The premise in question is that which states that Sycamore = *F* at *t.* According to the constitution view, Sycamore is *not* identical to *F*, at *any* time. Rather, *F* constitutes Sycamore at *t*, and no longer constitutes *F* at *t'*. This means that there is no problem of a single object having incompatible properties, because there are actually two distinct objects. One of those objects (*F*) can survive its parts being scattered far and wide, whereas the other (Sycamore) cannot.

I am not at all convinced by this response, for I think that it is problematic in many ways. First of all, there are all the traditional problems that go with the view, such as the fact that it entails the rather unsettling conclusion that two distinct material objects can be in the same place at the same time. In fact it entails the *ubiquitous* co-location of material objects.[121] But I think there are some much more damaging problems with the view than that. To draw these problems out, I want to pose a question about constitution that is closely analogous to the SCQ. I will call it the Special Constitution Question (SCNQ):

---

119 This term is taken from Wasserman (2004), 693

120 For defences of this view see Wiggins (1968a); Doepke (1982; 1996); Thompson (1998); Simons (1987); among others.

121 Interestingly, according to this view, the traditional problem of the statue and the lump would involve not two objects occupying the same place at the same time, but three: a mereological fusion of simples, a lump of clay, and a statue. But I imagine the constitution theorist would not be all that worried. After all, if two objects can happily be co-located, then why not three?

(SCNQ): For any mereological fusion, *F*, what are the jointly necessary and sufficient conditions required for *F* to constitute another object, *O*.

The precise way in which one answers this question will hinge significantly on the analysis one provides of the constitution relation itself. There have been a number of attempts to provide such an analysis, none of which are entirely without faults.[122] However, what seems pretty clear is that, regardless of the specific analysis of constitution one gives, to give the constitution view as a response to the identity argument at all, one must accept that constitution is *restricted*. To see why, just consider our previous example. The constitution view gets around the problem by saying that at *t*, *F* constitutes Sycamore, but at *t'*, *F* doesn't constitute Sycamore. Furthermore, it seems fair to say that at *t'*, when the parts of *F* are scattered far and wide, *F* doesn't constitute anything at all.[123] But that shows that constitution is restricted, since there are cases when constitution does occur and cases when it does not. Because the constitution view entails that constitution is restricted, I think that it suffers from many of the same problems that afflict restrictivism about composition. For instance, it will be difficult, I assume, to give any principled answer to the SCNQ that does not admit of any troublesome counter-examples. Most significantly, however, I think that restricted constitution, just like restricted composition, will feel the full force of the sorites.

To illustrate why this is so, let us revisit the vagueness game that was setup in the previous chapter. This time, involved in the game will be you, God, and our tree, Sycamore. When the game begins, God will take one of the many simples that compose *F* (which is the fusion that constitutes Sycamore) and place it in a far corner of the universe, millions of miles from where your game is taking place. After He does this, you ask Him: 'does *F* constitute Sycamore?', and God, playing the game in good spirit, will answer to the best of His omniscient ability. The game will continue, and after each question He answers, God will then take another simple from *F* and move it to a *different* location, that is also millions of miles from where you are. At the beginning of the game, *F* clearly constitutes Sycamore, so God's answer to your first question will be 'yes'. But if the game were to be played to its

122 See Wasserman (2004) for a good survey of some of the most significant attempts.
123 At least I cannot see what on earth such a fusion *would* constitute. Certainly not a tree.

conclusion, then it is fair to say that *F* will *not* constitute Sycamore at the end, since the parts of *F* will be scattered in all corners of the universe. So the dilemma should be fairly clear. There *must* reach a stage in the game at which God's answer to the question will change. It may simply change to 'no', which would indicate a sharp cut-off point between cases of constitution and non-constitution (or, perhaps, 'nonstitution'). Or maybe his answer will change to something less definite, indicating that there is a phase of indeterminacy, during which it is a vague matter as to whether constitution occurs. But either way, the conclusion is clear: the difference of a *single* simple, a truly miniscule existent, can supposedly make the difference between there being a tree there or not. And just as it was in the restricted composition case, this conclusion is surely unacceptable. It is simply absurd to suggest that the displacement of a single simple, an object so small that 3,000,000,000,000,000,000,000,000,00 of them can fit in a space the size of an average cat, could make any significant difference to the ontological status of a tree.[124] It is for this reason, then, that the constitution response should be rejected.

## *§3.2.6 Another Response to the Identity Argument: 4–Dimensionalism*

There is another, and much better, response to the identity argument than the constitution view. The response is to endorse a particular view of how objects persist through time: *four-dimensionalism*. Very roughly, four-dimensionalism is the view that material objects extend through time in very much the same way in which they extend through space. As such, in the same way that material objects are taken to have spatial parts at the different regions of space they occupy, they are also taken to have *temporal parts* for each of the different times at which they exist. On this view, then, material objects are not three-dimensional, as we may ordinarily think of them, but rather, four-dimensional hunks of matter.[125]

A consequence of four-dimensionalism is that objects are not *wholly present* at any particular moment of time. Rather, they merely have a temporal part that is wholly present at each moment of their existence. At first, this seems to clash a little with intuition. For as I

124 And this is only if one identifies simples with the fundamental particles of physics. One could take them to be even smaller, or even point-sized. In that case, each simple would be even less significant, and the argument even more forceful.

125 The term 'four-dimensional hunks of matter' is taken from Heller (1990).

sit here writing these very words, it *seems* as though it is the *whole* of me that is so doing, not merely a *part* of me. But this clash with intuition is fairly superficial. To see why, consider an analogy. Take the river Thames. The source of this river is somewhere near Cirencester, in South Gloucestershire. From there it meanders its way eastwards until it finally reaches the North Sea near Southend. On its way it flows through numerous villages and towns. But, of course, the river is not *wholly present* in any of these places, but merely *partly* so. It would be plain wrong to insist that the *whole* river was present at, say, Sunbury-on-Thames. Similarly, claims the four-dimensionalist, it would be plain wrong to insist that the *whole* river is present at any particular time. Rather, only a part of the river is present. The whole river is the sum of all its parts, both spatial *and* temporal.

So how does Four-Dimensionalism overcome the argument? Interestingly, just like the constitution theorist, the four-dimensionalist will deny the premise which claims Sycamore = *F* at *t*. Instead, Sycamore and *F* are taken to be distinct, four-dimensional objects that merely share a temporal part. Specifically, they share the temporal part at which all the parts of *F* are arranged tree-wise. So according to this view, there are not two distinct objects located in the same place at the same time. Rather, at *t*, there is a single object present, which is a (temporal) part of two distinct objects, Sycamore and *F*. Ted Sider sums it up nicely:

> Whenever distinct material objects coincide, they are never at that time wholly present, but rather overlap in a shared temporal slice or segment. Coincident objects are therefore no more mysterious or objectionable than overlapping roads.[126]

I concede that this response is enough to overcome the argument. But it comes at a cost. For what the argument shows is that *if* the universalist wants to hang on to the ordinary objects of common sense (and I am assuming that most universalists will), then she will also be *forced* into accepting some version of four-dimensionalism. This is not fatal to her thesis, of course, for four-dimensionalism is independently interesting in its own right. Indeed it is a thesis endorsed by a number of eminent philosophers.[127] However, it is certainly controversial. The debate over how objects persist through time is one of the most hotly

126 Sider (2001), 152

127 See, for instance, Heller (1990); Lewis (1986a), 218-220; Le Poidevin (1991); Sider (2001). And these are just a few from a much longer list.

contested in contemporary metaphysics, so for universalism to force one into a particular position regarding this debate must be viewed as a significant theoretical cost. If you are already sympathetic to four-dimensionalism, then this cost may not worry you a great deal. On the other hand, however, if you are inclined to disbelieve four-dimensionalism then you may see this argument as a reason to reject universalism outright. I will not go into the relative merits and demerits of four-dimensionalism here. That would be a different and lengthy project.[128] I will simply conclude by saying that if any metaphysical view to forces one to take sides on an unrelated and controversial debate, then regardless of your predisposition towards the side it forces one towards, it should be viewed as a cost of that metaphysical view. That universalism comes with such a cost is undeniable.

## §3.3. Concluding Remarks

I hope to have made a strong case against universalism. I have shown that a number of the key arguments that are often presented in support of universalism are unconvincing. And in the absence of any supporting arguments, of course, there is little motivation for endorsing a theory. I then went on to present a number of arguments against universalism. As I said at the outset, I don't envisage any of these arguments to be entirely conclusive when taken alone, but when taken together, they show that universalism comes replete with significant, and costly, baggage. Firstly, it is ontologically gratuitous. Secondly, it is seriously counter-intuitive. Thirdly, it forces one to endorse a controversial metaphysical view concerning the persistence of objects through time. Overall, then, the case against universalism is a strong one. And I hope it to get even stronger in the next chapter, when I will present arguments in support of one of universalism's rival views: compositional nihilism.

128 Although I should perhaps note that I do *not* endorse four-dimensionalism. One may not be surprised to learn that one as sceptical as I am about parts is even *more* sceptical about the possibility of *temporal* parts.

# §4. Compositional Nihilism

Having given a fair amount of consideration to both restrictivism and universalism, it is now time to consider the remaining candidate answer to the SCQ: nihilism. We can define it as follows:

> **Nihilism:** for any two or more *x*s, where the *x*s are material objects, there is *never* a further object which those *x*s compose.

Perhaps more simply put, according to nihilism, there are no material composite objects at all; all material objects are mereologically simple.

I believe that nihilism is true. In this chapter I will try to explain *why* I think it is true by providing arguments in its favour, and in so doing, I also hope to convince you, the reader, that it *is* true. Some of these arguments will be somewhat negative, in that rather than presenting a positive case *for* nihilism, they will argue against its alternatives. Some arguments will be indirect, in that they will argue against arguments against nihilism. Others will be more straightforward; arguments that make a positive case for nihilism being true. But before I get to the arguments, I want to clarify, more precisely, what the thesis of nihilism is, and how it admits of different variants.

## §4.1. Varieties of Nihilism

Nihilism is a very broad thesis; there are many different nihilisms. It states solely that material objects are mereologically simple. It says nothing more than that. But, of course, there is much more to a material object than its mereological structure. There are all sorts of other features that material objects may have, or may lack. And so it is that there are many different varieties of nihilism. All varieties of nihilism posit only material simples, but they can vary significantly as to what those simples they posit are like. Specifically, I want to draw a distinction between varieties of nihilism that vary in terms of the *size* of the simples

they posit. At one end of the spectrum, a nihilist could posit material simples that are *minimally small*, i.e. point-sized. I will call this view 'Punctal Nihilism'. At the other end of the spectrum, a nihilist could posit a single simple that is *maximally large*, i.e. world-sized. I will call this view 'Monistic Nihilism' (or just 'Monism'). In between these extremes there lies a whole host of possible varieties of nihilism, all of which that posit simples which are neither minimally small, nor maximally large. I will group all of these possible varieties under the heading of 'Intermediate Nihilism'.[129]

As one can see, then, nihilistic views can vary significantly. The difference between punctal nihilism and monistic nihilism, for instance, is stark. For the former will posit, presumably, a simply astronomical number of simples, whereas the latter will posit only one – the world itself. And intermediate nihilism could range from a variety which posits a vast amount of relatively small simples, to one which posits just a handful of relatively large simples, or it could even posit simples of all manner of different shapes and sizes. These would all represent vastly different ontologies, yet they are all nihilistic. In the following chapter, and indeed in the remainder of the thesis, I will argue that it is monistic nihilism that should be endorsed, but that is for later. For now it will suffice that we merely recognise that *there is* such a diverse taxonomy of nihilisms to choose from.

## §4.2. The Argument from Elimination

As I hopefully made clear in Chapter Two, restrictivism, universalism, and nihilism represent an exhaustive list of responses to the SCQ. Whatever answer you give to the SCQ, it must fall under one of these headings. When faced with a complete and exhaustive list of solutions to a problem, one always has the option of arriving at an answer to that problem through a process of elimination. If one knows that the list is exhaustive, and one can eliminate certain entries on that list, then one will also know that the correct answer must lie among those entries that remain. I think that if one follows this type of process when tackling the SCQ, one should arrive at the conclusion that nihilism is true. I have shown in

[129] Schaffer (2007) identifies a similar triumvirate of varieties of nihilism, although the terminology he uses is different from mine.

the last two chapters that there are a number of reasons why both restrictivism and universalism are problematic; *so* problematic, in my view, that they should be rejected as false. If I am right about this, then we have a very simple, yet very persuasive argument for Compositional Nihilism; an argument from elimination:

1. Either restrictivism, universalism, or nihilism is true.
2. Restrictivism is false.
3. Universalism is false.
4. Therefore, nihilism is true.

In my view, this is the strongest argument there is for nihilism. But that is because I have a very strong conviction in the truth of premises 2 and 3, *and* I think that nihilism is an independently plausible thesis. However, I anticipate that the reader may not be immediately convinced, and that is even if they agree with me that both restrictivism and universalism are, in many ways, problematic. The reason for this is that nihilism, too, is often seen to be problematic. And if the problems that face nihilism are seen to be equally troubling, or even more troubling, than those which face restrictivism and/or universalism, then the argument from elimination is unlikely to carry much weight.

Nihilism *is* commonly seen to be a problematic thesis. Many philosophers have rejected it out of hand as being not only false, but *patently* false.[130] The main reason for this seems to be that nihilism denies the existence of many things which obviously *do* exist. It is simply *obvious* that there are composite objects, like cats and dogs and tables and chairs, and so on. But nihilism denies this, so it *must* be false, or so the common objection goes. Nihilism, on this view, clashes with evident fact; but the facts are facts, so it must be nihilism that is wrong. This, then, represents the first challenge for the nihilist. In order for her thesis to be even taken seriously, she needs to overcome this objection. In the proceeding section that is precisely what I shall do.

[130] See, for example, Rea (1998); Markosian (1998b); Schaffer (2009), 358.

## §4.3. Defeating the Common Sense Objection

If nihilism is true, then there are no composite objects, there are only mereological simples. But there *are* composite objects! There are cats, dogs, trees and rocks; there are mats, logs, bees and socks. The world is populated by ordinary objects like this, yet they are all composite – *they all have parts.* Because of this, it is often said, nihilism is clearly false. We have already encountered this kind of objection from Markosian, who considered it a "fatal flaw" in nihilism. But the view is shared by many eminent thinkers, and is perhaps best summed up by Michael Rea: "It is just obvious that there are tables, chairs, computers and cars. The fact that some philosophical arguments suggest otherwise seems simply an indication that something has gone wrong with those arguments".[131] Nihilism, on this view, clashes with evident fact; but the facts are facts, so it must be nihilism that is wrong.

If I am honest, the level of credence that this view is often given never ceases to amaze me. For it is so uncritical; so uncharitable; so *naive.* Such a view blankly fails to see *exactly what it is* that the nihilist is proposing. As I said in chapter 2, just because nihilism denies that there are composite objects, like chairs, for instance, this does not entail that we have nowhere to sit down! If it *did* entail that then it *would* be absurd, and the common sense view set out above would be perfectly reasonable. But, of course, this is *not* what the nihilist claims. The nihilist agrees with the objector that there is plenty of matter in the world; she merely insists that there is no compositional structure to that matter. Certain consequences of nihilism, when stated bluntly, such as "there are no chairs", certainly *seem* a little peculiar, but they seem a whole lot less peculiar once one goes beyond their mere surface appearance and begins to see exactly what this thesis is that the nihilist is proposing.

To make the point more clearly, consider an analogy. A significant dispute in contemporary metaphysics concerns the existence of mathematical objects. Very roughly, some say they exist (we'll call them 'Platonists') and others say they don't (we'll call them 'Nominalists'). So according to the Nominalist, there is no such thing as the number three. But imagine the naivety of a Platonist who claimed, "But I have *three* children! And I am more than *three* feet tall! And I had *three* Weetabix for breakfast! So your view is obviously false". This objection fails to get to the heart of the matter. It is to conceive of the debate far

[131] Rea (1998), 348

too simplistically. For even if the Nominalist's view is correct, then she will still be able to have her Weetabix in the morning; she will still be the same height; she will still be able to count! Her view does not entail that she is unable to recognise, or quantify, different amounts of things. If it did, then it would be an absurd view, and the Platonist's objection would be perfectly valid. But it doesn't, and for her opponent not to recognise this is to be unreasonably obtuse. Of course, the nominalist must have some kind of way of making sense of statements that *seem* to refer to numbers, in a way that does not involve the actual existence of numbers, but there are ways of doing that.[132] In fact, we all do that kind of thing, in all sorts of circumstances, very often. When the weather forecaster tells you that the sun will rise at 5.46am tomorrow morning, you don't brand him a fool. You can make perfect sense of what he says, despite the fact that you know that the sun, in reality, will not rise at all.[133]

In fact, the 'rising' sun example serves to illustrate the problem nicely. Suppose you were transported back to a time at which science had not uncovered any of the true facts about planetary motion. For the people at that time, it was taken to be the case that the Earth beneath them was fixed and stationary, and that the sun moved around them through the sky.[134] As far as they were concerned, the sun did, quite literally, rise each morning and set each evening. Now suppose that you were to baldly assert to these people, "the sun does not rise". They would not believe you. And rightly so; they would have no reason to believe you, for what you said was so clearly at odds with the facts. Just wait until morning, they would say, and you will see it rise. But then suppose that you were to explain to them some rudimentary facts about the orbits of the planets around the sun, and the rotation of the Earth upon its axis. You could give them a perfectly good explanation of why it *seems* like the sun rises, even though it does not rise. Your initial assertion, *when coupled with this explanation,* no longer seems so contrary to the facts. In fact, it is perfectly compatible with the facts. For if what you said was true, then the experience of seeing the sun 'rise' would remain exactly the same. Now here is the point. Whether these ancient people would be

[132] Field (1980) provides what is probably the most well-known attempt to do just that. He claims to show how one can make sense of modern physics without positing mathematical objects.

[133] That is, the sun doesn't *literally* rise. It remains stationary whilst the rotation of the Earth makes it *look* as though it rises from the perspective of Earth-bound observers.

[134] I am presuming that there was a time when this was the case. But even if people did not think this, at *any* time, it really doesn't matter, just imagine that there was such a time, and there were such a people.

convinced by your explanation is up for debate. But what *is* clear is that if these people were at all rational, then they could not simply and flatly deny your assertion in the same way that they did before. They could not simply say, "you're wrong; wait until morning, and you'll see the sun rise". That would be an objectionably obtuse response, since as your explanation made clear, *your* story about the sun being stationary is equally as compatible with the facts as their story about the sun rising. Their initial objection (*of course the sun rises!*) has lost all of its bite, and it would be unreasonable for them to keep faith in it.

I think that the common sense objection to nihilism is almost perfectly analogous to the rising sun objection (i.e. the objection "of course the sun rises, I saw it rise this morning!"). If a nihilist were to flatly assert, "there are no composite objects: no chairs, no tables, no cats or dogs", and simply leave it at that, then one should expect, and accept, the common sense response: "you're wrong! *There is* a cat, and *there is* a chair", and so on. But if the nihilist were to give some coherent explanation of why it *seems* like there is a chair and a cat, when in fact there are no such things, then the common sense objection just won't do. One doesn't have to accept that the nihilist's explanation is *true,* of course, but one *does* have to accept that since their explanation is compatible with experience, it is no longer a reasonable objection to simply say, "you're wrong! There is a cat!" To continue with such an objection would be unreasonably obtuse, in the very same way that it would be overly obtuse to insist that the sun *does* rise, in light of a satisfactory explanation of planetary motion. Crucially, I think that the nihilist *can* provide such an explanation, as I shall shortly show. Because of this, the common sense objection should be forgotten. It no longer poses any threat to the nihilist.

The most well-known explanation that the nihilist could give would be that provided by van Inwagen.[135] The idea is that whilst there are no composite objects, like chairs, tables, cats and dogs, there *are* simples, and those simples are arranged in particular ways (to use van Inwagen's terminology, they can be arranged chair-wise, table-wise, cat-wise and dog-wise, and so on). The reason why this explanation works is because simples arranged $x$-wise would surely be sufficient to account for the appearance of there being an $x$. Suppose that you accept that macroscopic objects are made up of microscopic atoms. Therefore, you

[135] Van Inwagen (1990), 109. Of course, van Inwagen is not a nihilist, but he does reject the existence of most ordinary composites, which is why he needed to come up with the explanation.

would accept that a chair, for instance, is made up of lots of tiny atoms arranged in a particular way. But the nihilist, on this view at least, *allows* that there are those atoms, and *allows* that they can be arranged in that particular way. She simply denies that they compose a chair. Now suppose that the nihilist is right, and that those atoms exist, and they are arranged in that precise way, yet they do not compose a chair. It would surely still *seem* as though there was a chair in front of you, would it not? The atoms are surely sufficient to give the *appearance* of a chair, even if they don't in fact *compose* a chair? It seems very difficult to find any reason to suppose otherwise. But if this is right, then the nihilist's explanation should be taken seriously at the very least. For *in light* of this explanation, it should no longer seem absurd that the nihilist denies there is a chair, for she admits that there are atoms arranged chair-wise. Denying the existence of the chair is, as Ted Sider says, "no more absurd than denying holes exist in addition to perforated donuts, or denying that smirks exist in addition to smirking faces".[136]

This nihilistic explanation has enjoyed a certain degree of success in recent years.[137] But it is important to note that if one endorses it, then it forces one to accept a particular variety of nihilism (or, more precisely, it forces one to rule out certain varieties of nihilism). If, for example, the appearance of a chair is to be explained in terms of there being simples arranged chair-wise, then that entails that there must be a *plurality* of simples. So this type of explanation is not available to the monistic nihilist, since she posits only a single simple – the world itself. Moreover, this explanation entails that the simples arranged chair-wise must be fairly small. At least, they must be smaller than the alleged chair which they are taken to be arranged like. (At least, I *think* this must be the case. For I don't see how a number of objects, the $x$s, could be arranged $y$-wise, if each of the $x$s was bigger than $y$. I will simply assume that would be impossible). So this van Inwagen style explanation forces the nihilist into accepting a fairly specific variant of nihilism: one which posits a plurality of small simples.

But what about other varieties of nihilism? What about varieties that posit a small number of very large simples? What explanation can *they* provide to account for the

[136] Sider (forthcoming), 2

[137] The vast majority of philosophers who endorse nihilism in print, also endorse this explanation for the appearance of composite objects. For instance, see Sider (forthcoming); Dorr (2005); Cameron (2008a; 2010b). Even those who argue against nihilism seem to accept that it is an explanatory strategy that should at least be taken seriously. See for instance, van Cleve (2008), 330; Schaffer (2007), 176. However, see Uzquiano (2004) for reasons to object.

appearance of composite objects? For ease of exposition, I will focus purely on monism in what follows, but I think that the monistic explanation will carry over to any variant of nihilism which posits very large simples. Very roughly, I think that the monist can explain the appearance of composite objects (chairs and tables and so on) through an appeal to *ways the world is*. But this is *very* rough, so let me try to explain a little more precisely.

If monism is true, then the world is a single giant simple. It is a maximally large lump of matter that has no proper parts. But it *seems* like it has parts. For it seems like there are cats and dogs and chairs and tables, and so on, and that these are *parts of the world*. So the monist has to be able to explain these appearances. In later chapters I will go into a lot more detail about how, *precisely*, the monist should do this. But for now it should suffice to give a brief summary. Firstly, if you take monism to be a serious possibility, then it follows that you also take it to be a serious possibility that spatially extended, yet mereologically simple objects can exhibit qualitative variation. For if monism is true, then the world *is* a spatially extended simple, and the world clearly exhibits qualitative variation. But if you accept this, then it should be fairly easy to see how a monistic world could give rise to the appearance of its having parts. To illustrate, consider an extended simple that looks like the object depicted in *fig. 1.*

*Fig. 1*

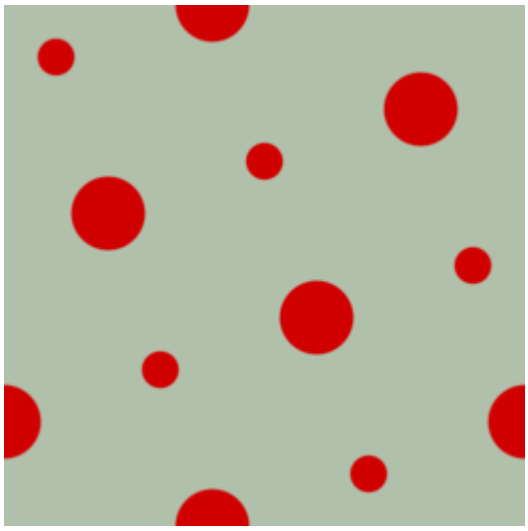

If you were to come across an object like this, it would be quite natural for you to think that it had parts. If you were to describe it to someone, for instance, you may say something like, "it is grey, and it has red spots". This description would seem to imply that it has *parts* which are grey, and *parts* which are red. But it doesn't. By stipulation, it is a mereologically simple object, with no parts whatsoever. However, the *way that object is*, or perhaps more precisely, the properties or features that it exhibits, give rise to the appearance that it *does* have parts. The object, as a whole, exhibits certain qualities which make it such that it is perfectly reasonable for us to presume that it is made of parts, despite the fact that it is not made of parts. You may say that it has a certain *colour variation property*, for instance. Or you may simply say it has the property of *being spotty*. But what we *call* these qualities is not really that important. What is important is that an extended simple object can exhibit properties that make it seem as though it has parts when it does not.

If you accept that this explanation is possible, then there seems no reason why the same explanation could not be applied on the grandest of scales – to the world itself.[138] Thus the world, on this view, would be a giant extended simple that exhibits certain properties which give rise to the appearance that it has parts, when in reality it does not. Thus where you think you see a chair, for instance, you do not, but the world does instantiate properties that suffice to account for the *appearance* of a chair. You might say that the world is *chair-ish* at that particular location, for example. Or you might say that whilst there is no chair, there is a chair-like thickening in the world-substance, or something like that. But however you choose to describe it, the important fact is that the world's exhibiting these particular properties is sufficient to account for the appearance of a chair (and of course, for all manner of sub-world objects that it may seem like one perceives). And thus once again, in light of this explanation, it is no longer justified for the objector to bang his fist and say "but *there is* a chair". Such an objection is obtuse. Of course, it *could* be a chair. But as has just been shown, it could also be some simples arranged chair-wise, or, indeed, it could be a chair-like thickening in the world-substance. Because of this, then, the common sense objection has no

[138] I envisage that certain readers will not agree that such an explanation *is* possible. For there is a fairly widespread opposition to the claim that extended simples are possible at all, let alone ones that exhibit qualitative variation. For these sceptical readers, I would simply ask you to withhold your incredulity for the time being. I will provide independent arguments for the possibility of extended simples in chapter 6, and I will go into greater detail about the properties required for such objects to exhibit qualitative variation in chapters 6 and 7.

force. In light of there being alternative, and coherent, nihilist explanations for the *appearance* of composite objects, the common sense objector can no longer flatly assert that nihilism is at odds with the evident facts. Quite simply, it is not.

## §4.4. Who Needs Composite Objects Anyway?

According to the standard scientific picture, the world is one of many layers. The macroscopic objects that we interact with from day to day are actually made up of much smaller parts: microscopic particles so small as to be far beyond the perceptual range of the naked eye. Moreover, according to this picture, accurate and complete explanations of phenomena at the macroscopic level can be given *solely* in terms of these microphysical particles and their interactions. To illustrate, consider that philosophers' favourite example of a macroscopic event: one billiard ball striking another. When describing such an event, we may commonly say things such as "the first billiard ball caused the second to move" or "the impact of the first ball with the second caused the 'click' sound". But these are not entirely perspicuous descriptions. An accurate scientific description of this event would make no reference to macroscopic objects like billiard balls. Rather, it would be framed in terms of a certain number of microscopic particles, exemplifying certain physical properties, interacting with a certain number of other microscopic particles, or something along those lines.[139] More generally, if we can explain, *precisely,* the activity and interaction of the particles, then we have explained all there is to explain. Moreover, we can even explain one's *experience* of the event, without once having to appeal to macroscopic objects like billiard balls. Such an explanation would consist of a story about how photons of light interact with the particles, and then reflect back into your range of vision, or something along those lines.

If this is right, then the *full causal story* behind all events at the macroscopic level does not involve any macroscopic objects. But this invites a question: what role, precisely, are macroscopic objects supposed to play in the evolution of the physical world? If the microphysical particles do all the work on their own, so to speak, then why should we

[139] It is likely to be a lot more complex that this, of course. But the central import should be clear: whatever the precise scientific explanation *is,* it will not involve any macroscopic objects like billiard balls.

suppose that there are any macrophysical objects that they compose at all? What would the point of such objects be? Let me put that another way. Consider the billiard balls again. Let's call the event of one of them striking the other, *E.* If, as the scientific picture supposes, the full causal story behind *E*, including your perceptual experience of *E*, is explainable entirely in terms of the activity of microphysical particles, then it is surely fair to assume that *E* would be exactly the same, as would your perceptual experience of *E*, whether or not the particles involved did compose macroscopic objects, like billiard balls, or not. For if they did *not* compose billiard balls, then *E* and your experience of *E*, would be just the same, since, *ex hypothesi,* the particles alone are all that are needed to explain *E* and your experience of it. And if they *did* compose billiard balls, then those balls would not change, *in any way*, *E* or your experience of *E*, for if they did, then we would have been wrong to say that the full story could be told in terms of the particles alone.

This suggests that ordinary composite objects, if indeed there are any such things, are not quite as ordinary as we may have thought. For if the above reasoning is right, it looks as though composite objects have no causal powers; they have no causal impact on the physical world whatsoever. They are, to use the technical jargon, *epiphenomena.* In the next section I will talk at greater length about the severe problems that arise if one views composite objects to be epiphenomena, but for now I just want to question the motivation for positing such things in the first place. For if ordinary composite objects have no causal impact on us at all, and if events and appearances can be explained perfectly well in terms of only microphysical particles, then why should we bother positing such objects at all? There just doesn't seem to be any good reason why we should. Cian Dorr effectively summarises this view:

> If all the plates in my kitchen dresser were to cease to exist, but all the molecules in my dresser were to stay arranged exactly as they are, I wouldn't care very much. My guests would have no new reason to worry about their food getting all over the tablecloth. In fact, they would never know unless I told them – but come to think of it, I would never know either.[140]

[140] Dorr (2002), 42-3

Once this is recognised, it really seems quite straightforward. There is just no need to posit the particles *and* the billiard balls. Or more generally, there is no need to posit parts *and* wholes. Since the activity and interaction of microphysical particles is enough to adequately explain all macrophysical phenomena, it would be nothing short of gratuitous to posit epiphenomenal macroscopic objects as well.

These thoughts form the beginnings of what are generally known as 'causal redundancy arguments', and variations of them have been advanced by a number of thinkers over recent years.[141] Even in a mere rough outline, as I have sketched above, I think these types of arguments very powerful. In the next section, however, I will follow Trenton Merricks, and show how appeals to causal redundancy, and causal overdetermination, can be used to create a *fatal* argument against the existence of composite objects.

## §4.5. The Causal Overdetermination Argument

In most instances, causal redundancy arguments have been used to show merely that there is no need to posit composite objects.[142] But in his 2001 book *Objects and Persons,* Trenton Merricks, uses a similar style of argument to argue for a much stronger conclusion: that there *can't* be any composite objects.[143] Merricks's argument represents one of the most persuasive reasons to endorse nihilism. I will set out the argument below. Whilst I will stick fairly closely to Merricks's original formulation of the argument, I will alter it somewhat. This is for increased clarity and will, I hope, make the argument even more compelling. To begin, consider a standard causal event: a baseball hitting a window and causing it to smash. Now let me set out the argument.

[141] See Dorr (2002), 41-45; Schaffer (2007), 176-178; van Inwagen (1990), 122; Merricks (2001), 56. This style of argument is much more well-known within the philosophy of mind to make the case for physicalism, by arguing that non-physical mental states, if there were any such things, would be causally redundant.

[142] E.g. Dorr (2002), van Inwagen (1990)

[143] Merricks (2001), 56. This is not strictly true. Merricks actually endorses a similar view to that of van Inwagen, such that there are *some* composite objects, namely, those which are not causally redundant. Among the composite objects he admits are human beings. However, he eschews the existence of the vast majority of composite objects, and his argument could be extended to argue for full-blooded nihilism.

***The Causal Overdetermination Argument (COA):***

1. The baseball – if it exists – is causally irrelevant to whether its constituent atoms[144], acting in concert, cause the shattering of the window.
2. The shattering of the window is caused by those atoms, acting in concert.
3. The shattering of the window is not causally overdetermined.
4. Therefore, if the baseball exists, it does not cause the shattering of the window.[145]
5. Therefore, if the baseball exists, it is causally inert; it has no causal powers.

So far, Merricks's argument looks pretty much the same, albeit a little more formally structured, as the argument considered in the previous section. For it leads to the same conclusion that if there were any composite macroscopic objects, like baseballs, then they would be causally inert; they would be mere epiphenomena. And that is because all the causal powers we normally ascribe to (alleged) composite objects, like baseballs, are in fact sufficiently accounted for by those objects' constituent atoms acting in concert.

But it is at this point that Merricks begins to turn the screw. For he claims that if one accepts that inanimate, composite objects have no causal powers, then one cannot accept that such objects even exist, on pain of inconsistency. For if baseballs (and other inanimate, composite objects) *did* exist, then they *would* have causal powers – or so Merricks insists. I.e. they *would* cause windows to shatter; they *would* cause visual, tactile sensations; etc. Thus if one accepts that baseballs exist, *and* one accepts the overdetermination argument, one must also accept the contradictory position that baseballs (and other inanimate, composite objects) simultaneously *do* have causal powers and *don't* have causal powers. Since that is clearly absurd, we must conclude that baseballs *et al* do not exist. So we have an extra premise or two now being added to the argument. Let me make those premises a little more explicit:

6. If baseballs (and by extension, any inanimate composite objects) *did* exist, they *would* have causal powers (i.e. they would cause windows to smash, etc.).

[144] It should be noted that we are not talking about the atoms of modern physics here, but rather, mereologically simple minima, whatever they may be.
[145] Merricks (2001), 56

7. Therefore, from 5 and 6, it follows that *if* there were composite objects like baseballs, they would have to be, simultaneously, causally inert *and* causally efficacious.
8. Nothing can be causally inert *and* causally efficacious.
9. Therefore, there can be no baseballs (and by extension, there can be no inanimate composite objects)

So this is Merricks's argument in full. The argument is valid, but as is normally the case, a number of the premises require some independent support in order for it to be convincing. Firstly, one may well have noticed that there are two key terms involved in the argument which are in need of clarification. The first is 'causal irrelevance', and the second is 'overdetermination'. Let me explain, in a little more detail, what Merricks takes these terms to mean. The *general notion* of causal irrelevance should be fairly obvious and easy to grasp – in certain instances at least. My having sausages for breakfast, for example, is clearly causally irrelevant to whether it rains in the afternoon. It may not be so obvious, however, that an (alleged) entity is causally irrelevant to the actions/effects of its constituent parts acting in concert. Merricks attempts to clarify the issue, and defines causal irrelevance in the following way. An object, *O*, is causally irrelevant to whether some *x*s cause an effect, *E*, if and only if, none of the following four criteria are met:

A) *O* is one of the *x*s.
B) *O* is a partial cause of *E* alongside the *x*s.
C) One or more of the *x*s causes *O* to cause *E.*
D) *O* causes one or more of the *x*s to cause *E.*[146]

Merricks then introduces the following general rule:

> ***Causal Principle.*** Suppose: *O* is an object. The *x*s are objects. *O* is causally irrelevant to whether the *x*s, acting in concert, cause a certain effect, *E*. The *x*s, acting in

[146] Merricks (2001), 57

> concert, do cause *E*, and *E* is not overdetermined. It therefore follows that *O* does not cause *E*.[147]

It is on the basis of this causal principle that Merricks draws his conclusion, 4, from the premises 1 – 3. There are possible concerns one could have with Merricks' causal principle, and his notion of causal irrelevance, but they are at least *prima facie* plausible, so for the time being I shall grant them as true.[148]

Secondly, we have the notion of 'causal overdetermination'. Merricks claims that overdetermination should be understood in "the most literal, straightforward, and natural sense possible".[149] By this, then, we can take it that an event, *E*, is overdetermined if it has two (or more) distinct and independent causes, each of which would have been quite sufficient on its own to cause *E*. This is not to say, of course, that all events with more than one cause are overdetermined. In many cases more than one cause is required for an event to occur (e.g. the depression of the accelerator pedal *and* the spark of the sparkplug are both causes of the car's moving), but in most cases, all the causes are *necessary* for the resulting effect to occur, even though they may not be individually sufficient (e.g. neither the sparkplug *nor* the depression of the accelerator pedal *alone* are sufficient to cause the car to move, but both are necessary). Overdetermination only occurs when there are multiple, *individually sufficient*, causes for an event. It is widely agreed that causal overdetermination is objectionable.[150] That is to say, it is widely agreed that we shouldn't posit entities or events that would imply, or result in, causal overdetermination unless absolutely necessary. Merricks shares this distaste for causal overdetermination. Indeed, it is a crucial step in his argument. For if one were happy to accept that there was widespread causal overdetermination, then the argument would simply not go through. Rather, one could accept that there are composite objects *and* their parts, and that the causal effects of both are systematically overdetermined. I share with Merricks the view that causal overdetermination is highly objectionable, particularly when on such a widespread and

[147] Merricks (2001), 58
[148] In particular, see Thomasson (2006) for arguments to the contrary. I will address Thomasson's responses below, in §4.6.
[149] Merricks (2001), 58
[150] Aside from Merricks, see Bunzl (1979); Loeb (1974); Kim (1993).

ubiquitous scale as is being suggested here. Thus I accept premise 3 of the argument as true.[151]

The final premise that requires some independent support is premise 6: the claim that baseballs, if they *did* exist, *would* have causal powers. This claim is most plausible, to the extent that I don't see how one could reasonably reject it. Firstly, it has the weight of common sense on its side. After all, people *talk* as if composite objects have causal effects all the time. It's common to hear things said such as "the dog knocked over the vase", or "the bomb caused utter devastation". These types of statements are clearly describing *causal* events. But if there *were* composite objects, like dogs and bombs, to which these utterances refer, then surely it would make sense to suppose that they *did* have the causal powers that are so often ascribed to them. Moreover, if a baseball, for instance, were to exist yet not have causal powers, it would not be able to cause light to deflect into one's eyes, thus essentially, one would not be able to see it. But it would surely be absurd to suggest that in addition to the particles, there *was* a baseball which they compose, but that you couldn't see the baseball, but only the particles!

But there is a much stronger reason to believe that 6 is true, and that reason is scientific. Firstly, if there *were* composite objects like baseballs, dogs, and bombs, then they surely must be *physical* objects. They would have physical properties after all, like shape, size, mass, and location. So I don't see how anyone could possibly say that they were *not* physical objects. But if they are physical objects, then they must be subject to the laws of physics. That is, they must be capable of being acted upon by external forces, and likewise, they must act upon other objects. If $x$ is a physical object, then it is always in principle possible to exert a force upon $x$, but Newton told us that for each and every action there is an equal and opposite reaction, so if we can, even in principle, exert a force upon $x$, then $x$ must be able to exert a force back. But to exert a force is to interact *causally*. Causally inert objects cannot exert forces on physical objects. So it seems that composite objects, if they *were* to exist, *could not possibly* be causally inert. Premise 6, then, is true.

To conclude, therefore, Merricks's COA represents a truly compelling case for nihilism, for it highlights some seemingly insurmountable problems with composite objects.

[151] There have been attempts to claim that overdetermination is nothing to worry about. See Sider (2003) and Thomasson (2006). I will address, and reject, these claims in §4.8.

It just does not seem possible that a physical, composite object could be entirely causally inert. Yet if the first stage of the COA is right (and I hope to have shown that it is), then composite objects, were they to exist, *must* be causally inert. This is a worrying contradiction, and it is one that has to be dealt with by the believer in composite objects. For the nihilist, by contrast, the contradiction never arises, and that is why the COA represents such a strong case to believe that nihilism is true.

# §4.6. Potential Responses to COA

## *§4.6.1 Presupposition Failure*

Amie Thomasson has argued that whilst Merrick's overdetermination argument may be valid, it is not sound.[152] Thomasson targets premise 3 of COA, which claims that the shattering of the window is not overdetermined, and suggests that there are good reasons to believe it is not true. Very roughly, she claims that the third premise involves a particular kind of presupposition failure that is significant enough to alter its truth-value. In what follows I will show how Thomasson reaches this conclusion, and also show why I think her conclusion should be resisted.

Thomasson agrees that premise 3 *looks* very plausible. The reason for this, she claims, is that if we were to admit that the shattering of the window *was* overdetermined, then we would have to accept the truth of the following conjunctive statement:

(1) The atoms caused the window to shatter *and* the baseball caused the window to shatter.

Thomasson claims that (1) just immediately strikes us as objectionable, and she admits that there is something "clearly amiss" with it.[153] Moreover, she claims that it is our disinclination to accept (1) that inclines us to accept premise 3. However, and this is where her argument

[152] Thomasson (2006). Later in the same paper, she actually goes on to suggest that it is not valid either. I will consider her reasons for this in the next section.
[153] Thomasson (2006), 344

springs into life, it is not because (1) is *false* that we are disinclined to accept it, she insists, but rather, it is because it conjoins items together in a list in a wholly inappropriate way.

Thomasson claims that when one conjoins items in a list with the term 'and', one presupposes that the items one is conjoining are "separate and independent entities". This is why, we are told, some of Gilbert Ryle's famous examples of category mistakes (such as 'she bought a left-hand glove, a right-hand glove, and a pair of gloves') strike us as strange, or even absurd.[154] The reason why they seem strange, she claims, is that the presupposition of separateness and independence is violated. The pair of gloves is not a separate and independent entity from each of the gloves it consists of, and this is why it seems strange to say that you bought two gloves *and* a pair of gloves. The same is the case, it is supposed, with the baseball and its constituent atoms. They are not independent and separate entities, thus conjoining them in a list of causes can create strange, counter-intuitive, or even absurd, results.

Drawing on previous work on presupposition failure, by the likes of Strawson, Grice, and Stalnaker, Thomasson suggests that cases of presupposition failure like this can have a significant impact on the truth-value of the statements in which the presupposition has failed. Strawson, for instance, claimed that statements involving presupposition failure lacked a truth-value altogether.[155] Grice has claimed that statements involving presupposition failure could actually be true, yet merely misleading.[156] Thomasson points out that if either of these two views are right, then they will have a significant impact on the causal overdetermination argument. Recall, if you will, the key statement that is the object of our current concern:

(1) The atoms caused the window to shatter *and* the baseball caused the window to shatter.

Now it is the perceived falsity of this statement that is responsible for our inclination to accept premise 3 of COA. But say Strawson's analysis of statements involving presupposition failure is correct. In that case, (1) would not be false, but rather, would lack a truth-value. Moreover, its *negation* would also lack a truth-value, since its negation would

[154] Ryle (1949), 16-42
[155] Strawson (1950). Frege (1892) originally put forward this view.
[156] Grice (1989), 5-21

also involve the same presupposition failure. But if that were the case, then it looks like premise 3 would also lack a truth-value, in which case, COA would be rendered unsound. Likewise, if we took the Gricean view that statements involving presupposition were capable of being true, yet merely misleading, we would get a similar result. For if (1) were taken to be true, then even if it were misleading, it seems that premise 3 must be false. For if the shattering of the window was caused by the baseball *and* the atoms, then it surely *was* overdetermined, which would directly contradict premise 3. It is for these reasons that Thomasson suggests that COA may actually be unsound, since the truth of premise 3 is in doubt.[157]

I think there are a number of reasons why one might want to resist Thomasson's argument. Firstly, I am not entirely convinced that when conjoining items in a list one *does* actually presuppose what Thomasson claims one does. She claims that when one conjoins items in a list, it is presupposed that those items are "separate and independent".[158] Let's consider these terms one by one. I think it is fair to say that items on a list *are* commonly presupposed to be separate. And by separate I can only presume that Thomasson means *distinct.* For one cannot make a list of just one item. And if one has a list of multiple items, but those items are not distinct, then that is not a list either, but merely a repetition of one and the same thing. For example, it would be quite misleading, in fact it would surely be quite wrong, if you were to tell someone that you had just bought a new item of clothing *and* a mackintosh *and* an anorak, if in fact you had simply bought a new coat. So I am willing to accept that it is presupposed that conjoined items on a list are *separate, i.e. distinct.* But, of course, the baseball and the atoms *are* separate items. They are quite distinct – they have different properties.[159] So no presupposition has failed there.

The other feature that Thomasson claims is presupposed is *independence.* Now it is not immediately clear to me, or at least, it is certainly not obvious, that we *do* in fact

[157] It is perhaps quite revealing to note that Thomasson does not make any mention at all of what is maybe the most well-known analysis of statements involving presupposition failure – that of Bertrand Russell. Russell (1905) claimed that statements involving presupposition failure should normally be considered false. This analysis, of course, would not give Thomasson the result she requires – quite the opposite, in fact – as it would mean that (1) is indeed false. I won't push this line of response, however, since I don't think the falsity of (1) has anything to do with presupposition failure. Rather, I think it is false simply because things are not as it claims them to be.

[158] Thomasson (2006), 344

[159] Unless, of course, you endorse the thesis that Composition is Identity. However, that is a most controversial thesis, and I am pretty sure that Thomasson does not want to rely on its truth in order for her argument to go through.

presuppose that items on a list should be independent. In fact, the first thing to note is that it is not immediately clear what 'independent' is actually meant to mean. For there are lots of different ways in which one thing can depend upon another. A child can depend on her parents for instance; the course of action one chooses may depend on the circumstances; the colour that an object appears can depend on the lighting conditions; and so on. But I don't think that Thomasson intends the word 'independent' to be taken in any of *these* contexts. I can only presume that she has some kind of *existential* dependence in mind. Specifically, the kind of dependence such that, if *x* depends on *y*, then *x* could not exist *unless y* exists, or something along those lines. For it is reasonable to suppose that this kind of dependence holds between baseballs and their constituent atoms, or indeed, between any whole and its parts. For if there was a baseball composed of atoms, then surely that baseball could not exist *unless* those composing atoms also exist. I.e. if you take away the atoms, you take away the baseball too.

The question is: is it really that obvious that we do presuppose that items on a conjoined list must not stand in the relation of existential dependence just described? I'm really not sure that it is. Firstly, it seems a bit of a stretch of the imagination to suppose that we regularly make presuppositions of quite technical dependence relations obtaining when we engage in the simple task of providing a list. Secondly, there are surely instances where it seems perfectly reasonable to conjoin items in a list that *are* existentially dependent in the way described, and moreover, that do stand in the relation of part and whole. If someone were to tell you that there were three trees in the field and thousands of leaves, for example, that doesn't sound in the least bit peculiar or misleading, yet the leaves are surely parts of the trees, and they surely couldn't have existed *without* the trees. Of course, one can just as easily come up with examples that *do* sound peculiar (e.g. "what an astounding catch! Not only did he catch the baseball, but he caught the atoms arranged baseball-wise as well!"), but I am merely trying to impress two important points: one is that it is not obvious that we do actually presuppose the things Thomasson claims we do, and the second is that even if we do make such presuppositions, it is not always clear that they fail in such a way as to make the statement misleading at all.

There is a more compelling reason to be suspicious of Thomasson's argument, however. It is that the *type* of presupposition failure she highlights is very different to the

type of presupposition failure that was the concern of the likes of Strawson *et al*. Strawson, like Frege before him, was concerned specifically with *existential* presupposition failure. This occurs when a statement presupposes that something exists, when in fact it does not. More often than not, these types of statement take the form "the *a* is *F*", where there is actually no referent of *a.* My concern is that this is a much more significant type of presupposition failure than that which is the concern of Thomasson, and as such, it may not be legitimate to suppose that any conclusions one draws about the former will automatically apply to the latter. The most famous example of a statement involving existential presupposition failure must be:

(2) The present King of France is bald.

Since it is the case that there is no present King of France, what are we to make of this statement? Is it true or False? This is where philosophers have been divided. Some think that the sentence is not apt for truth, and thus is neither true nor false.[160] Others think that it is plain false.[161] (I am unaware of anyone who believes sentences like (2) to be true). I am not overly concerned, at this juncture at least, to establish which of these views, if any, is correct. What I do want to emphasise, however, is what I see as a stark disanalogy between a sentence like (2) and a sentence like (1). The presupposition failure in (2) is clearly going to have a significant effect on the content, and thus the truth-value, of the sentence in which it occurs. As Stephen Yablo has said, this kind of presupposition has "failed catastrophically".[162] For in uttering a sentence like (2), one is presupposing the existence of something that is non-existent; one is trying to describe the qualitative features of an object that doesn't exist. So it is no wonder that one runs into trouble!

But the presupposition failure that Thomasson highlights – if indeed there is any such failure at all – is not nearly so acute; it is far more subtle. For in uttering a sentence like (1), one is not presupposing the evident falsity that something non-existent exists, but rather, one is merely presupposing that the items referred to in each conjunct of the sentence are not linked by some specific relation of existential dependence. But what is the real effect of such

[160] Frege (1892); Strawson (1950)
[161] Russell (1905)
[162] Yablo (2006, 165; 2009, 515)

presupposition failure? It is surely not *catastrophic* in the way that it is to presume that there is a present King of France and then go on to describe what he is like. Perhaps we need a clearer idea of what it is for a presupposition failure to be catastrophic or not in order to assess the full extent of the damage – if any – caused by the presupposition failure in (1). Stephen Yablo explains the distinction nicely:

> Failure is catastrophic if it puts a thing out of commission, renders it unable to perform its primary task. Non-catastrophic presupposition failure occurs when a sentence's false presuppositions do not prevent it from performing its primary task. [...] The primary task of an indicative mood sentence is to make a claim [...] that is true or false according to whether matters are as described.[163]

I think it is fairly clear that the kind of presupposition failure that Thomasson claims to be involved in (1) is *not* catastrophic. It does not render the sentence unable to perform its primary task. That is, (1) surely still makes an identifiable claim that one can evaluate. Even if one recognises that (1) presupposes a falsehood (i.e. the false claim that one can conjoin existentially dependent items in a list), it is just not plausible for one to insist, 'I have no way of identifying what claim (1) is making about the world, and indeed, whether matters are as it claims them to be'. Such an assertion just doesn't wash. *Of course* it makes a claim about the world! It claims that the baseball *and* the atoms caused the window to smash. And this is perfectly evaluable for truth or falsity. Just *how* one evaluates it will depend on whether one believes that the shattering of the window was caused by both the atoms *and* the baseball or not.

The upshot of these remarks is, I hope, that the presupposition failure that Thomasson highlights is really of very little significance. Firstly, it is questionable whether we do actually presuppose the things she claims we do. Secondly, if we *do* so presuppose, it is questionable as to whether *what* we presuppose is false. Lastly, and most importantly, even if it is accepted that (1) does involve a presupposition failure, it clearly isn't a catastrophic failure, for it is still quite easy to identify the claim the sentence is making: that the baseball *and* the atoms cause the shattering of the window. At base, then, it seems that Thomasson has not really provided a counter-argument to COA, but rather, she just

[163] Yablo (2006), 165

disagrees about the truth of premise (3). That, in itself, is fine. But in the absence of any reason to suppose that premise (3) *is* false, it is not going to convince the proponent of COA. For the very driving force behind COA is that the shattering of the window is *not* overdetermined, and thus that premise (3) is true. I hope to have shown that the presupposition failure that Thomasson highlights does not provide any such reason, and therefore, I think her argument should be resisted.

### *§4.6.2 What's So Bad About Overdetermination Anyway?*

A few philosophers have claimed that we should simply accept that macroscopic events like the shattering of windows *are* causally overdetermined by objects and their parts. Coupled with this claim, however, is the insistence that the type of overdetermination involved is really nothing to worry about.[164] It is, in Thomasson's words, not "*real* overdetermination".[165] The reason for this claim appeals to the fact that there is some kind of inextricable link between objects (like baseballs) on one hand, and their parts (i.e. atoms arranged baseball-wise) on the other. Sider, for instance, says there is an "intimate relationship" between the two.[166] Thomasson, meanwhile, claims that there is a certain "independence [which] seems to be lacking between the causal claims of the baseball and the atoms arranged baseball-wise".[167]

Now I concede that there is, of course, a certain sense in which a baseball and its constituent atoms are not completely independent. One could not throw a baseball at a window, for instance, without also throwing its constituent atoms. The atoms and the baseball seem to follow one another around, so to speak.[168] But one should not forget that they are still *distinct* entities. Or at least, providing composition is not identity, they are distinct entities, and given what I have said about that in the last chapter, I will assume that composition is *not* identity. So the question is, even if one concedes that there is a certain dependence that holds between objects and their parts, given that they are distinct objects, why should we be unconcerned if wholes and their parts systematically overdetermine

[164] Thomasson (2006); Sider (2003)
[165] Thomasson (2006), 349
[166] Sider (2003), 1
[167] Thomasson (2006), 347
[168] Although see Saucedo (2011) for an alternative view.

events? Let me try and unravel the reasoning behind such a claim. In so doing, I hope to expose it as false.

Causal overdetermination, and here I mean *real* causal overdetermination, can be defined as follows:

**Causal Overdetermination:** An event, *E*, is causally overdetermined iff it has two (or more) *independently sufficient* causes (or causal stories).

This is, I think, a fairly standard definition. Moreover, it is a definition that Thomasson explicitly endorses, and one which Sider seems happy to accept. So in accepting this definition I cannot be accused of adopting a biased starting point against my opponents. Now we have a set definition in place, let us return to our example case of a baseball shattering a window. If we are to conclude that this event is not overdetermined (in the *real* sense), then we must accept that it does not have more than one *independently sufficient* cause (or causal story). So the thing we really need to establish is what the cause (or causes) of the shattering *is*. Both Thomasson and Sider seem willing to accept that both the baseball *and* the atoms arranged baseball wise *do* cause the shattering of the window. If this is right, then we have two causal stories to explain the shattering:

(A): The baseball caused the window to shatter
(B): The atoms arranged baseball-wise caused the window to shatter

So, given our definition of *real* causal overdetermination, the only way that Thomasson and Sider can claim that the shattering of the window is *not* overdetermined (in the *real* sense), is for them to show that A and B are not both *independently sufficient* causal stories. But they surely *are* both independently sufficient! To see why, let us consider them one by one.

Firstly, take B. It is surely beyond doubt that atoms arranged baseball-wise are independently sufficient to cause a window to shatter. Let us suppose that nihilism was true, for instance. In that case, atoms arranged baseball-wise would not compose a baseball; there would be no baseballs. But one could not seriously claim that those atoms would be causally incapable of shattering a window. And one does not have to accept that nihilism is

true to accept this conclusion. Rather, one merely has to accept that *if* it were true, then atoms arranged baseball-wise, acting in concert, would be sufficient to cause a window to shatter. And since no baseballs would be involved in such an event, their causal sufficiency would be entirely independent.

So if Sider and Thomasson's view is to be vindicated, it must be A that is taken to be not an independently sufficient causal story. But what could one mean by saying that a baseball is not independently sufficient to cause a window to shatter? One possibility would be to say that the baseball's causing the window to shatter is the *very same event* as the atoms's causing the window to shatter. But this seems to indicate that there are not two causal stories behind the event at all. Rather, there is a *single* causal story that has been described in two different ways. Perhaps this *is* what is being implied? But if it is, I don't see how it would avoid the problems. Because it just immediately invites the question: *but which causal story is the right one? I.e. what causes the window to shatter – the baseball or the atoms?* Now if one endorses the existence of the baseball and the atoms, and one also takes both to be causally efficacious (as Thomasson and Sider do), then the answer here would surely have to be: *both.* (For to admit it was *only* the atoms which caused the shattering, or *only* the baseball, would be to admit that only one of them was causally efficacious). But then we are back to where we started! For if the baseball and its constituent atoms are distinct entities, and they both cause the window to shatter, then we *do* have two distinct causal stories behind the shattering. And this then re-invites the original question of why the baseball is not independently sufficient to cause the window to shatter.

The only way I can make sense of the response under consideration would be if one was taking composition to be identity. For if the baseball *just is* the atoms arranged baseball-wise, then it is clear why it is not independently sufficient to cause the window to shatter. For it *just is* its constituent atoms, so it would make no sense of talking of the baseball *apart* from those atoms. But I am assuming that composition is *not* identity, and, moreover, that Thomasson and Sider would not want to rely on the fact that composition *is* identity in order for their argument to go through. But if composition is not identity, then the baseball is a distinct, physical object in its own right. It has physical properties. It has a shape, a size, a mass, and a density. In fact, it has all the right physical properties an object would need to be capable of shattering a window. But if such an object exists, then I need to be told why

such an object is not independently sufficient to shatter a window! But these arguments do not tell me why this is. Indeed they appear to have all the merits of theft as opposed to honest toil. For on the one hand I am told that atoms arranged baseball-wise really do compose a physical object; an object with just the right kind of physical properties required to smash windows. But on the other hand I am told that such an object is not independently capable of smashing windows! Something has gone wrong here, and that is why I do not accept the argument. If both the atoms and the baseball exist, *and* they are causally efficacious, then the shattering of the window is overdetermined in a very real sense. There is no other plausible option.

## §4.7. Dispensing with Age-Old Problems

Ordinary objects are so called because they are taken to be just that: *ordinary.* What could be more ordinary than a tree or a rock, for instance, or a house or a sock? Natural or man-made, ordinary objects are instantly recognisable, easily describable, and seemingly unproblematic, or so it might seem reasonable to think. But a little bit of philosophical reflection casts doubt on all that. If one tries to establish precisely what makes an ordinary object an object, what its identity conditions are, or how it persists and changes through time, one will inevitably run into serious trouble. We have just seen how ordinary objects are *causally* problematic, in that they either force one into accepting widespread causal overdetermination, or they must be seen as being mysterious, epiphenomenal entities. But this is not the only way in which ordinary objects can be problematic. There are a number of long-standing puzzles about ordinary objects that, despite having been subjected to intense scrutiny over the years, remain unsolved. These puzzles, as I will explain below, make ordinary objects seem problematic to the point of being paradoxical. In fact, I think it is fair to say that the term 'ordinary object' is a complete misnomer. A little reflection shows that such objects are anything but ordinary. Nihilism, however, avoids all these problems, since it denies that there are any 'ordinary' objects at all. This, I think, should count significantly in nihilism's favour.

### *§4.7.1 The Statue and the Lump*

Let's suppose that you decide to make a statue. First of all, you get yourself a lump of clay. Let's call it 'Lump'. Then you carefully go about sculpting the clay into your shape of choice, say, a horse-shape. Let's call the finished statue 'Horse'. Now suppose that on your way home, whilst clutching the statue, I was to ask you how many objects you were holding. What would the right response be? The natural response, I imagine, would be 'one' – you are holding the statue, Horse. But what about Lump? You are holding Lump as well. So are there not two objects you are holding – Horse and Lump? Now the natural response *here* would be, in addition to a certain exasperation, that they are not two different objects; Horse *is* Lump; they are one and the same thing. So you are only holding one object, but that object has two names – Horse and Lump. But herein lies the puzzle. We know that if Horse and Lump are the same thing (i.e. identical) then they must share all the same properties – Leibniz told us that. But on closer inspection, it seems that Horse and Lump do *not* share all the same properties. In fact, it seems as though they differ in a number of very important respects. For a start, Lump existed well before Horse existed. This immediately suggests that they can't be the same object. For it is surely absurd to suggest that something could have existed long before it existed! Secondly, they differ in terms of their persistence conditions, for if you squashed the Statue into a shapeless ball, then it seems Horse would be destroyed, whereas Lump would survive.

So what, *precisely*, is the puzzle? It is that the story of Lump and Horse leaves one on the horns of a particularly prickly dilemma. On the one hand you could accept that Horse and Lump are distinct things, since they have different properties, but that would mean that you have two distinct physical objects occupying exactly the same space at exactly the same time, which sounds all wrong. On the other hand, however, you could maintain that they are one and the same object, only to then face the somewhat galling consequence of denying Leibniz's Law, and asserting that $a$ and $b$ can be identical *even though* they do not share all the same properties. Neither option is attractive. In fact, both options are positively abhorrent. And that is precisely why the case of the statue and the lump presents such a problem.

The puzzle of the statue and the Lump is not insurmountable; there have been a number of different solutions proposed by a number of different philosophers. One popular

response, for instance, is to allow that distinct objects *can,* in fact, occupy the same space at the same time.[169] Very roughly, on this view, the lump and the statue are said to stand in the relation of *constitution.* The lump *constitutes* the statue, but is not *identical* to it, and they both occupy the same space at the same time. Another common response is that objects are actually four-dimensional (i.e. extended through the three dimensions of space *and* the fourth dimension of time), and composed of *temporal parts,* as well as *spatial parts.*[170] According to this view, then, Lump and Horse are two distinct four-dimensional objects which share a certain temporal part. So once Lump has been sculpted into Horse, there are not two distinct physical objects in the same place at the same time, but rather, a single object – a temporal part – which belongs to two distinct four-dimensional objects: Lump and Horse. Lump and Horse *temporally overlap,* so to speak. Yet a third solution to the problem has been offered by Peter Geach, who claims that it is our (mistaken) notion of identity which is the cause of the puzzle.[171] Geach claims that identity is never *absolute,* but rather, it is *relative to a kind.* So we can't say that Horse and Lump are *identical per se* – in fact, we can't say that any things are identical *per se.* But we can say that Horse is the same bit of clay as Lump, or the same statue as Lump, and so on.

So there *are* possible solutions to the puzzle. But they are all, I would suggest, theoretically costly. For whilst they may overcome the puzzle, they all represent significant crosses for their proponents to bear. The nihilist, however, has a much easier option. Indeed, for the nihilist, the puzzle never even arises. For according to the nihilist, there are no such objects as Lump or Horse.[172] So it is wrong to say that there are two objects occupying the same space at the same time, and it is wrong to say that these objects have different, incompatible properties. There are no such objects in the first place. There may be simples arranged Lump-wise, for instance. And those simples may then be arranged Horse-wise, and then later re-arranged Lump-wise once more. Or there may be certain Lump-shaped 'thickenings' in the fabric of the world, which become Horse-shaped and then revert to being Lump-shaped. But there is no problem in that; no new objects have come into, or gone

[169] Perhaps the most well-known defender of this view is Wiggins (1968a; 1968b). But see also Shoemaker (1999); Simons (1987), 237-240; among many others.
[170] See, for instance, Lewis (1986a), 202-204; Heller (1990); Sider (2001); among others.
[171] Geach (1962; 1967)
[172] Unless, of course, you were a nihilist who took Lump and/or Horse to be mereological simples. But this would be a rather strange and unmotivated view, so I will ignore it here.

out of, existence. So the nihilist bypasses the whole problem easily and elegantly.[173] She who endorses composite objects, however, must bear a significant cross. True, she has a choice of crosses to choose from – either the ubiquitous collocation of objects, the endorsement of temporal parts, or a relativised notion of identity – but she has to choose one all the same. This represents yet another reason to favour nihilism over universalism.

## §4.8. The Neutrality of Nihilism

It is often assumed that if nihilism is true, and all material objects are mereological simples, then those simples must be very large in number, and very small in size. These tiny little simples are often identified with the sub-atomic particles of physics, or even extensionless points of spacetime.[174] But nothing about the thesis of nihilism *entails* that this is the case. As I have already pointed out in §4.1., nihilism is a very broad thesis, and it is perfectly neutral over what *size* material objects are, so long as those objects are mereologically simple. Thus it is perfectly *consistent* with nihilism to suppose that simples are microscopic or point-sized, but it is also perfectly consistent to suppose that they are much bigger, say, $1km^3$. Indeed it is perfectly consistent with nihilism to say that there is only one, *maximally large*, simple in existence – the world itself. (This latter thesis is, of course, the thesis of monism which is the central topic of this dissertation). So there are many different variants of nihilism; there are many different *ways* that simples could be. But this is often ignored. Indeed, in the next chapter I introduce what I call *the Nihilistic Fallacy*, to expose the evident, and unjustified, bias in the literature towards those nihilisms which posit simples that are many and small. But for now I want to point out that nihilism, in its broadest sense, is quite neutral as to which of these variants is the correct one.

What is also important to note is that all the arguments that have so far been given for nihilism argue only for nihilism in this broadest sense, and they lend no more

[173] And the same goes for other puzzles of constitution. There is no puzzle as to which ship is the Ship of Theseus, for instance, if there are no ships at all. And there should be no concern as to whether Tib existed at the same time as Tibbles, for instance, if neither of them *ever* existed. And so on.

[174] Sider (forthcoming); Cameron (2010b); and Dorr (2005) all seem to favour the view according to which simples are numerous and tiny, although none explicitly state whether they take them to be point-sized or not. There seems to be a general *gesture* towards the view that simples *are* point-sized, but again, this is not made entirely explicit.

plausibility to one variant of nihilism than they do to another. In some cases this is quite obvious. The argument from elimination, for instance, says only that universalism and restrictivism should be rejected, thus leaving nihilism as the only option. It is quite clear that *this* argument gives no indication of what variant of nihilism should be endorsed, but only that *some* variant of nihilism should be. But it is not so clear in all cases. In particular, the main argument I have given for nihilism – the causal overdetermination argument – *does* seem to point towards a nihilism of microscopic particles. After all, it says that we should not posit baseballs because the particles arranged baseball-wise do all the causal work on their own. But that surely entails that the particles must exist in the first place! So it might reasonably be thought that the causal overdetermination argument *does*, in fact, indicate that a specific variant of nihilism is the correct one: a nihilism that posits a plurality of tiny little simples.

I can see why someone might take this view, for the way in which the argument is set out seems to point towards it. Nonetheless, it is wrong. The causal overdetermination argument is an argument for nihilism in its broadest possible sense. It lends no more support to one variant of nihilism than it does to any other. In the next section I will show exactly why that is.

### *§4.8.1 The Causal Overdetermination Argument Revisited*

The reason why the causal overdetermination argument *seems* to point towards a nihilistic ontology of microscopic, or point-sized, particles is solely down to the somewhat misleading way in which Merricks sets it out. His set up of the argument rests on a number of key assumptions that give it something of an unfair bias. At first glance, however, these assumptions seem perfectly reasonable, to the extent that one may not have even noticed that they were being made at all. If we restate the argument, however, and make these assumptions clear and explicit, the unjustified bias of the argument soon becomes apparent. To begin, reconsider the first premise of the argument:

1) The baseball – if it exists – is causally irrelevant to whether its constituent atoms, acting in concert, cause the shattering of the window.

The implicit assumption in this premise is that it is the *baseball* whose existence is in doubt. The existence of the atoms-arranged-baseball-wise is simply taken for granted. As already mentioned, this assumption may seem perfectly reasonable at first glance, particularly in light of the alternatives. For the alternative would be to accept the converse: to take the baseball's existence for granted and bring into question the existence of its (alleged) constituent atoms. I can anticipate a raised eyebrow at this point, for it would be unusual, to say the very least, to suggest that the baseball could exist whilst its (alleged) constituent atoms did not. However, the mere peculiarity of this supposition should not be reason enough to reject it out of hand. This is particularly so, since one should remember that Merricks presents his causal overdetermination argument as direct support for his thesis of *eliminitavism.* This thesis, let us not forget, states that inanimate, macrophysical objects (like baseballs) do not exist, and that the only things which *do* exist are microphysical simples.[175] Surely, then, we should be wary of any supporting argument whose premises seem so biased in favour of the microphysical and against the macrophysical. For any argument to be convincing, it must surely begin on neutral ground with regard to the central tenets of the theory it is trying to prove. In this case, the clear bias in favour of the microphysical would suggest that the starting ground is far from neutral. This bias continues in premise two:

2) The shattering of the window is caused by those atoms, acting in concert.

The assumption is no longer implicit here. It is baldly asserted that it is the microphysical atoms-arranged-baseball-wise which do all the causal work in shattering the window. Since it is the atoms which are causally responsible for the shattering of the window, we are told, then if there were a baseball as well, which also caused the shattering, then that shattering would be overdetermined. As a result, we are invited to conclude that there must be no baseballs. The bias in favour of the microphysical over the macrophysical is plain to see. The question, then, is this: *is such a bias justified?* For if it is not, Merricks' argument runs the risk

[175] It should be noted that Merricks does accept that *some* composite entities exist, but not *inanimate* ones, only those which are not causally redundant. This acceptance has no impact on the point in question.

of begging the question. In order to answer this question, we should start by considering the alternatives.

Firstly, let's consider the converse. Let us rebalance the scales in favour of the *macro*-physical. In so doing, we could restate the argument as follows (let us call it COA*):

1) The atoms-arranged-baseball-wise – if they exist – are causally irrelevant to whether the baseball causes the shattering of the window.
2) The shattering of the window is caused by the baseball.
3) The shattering of the window is not overdetermined.
4) Therefore, if the atoms-arranged-baseball-wise exist, they do not cause the shattering of the window.

We can then add Merricks' subsequent reasoning, but in reversed form, and say that if the atoms-arranged-baseball-wise *did* exist, then they *would* have causal powers; they *would*, acting in concert, cause windows to shatter and cause visual and tactile sensations, and so on and so forth. Therefore, since we have shown that the shattering of the window is not caused by the atoms-arranged-baseball-wise, we can conclude that there are no such atoms.

What are we to make of this argument? Let us first put to one side any intuitions we may have about the truth of the conclusion, or indeed of any the premises. Instead let us consider the argument's logical form. COA* is, clearly, of precisely the same form as Merricks' original argument. Therefore if one accepts Merricks' original argument as valid, then one must also accept that COA* is valid. So if one wants to raise concerns about COA*, which I presume Merricks would, one could not attack its validity; one must take an alternative critical route.

The obvious route, and indeed the one which Merricks (albeit briefly) explores, is to judge the arguments on the plausibility of their conclusions.[176] The conclusion of COA is that only microphysical simples exist, and that baseballs and other inanimate composites do not. This may be seen as a somewhat implausible conclusion, particularly when contrasted with the dictates of common sense. The conclusion of COA*, however, is *far more* implausible. For COA* suggests that there are no microphysical atoms, and even more strangely, that

[176] Merricks (2001), 63

baseballs have no proper parts. COA* would have us believe that baseballs are simples! Merricks claims that this conclusion is sufficiently absurd for it to be rejected out of hand: "The claim that atoms arranged baseball-wise fail to compose a baseball might be hard to swallow. But it goes down like draught Guinness compared to the claim that baseballs are simples".[177]

At first glance, it seems that Merricks has a point. For it would be a brave philosopher indeed who was willing to reject the existence of atoms in favour of baseball-sized simples. However, this seemingly absurd conclusion is somewhat superficial; it can be explained away. For the conclusion results, I would argue, not from giving equal existential credence to both the micro- and macro-physical, but rather, from the specific and arbitrary nature of this particular example. Let me explain. When we do ontology, we aim to determine what exists; we aim to discover the basic building blocks of reality; the fundamental existents of the world. Now microphysical atoms – the fundamental particles of physics – seem to be a quite plausible candidate for these existents. Baseballs, in contrast, do not. Philosophers have been known to argue for some peculiar theses, but none, to my knowledge, have ever claimed that baseballs are among the fundamental constituents of reality. And I have no intention of making such a claim here.

As shown above, Merricks' argument is valid. But it has also been shown that it remains valid whether one's bias is towards the microphysical *or* the macrophysical. That is to say, COA* and COA are *both* valid. Now of course, when the argument is framed in terms of baseballs and their constituent atoms, one will inevitably be drawn towards the conclusion of COA rather than COA*. That is because the very nature of the example leaves one with the choice of either atoms or baseballs as being mereological simples. And as I have already said, it would be a brave (or perhaps crazy) philosopher who opted for baseballs over atoms. But the point is that once the argument has been properly unpacked, it is clear that the baseballs and atoms are entirely irrelevant. What is important is the logical form of the argument, and the implications it has for mereology *in general*. For the argument to work, it needn't necessarily be about baseballs, atoms and windows. To illustrate this, what follows is a restatement of the argument in *general terms* (let's call it $COA^{GEN}$):

[177] Merricks (2001), 63

1) X – if it exists – is causally irrelevant to whether the *y*s, acting in concert, cause event, E.
2) Event E is caused by those *y*s, acting in concert.
3) Event E is not overdetermined.
4) Therefore, if X exists, it does not cause event E.

As far as I can see, these variables (X, *y*, E) can be substituted for *any* entities/events, and the argument would remain valid, provided we recognize *one* crucial stipulation. That stipulation is that the *y*s must represent the (alleged) constituent parts of X. If one were to substitute, for X and the *y*s, entities which did not stand in the part-whole relation, the argument would still be logically valid, but it would express only trivial platitudes and would shed no light whatsoever on the question of mereological composition.[178] With that stipulation in mind, one can substitute the variables in any way one likes without affecting the argument's force. For example, we could say:

X = the baseball
The *y*s = two arbitrary halves of the baseball (e.g. the left and the right half)
E = the smashing of the window

This example would give the same result as before, in that it would have us conclude that the baseball is causally redundant. But it would seem a little odd, I would suggest, to eschew the existence of baseballs whilst allowing that baseball-halves exist. In fact, in a battle between baseballs and their halves, it is surely the baseball which is a more plausible candidate for doing the causal work. Either way, what I am meaning to impress here is that the nature of the specific example which one plugs in to the argument will have a big effect on the plausibility of the conclusions.

Now we have also seen that the validity and force of the argument remains unaffected if we reverse the position of X and the *y*s. That is to say, that if you accept COA

[178] It is interesting to note that Merricks himself provides an example that does exactly that, whereby he substitutes himself for X, an unruly mob of vandals (of which he is no part) for the *y*s, and the vandalism of a park for E. See Merricks (2001), 57.

as valid, you must also accept COA* as valid. So just as we have done above, we can now restate COA* in more general terms (lets call it COA*$^{GEN}$):

1) The *y*s – if they exist – are causally irrelevant to whether X causes event E.
2) Event E is caused by X.
3) Event E is not overdetermined.
4) Therefore, if the *y*s exist, they do not cause event E.

Just as before, we can now substitute anything we like for those variables - provided that the substitutes for the *y*s and X stand in the part-whole relation – and the argument still goes through. The only difference being that it is now the *y*s, not X, which come out as being causally, thus existentially, redundant. And just as I noted above, I think one could come up with some specific variables to plug in to the argument that do *not* lend any more credence to the microphysical than they do to the macrophysical. For instance:

X = The world itself
The *ys* = The world's microphysical parts
E = The evolution of the world

If we were to plug these variables in to the argument we would be left with a choice. Either the world should be seen as causally redundant, and thus dismissed, or the microphysical particles should be seen as being causally redundant, and thus dismissed. Which choice to make, however, is not at all obvious. Both the world itself *and* fundamental microphysical atoms, seem like fairly plausible candidates for being fundamental existents. In Merricks' example, it is easy to eschew the existence of the baseball, but in this example it is not so easy to eschew the whole world, or so it seems to me. To choose between the world and its parts, one would surely need some further, independent arguments. But in the absence of such arguments, the wise move would surely be to reserve judgement.[179]

[179] In the following chapter I will provide some arguments to show that we should favour the world over its parts, but that is for later.

So what does this all show? First and foremost, it shows that the argument has nothing to do with baseballs. The argument is, in fact, of much wider scope, and indeed much greater interest than that. It is an argument against *parthood in general*. It shows that if one accepts composite entities (i.e. entities with proper parts), then one has no choice but to accept widespread and systematic causal overdetermination. This is because the causal powers we ascribe to the object can be sufficiently accounted for by the causal powers of its parts, *and vice-versa*. Given the objectionable nature of such overdetermination, the argument gives us good reason to believe that there are no composite objects. What we have, then, is an argument for nihilism, and a compelling argument at that. However, once unpacked, and stated in general terms, it is clear that it is an argument for nihilism *in its broadest sense.* That is to say, the overdetermination argument gives us no reason to favour one variant of nihilism over another; it gives us no reason to favour a nihilistic ontology of microscopic particles over a monistic nihilism that posits only the world, or indeed any other form of nihilism. What it *does* do is give us good reason to reject the notion of parthood.

The only way one could use the overdetermination argument to support a nihilistic ontology of microphysical particles would be to take an initial bias towards the microphysical. I.e. to opt for COA rather than COA*. But this is to beg the question! One cannot just pick and choose one's premises in order to arrive at one's preferred conclusion. It must be left as an open question as to whether COA is preferable to COA* unless one can provide *independent* argument to support one over the other.

Merricks does, albeit briefly, recognize this possible reading of his overdetermination argument. As a pre-emptive response he offers a reason why it should be rejected, and why the microphysical atoms *should* be given existential priority over the macrophysical objects they allegedly compose, i.e. why we should accept COA rather than COA*. He claims that "there will be some things that the atoms seem to cause for which the baseball cannot account".[180] To illustrate this, he asks us to consider the activity of the atoms *before* they were arranged baseball-wise. At one time, we are asked to presume, those very atoms were floating free in the early universe, independently and no doubt incontiguously. Yet they still would have had causal powers back then; they would, we can assume, have causally affected any other atoms they may have collided with. But these causal effects cannot be

[180] Merricks (2001), 63

accounted for by the baseball! To suggest so would be absurd, Merricks continues. But the converse does not hold, however. According to Merricks, *all* effects that are allegedly caused by a baseball *can* be accounted for by the causal powers of the atoms which allegedly compose that baseball. Thus there is an *asymmetry* involved. The causal powers of the atoms (if they exist) can sufficiently account for *all* the effects of the baseball (if it exists), but the causal powers of the baseball *cannot* sufficiently account for *all* the effects of its constituent atoms. Thus Merricks concludes that "this asymmetry gives us strong reason, when forced to choose, to favour the causal powers of the atoms over those of the baseball".[181]

On the face of it, the 'asymmetry argument', if I may so label it, seems to provide quite a compelling reason to give existential priority to the microphysical over the macrophysical. For it would be rightly absurd to suggest that a baseball could sufficiently account for the causal efficacy of its constituent atoms long before it had even existed. But once more, I argue, the force of this argument is entirely superficial, and stems from the arbitrary nature of the specific example. If one were to consider the case of *the world* and its microphysical parts, rather than the case of a baseball and its microphysical parts, then the argument clearly doesn't go through. For the atoms did not exist before the world existed! Thus it is equally plausible to say that the causal powers of the entire world can account for those of its alleged parts, as it is to say that the causal powers of the atoms can account for those of the alleged whole that they make up. And, moreover, it would be equally one-sided and biased if a monistically inclined thinker were to say that since the world existed long before any alleged baseball did, it is monism that must be true, because the causal powers of the world can fully account for those of the baseball, whereas the causal powers of the baseball cannot account for those of the whole world.

In light of this, it should be clear that Merricks' asymmetry argument just doesn't wash - it is unjustifiably biased towards a microphysical nihilism. It assumes the existence of microphysical atoms, and then tells a story of how they existed long before someone (or, more correctly, some atoms arranged person-wise) arranged them baseball-wise. But this will do nothing to convince the monistic nihilist, for she could just as plausibly assume the existence of the entire world, and tell a story of how it existed long before it developed any baseball-shaped thickenings. So given this analysis, I claim that whilst the

[181] Merricks (2001), 63

overdetermination argument is a strong and compelling argument in favour of nihilism in general, it lends no support whatsoever towards one variant of nihilism over another. Rather, it lends support to the more general claim that there is no mereological complexity to the world, i.e. that nihilism is true. To choose between *variants* of nihilism we will need further, independent arguments. But that is for the next chapter.

## §4.9. Concluding Remarks

In this chapter I hope to have shown that there are a number of reasons to believe nihilism is true. The first crucial step was to overcome the perennial thorn in the nihilist's side: the common sense objection. But I have shown that nihilism does *not* conflict with common sense, or at least, not in the objectionable way that it is often assumed. Since the nihilist has a number of ways to explain the appearance of composite objects, her assertion that there are no such objects is no longer at all difficult to swallow. The second step was to show that composite objects, were they to exist, would be hugely problematic entities. Either they would be mysterious epiphenomena, or they would be the source of systematic and widespread causal overdetermination. Nihilism, it was shown, avoids all these problems, which should count as a compelling reason to believe it is true. Finally, and importantly, I hope to have impressed that nihilism comes in many guises. At this stage in the dialectic, none of the arguments considered provided reason to favour any one of those guises over another. In what follows, however, I will show why it is *monism* that should prevail.

# Part II

# Monism

# 5. Why Monism?

So far I have put together a case for believing that compositional nihilism is true; a case that I hope the reader has found compelling. But I have also shown that nihilism in its basic form is a very broad thesis. There are different variants of nihilism, with the only thing upon which they agree being that all material objects are mereologically simple. Specifically, I have identified a tripartite division between nihilisms. At one extreme there is punctal nihilism which supposes material objects to be minimally small (i.e. point-sized) and plentiful; and at the other there is monistic nihilism (or just 'monism') which takes there to be a single material object – maximally large yet lonely. In between lies a multitude of views which posit material simples which are neither maximally large nor minimally small. Any of these views will be subsumed under the heading 'intermediate nihilism'. The arguments I have presented so far have been in support of nihilism *in general*, but have been quite neutral as to which variant of nihilism should be preferred. In this chapter I will present a number of arguments to show that of the three variants identified, it is monism which should emerge as the front runner.[182] Monism, I claim, is the best form of nihilism.

## §5.1. The Nihilistic Fallacy

I think it is fair to say that the overwhelmingly common assumption among contemporary philosophers is that *if* nihilism is true, then it is *punctal* nihilism that is true. It seems often to be simply assumed that if all material objects are mereologically simple, then they must also be minimally small. Markosian, for instance, says that on a "natural way of thinking [...] simples are point-sized objects", and elsewhere McDaniel has called this "the traditional

[182] I should note at this point that Jonathan Schaffer has, fairly recently, published a paper with the very same goal as this chapter – to show that "nihilism culminates in monism" (Schaffer (2007), 175). Shaffer's paper is a very good one and has greatly inspired this chapter. In particular, the arguments I present in §5.3. and §5.5. both appear in Schaffer's paper. In both cases, however, I think I have expanded upon, and hopefully improved, the arguments in question. Coupled with the other arguments I will present in this chapter, I hope to have produced an even more compelling case for monism than that presented by Schaffer. (To my knowledge, no-one else has argued from compositional nihilism to monism).

view of simples".[183] I should also note, however, that this assumption is not always made entirely explicit in the extant literature. Many defenders of nihilism are somewhat vague about the precise nature of the simples they posit, other than to maintain that they are material objects that lack proper parts. But the implication that material simples must be minimally small remains evident. For instance, most of those who defend nihilism explicitly employ the van Inwagen-style strategy of paraphrase. That is, they claim that talk of composite objects (like tables and chairs) can be paraphrased into talk of simples arranged in certain ways (table-wise and chair-wise).[184] But, of course, implicit in the use of this strategy is the belief that there are *many* simples, and indeed that they are very small (or at the very least, that they are much smaller than the chairs and tables that they are alleged to be arranged like). It is true that this strategy does not entail that the simples being referred to are *point-sized,* but it does entail a nihilistic view on which simples are taken to be small and numerous; what one might call a generally *Democritean* nihilism.

This assumption is common even among those who reject nihilism. Indeed, it is so common that the fact that there are other varieties of nihilism even available is often entirely overlooked. To illustrate this, consider the following quote:

> Nihilism is the view that there are no composite objects (i.e. objects with proper parts); there are only mereological simples (i.e. objects with no proper parts). The nihilist thus denies the existence of statues, ships, humans, and all other macroscopic material objects. On this view, there are only atoms and the void.[185]

This quote is taken from the Stanford Encyclopaedia of Philosophy, an online resource designed to give unbiased and informative overviews of various important philosophical issues. But even here, the bias towards punctal nihilism is clear to see. We are told that nihilism denies the existence of *all macroscopic material objects.* That is simply not true! Monistic nihilism, for instance, states quite the opposite: it states that the *only thing there is* is a macroscopic material object, the world, and it denies the existence of any *micro*scopic material objects. Intermediate variants of nihilism, too, need not deny the existence of macroscopic material objects, and may well deny the existence of microscopica. There is

[183] Markosian (1998a), 216; McDaniel (2003), 266. I should point out that neither Markosian nor McDaniel actually endorse the view that all simples must be point-sized.
[184] Dorr (2005); Sider (forthcoming); and Cameron (2010b) all make use of this paraphrase strategy.
[185] Wasserman (2009)

nothing to stop a nihilist accepting the existence of statues, ships, humans, and so on, as long as he also maintains they are all mereologically simple.

The common assumption, then, seems to be that compositional nihilism somehow entails punctal nihilism. In other words, *if* nihilism is true and all material things are mereologically simple, then those material simples will be minimally small. This mistaken assumption is what I shall call *the nihilistic fallacy*. It is not clear to me why so many have succumbed to this fallacy. Philosophers are usually such a careful and rigorous breed, so it is most surprising that many have ignored what should really be quite obvious. The fact that punctal nihilism has assumed the default position is one thing, but the fact that the other options are rarely even *recognised* is little short of astonishing. But the fact remains that, a few rare exceptions apart, this is how nihilist landscape currently sits.[186]

My best guess as to why philosophers so often succumb to the nihilistic fallacy is that they have been duped, or perhaps they have duped themselves, into thinking that their position is somehow *scientific*. The kind of reasoning that might lie behind such a thought could be presented as follows:

> *Modern science tells us that the universe is composed, ultimately, of fundamental particles. Tiny little bits of matter, perhaps even dimensionless, believed to be fundamental (i.e. to not have any smaller parts). These particles are called quarks and leptons. Compositional nihilism states that the only material things are mereological simples. Therefore, it would make sense to identify the simples posited by nihilism with the simples posited by science (i.e. leptons and quarks). In that sense, modern science and compositional nihilism are mutually supportive – they both have the same posits.*

This kind of reasoning – what I will call *naive scientism* – is flawed. The claim that punctal nihilism concords with modern science is a mirage, and any suggestions that science points towards punctal nihilism are false. Punctal nihilism posits *only* microscopic simples, whereas science posits a whole lot more. The chemist, for instance, will tell us that there are atoms and molecules, ions and isotopes. The biologist will tell us that there are living organisms. The cosmologist will tell us that there are stars and galaxies. All of these are composite objects, and none of them have any place in a nihilistic ontology, punctal or otherwise. Try

[186] Exceptions include Horgan & Potrč (2000; 2008); McDaniel (2007; 2009); Markosian (1998a); Parsons (2000). Note that of those listed, only Horgan & Potrč actually endorse compositional nihilism; the others merely endorse the possibility of spatially extended mereologically simple entities.

telling a biologist that your metaphysics is in concurrence with her science whilst insisting that there are no biological organisms – I guarantee she will think you quite mad.

I can envisage a response to this claim, however. The punctal nihilist will respond that whilst they do not admit many of the posits of science (indeed they will not admit any posits which are taken to be composite), their ontological picture is still perfectly *compatible* with science. Sure there are no molecules, or organisms, or galaxies, the punctal nihilist will say, but there are punctal simples arranged molecule-wise, organism-wise and galaxy-wise. The punctal nihilist can agree wholeheartedly with all the theories of modern science, it is just that she will have to substitute any quantification over composite objects with plural quantification over simples arranged composite-wise. So any objection that punctal nihilism contradicts modern science is unfounded.

This may well be so, but it is to miss the point entirely. Being compatible with science is a different thing entirely from being *vindicated* by science. And the fact that two theories are consistent does not entail that they are mutually supportive. For one should remember that the monist can tell a very similar story to that of the punctal nihilist. As we saw in the last chapter, the monistic strategy for dealing with talk about composite objects is to paraphrase it into talk of the world's properties. So whilst the monist will also deny that there are molecules, organisms, and galaxies, she will say that the world is moleculey, organismish, and galactical, or something to that effect. The monist posits the entire universe, and explains the illusion of its having parts in terms of its properties, whereas the punctal nihilist posits the fundamental particles and explains the illusion of there being composite objects in terms of the arrangement of those particles. But science posits particles *and* the universe (not to mention everything in between). To say that science, with its sub-atomic posits, points towards the truth of punctal nihilism would be objectionably one-sided. It could just as legitimately be said that with its positing of the entire universe, it points towards the truth of monism. Even an intermediate nihilist could claim parity on this score. An intermediate nihilist could say, for instance, that it is the atoms (and here I mean 'atom' in the scientific sense) which are mereologically simple. The illusion of larger composite objects would then be explained in terms of the arrangements of those atoms, and the illusion of smaller, sub-atomic objects would be explained in terms of the atoms's properties (e.g. an atom might be quarky or leptonic). So it is a mistake to suggest that science points

towards a particular variant of nihilism. All variants of nihilism can be made *compatible* with science, but none of them are vindicated by science. The reasoning underlying naive scientism, then, is fallacious.

Having exposed the nihilistic fallacy, one should hopefully now recognise that the evident bias towards punctal nihilism (or towards some form of Democritean nihilism) in the extant literature is unfounded. The arguments that have been presented for nihilism in the last chapter are neutral with respect to which variant of nihilism they support, as are the arguments customarily presented in the literature. Thus in the absence of any further arguments that support one variant over another, one should view all three variants of nihilism (punctal, intermediate and monistic) as being on an equal footing. In what follows, I will present a number of just such arguments and show that of the three variants of nihilism it is actually monism that should be preferred.

## §5.2. The Argument from Elimination

As has now been established, compositional nihilism must take one of three forms. There is punctal nihilism, which takes simples to be minimally small, i.e. point-sized. At the other end of the scale there is monistic nihilism (or just monism) which takes there to be just one simple, which is maximally large, i.e. world-sized. And there is intermediate nihilism, which covers any view in between the two, i.e. any view according to which simples are taken to be spatially extended, yet not maximally large. I think that there are some good *prima facie* reasons for rejecting both intermediate nihilism and punctal nihilism. If I am right, then this gives one a *prima facie* reason to endorse monism – since it is the only remaining option for the nihilist to take. Let me spell out these *prima facie* reasons below.

## §5.2.1 Eliminating Intermediate Nihilism

Intermediate nihilism, in almost any of its possible forms, should be rejected on grounds of arbitrariness.[187] You may recall from chapter 2 that a big problem for restrictivist theories of composition was that it was seemingly impossible to draw a line between cases of composition and cases of non-composition that was not entirely arbitrary. A similar problem afflicts intermediate nihilism. For once you admit *some* spatial extension to the simples that make up your ontology, it is hard to justify why a particular extension should be preferred over another. As I noted in the last chapter, it would be a brave philosopher who posited baseball-sized simples, for it would be very hard to know how to defend such a position. I for one would not know where to start. But this is exactly the kind of position that the intermediate nihilist *would* have to defend. Admittedly, there may not be a knock-down argument to such a position. But it simply seems crazy! Why are the simples baseball-sized? Why not golf-ball sized? Or indeed empire-state-building-sized, or even galaxy-sized? There is just no reason to suggest that any of these options are any better, or more likely, than any other, so to plump for one must be entirely arbitrary. Arbitrary assertions such as these, when there is no evidence whatsoever to support them, are truly objectionable, and thus should be avoided at all costs. Intermediate nihilism, then, should be rejected.

But perhaps that was a little hasty. For I can envisage one or two versions of intermediate nihilism that could lay claim to being somewhat less arbitrary and objectionable than others. So I will spend a little time here explaining what those versions of intermediate nihilism are, even though I will claim that, ultimately, they too should be rejected. Firstly, one could endorse a nihilistic view which admits objects of *all* sizes, yet maintains they are all simple. On this view, which I will call 'omni-magnitudinous nihilism', there will be point-sized simples *and* a world-sized simple, as well as myriad other-sized simples in between (e.g. table-size simples; earth-size simples; galaxy-size simples; and so on). This view avoids any accusations of arbitrariness, since it makes no arbitrary decisions concerning the size of material simples. Rather, it admits simples of *all* sizes. However, this theory will face a number of other problems; problems of arguably much greater severity. The main problem, I think, is that it would entail the ubiquitous *co-location* of objects. That is

[187] I say 'almost any' because there are one or two formulations of intermediate nihilism which may not be so susceptible to accusations of arbitrariness. I will consider, and ultimately reject, these formulations below.

to say that since there is a world-sized simple on this view, then all the other simples the theory posits will be partially co-located with it. I.e. there will be a multiplicity of distinct material objects (all of which are mereologically simple) occupying the same place at the same time. But co-location is weird at the very best and incoherent at worst, so on such a grand scale as this it would appear to represent a weighty theoretical cost. It would also appear that this widespread co-location of objects would lead to widespread causal over-determination. After all, if there are baseball-sized simples *and* much smaller simples arranged baseball-wise, then which of them is it that causes windows (and simples arranged-window-wise) to smash? For these reasons (coupled with the fact it is seemingly unmotivated, not to mention bizarre) I will rule out omni-magnitudinous nihilism as a serious possibility and not consider it any further.

But there is another, somewhat more plausible, option for the intermediate nihilist to pursue, which could be determined by one's view of the nature of space itself. One fairly natural way of thinking about space is that it is *continuous.* By that I mean that, for all intents and purposes, it is *infinitely divisible.* According to the view that space is continuous, for *any* line in space that is of a non-zero extension, it can always be divided into two smaller lines. Another way of thinking about space being continuous is to say that its extension in all three dimensions can be accurately modelled on the real numbers; for any two real numbers you care to choose, there will *always* be a further number that belongs between them. The real numbers present a continuous series; there are, as it were, no gaps. The same can be said of a continuous space; it has no gaps.

However, it is not universally accepted that that space *is* actually continuous; for there is a contrasting view that space is *discrete.* On this view, space is taken to be granular. It is said to consist of *spatial atoms,* regions of space with a non-zero extension which are nonetheless indivisible. (An intuitive way to think about discrete space is by analogy with your TV or computer screen. Just as your TV screen is made up of many small pixels, each of which is minimally small yet still extended, space is taken to be made up of many small spatial atoms). Now *if* one took space to be discrete, one may be in a position to adopt an intermediate nihilism that is not susceptible to the arbitrariness objection I raised above. For it is natural to suppose that if space is discrete then since material objects must be located *in* space, they could not occupy any region smaller than a single granule of space (for if they

*did*, then where, *precisely,* would they be?). Therefore, this view would have it that there is an absolute minimum size that any given material object can have – the size of a single spatial atom. Moreover, it would be natural to presume that these minimally small material objects must be mereologically simple, since if they were to have any proper parts those parts must be smaller than the object itself. But that is impossible, since the object itself is, by definition, minimally small. Thus contrary to my earlier assertion, it looks as if there *is* a version of intermediate nihilism which cannot be rejected on grounds of arbitrariness, and which, at the very least, deserves to be taken seriously.[188]

Whilst I do take such a view seriously, I think that there are a number of costs associated with it that provide reason (at least *prima facie* reason) to reject it. The first cost is that it would seem to entail a substantive realism about space. By a substantive realism about space, I mean *Substantivalism*: the theory which takes space to be a substance, a real thing, in which material objects somehow reside.[189] Substantivalism is often presented in opposition to a contrasting view: *relationism*[190]. Relationists deny that space is a substance, but instead claim that there are only distance relations which hold between material objects. Space is nothing over and above this network of distance relations. The debate between relationists and substantivalists is a long-standing and controversial one. But the version of intermediate nihilism currently being considered entails, I claim, a Substantivalist view about space. This should be considered a cost, not because Substantivalism is false (indeed, I want to remain neutral on that for the time being), but merely because it forces one to take sides on a debate on which the other versions of nihilism appear to be neutral.[191] This is admittedly not an unacceptable cost – indeed, those who already lean towards Substantivalism will be quite untroubled by it – but it is a cost all the same.

So why does the current view entail Substantivalism? For it is natural to think that one could believe in a discrete space (of sorts) *and* still be a relationist. In other words, one could be a relationist about space yet hold that there is a minimum magnitude across which

[188] Braddon-Mitchell & Miller (2006) actually defend a view similar to the one currently being espoused. The view they defend is not explicitly nihilist, in fact it is quite neutral over the debate on mereological composition. However, they argue that modern science suggests that there are minimally small regions of space (specifically, "Planck squares", two-dimensional regions with an area of $10^{-66}cm^2$), which should be considered non-point-sized, yet minimally small, simples.

[189] For a defence of substantivalism, see Graham Nerlich's classic book *The Shape of Space.*

[190] Relationism's most famous proponent is probably Leibniz. See, for instance, Alexander (1956), 25-26

[191] Later on, however, I will argue that the monist *should* reject substantivalism.

distance relations could obtain. This is true, and it is a coherent view.[192] But it does not rule out the possibility of point-sized objects. For on this view, there is nothing to prevent there from being point-sized objects, but only to prevent them from being closer to one another than the minimal distance relation allows. So the discrete *relationist* view of space is not enough to support an intermediate nihilist view. It is only if one takes a discrete space to be composed of substantive minimal spatial regions, that one can reasonably uphold the intermediate nihilist view being considered, and to take space to be like that is to be a committed substantivalist. So that is the first cost associated with this particular variant of intermediate nihilism – it commits one to substantivalism about space.

The second, and I think more damaging, cost to the current view is that there are a number of independently problematic consequences of supposing that space is discrete. One of these consequences is that certain metrical relationships that we would normally take to hold with necessity, appear to break down in a discrete model of space. In what has come to be known as 'Weyl's Tile Argument', Hermann Weyl points out that Pythagoras's theorem fails to hold if space is taken to be granular or discrete. He says: "If a square is built up of miniature tiles, then there are as many tiles along the diagonal as there are along the side; thus the diagonal should be equal in length to the side".[193] Of course we know that in a right-angled triangle the hypotenuse is *not* equal in length to either of the other sides, so something has clearly gone wrong. The discrete model of space, it seems, contradicts one of the most famous and long-standing mathematical proofs in human history – a significant cross for the model to bear!

To show why this is, consider *fig 2.* Suppose space is discrete, and that the minimum unit of distance is represented by the length of the side of one of the sixteen squares, let's call it *m*. Now since this space *is* discrete, then *ex hypothesi,* it is impossible to have any unit of length *less* than *m*. Moreover, all *possible* distances must be expressible in exact multiples of *m*. As one can see from *fig. 2.*, then, if you create a right-angle triangle, then there will be the same amount of squares on the hypotenuse as on the other two sides. And since there is no way of measuring distances that is more fine-grained than in terms of *m*, it follows that this

[192] Indeed, Forrest (1995) endorses such a view.
[193] Weyl, H. (1949), 43

right angle triangle has three sides of the same length (4*m* each). That is a violation of Pythagoras's Theorem.

*Fig. 2.*

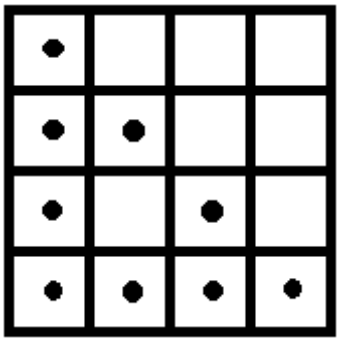

Now I should note that responses have been given to Weyl's Tile Argument, but they either involve mathematics that is in tension with the discrete model of space, or they no longer support the variant of intermediate nihilism that is being proposed.[194] So here we have our second cost associated with this particular variant of intermediate nihilism. From its base assumption that space is discrete, it suffers from all the significant problems that afflict that very model. This is not a completely insurmountable problem, for there are many reputable thinkers who *do* endorse a discrete model of space, but it represents a substantial theoretical cost. And when it is recognised that the other variants of nihilism under consideration appear neutral with regard to the nature of space, this cost gives us at the very least a *prima facie* reason to reject this version of intermediate nihilism.

There is perhaps one last line of response that I need to consider here before moving on. For it has been suggested to me that an intermediate nihilist could simply assert that there is an absolute minimum, yet non-zero, size that a physical object could be, and that

[194] Van Bendegem (1987) shows that the tile argument can be overcome if we take lines to have a non-zero and finite width (thus essentially as rectangles). However, the construction of van Bendegem's solution draws on classical Euclidian geometry as it operates in a *continuous* space. But of course, his solution is meant to apply to a *non-continuous*, discrete space. Specifically, there is no definition of a line in discrete geometry as there is in classical geometry. Van Bendegem has assumed that the definition of a line from classical geometry can simply be transferred to apply in discrete geometry, and this assumption is questionable to say the least. Forrest (1995) overcomes the tile argument by supposing that discrete space is not composed of some regularly-shaped spatial atoms, as is often supposed, but rather, composed of spatial points related to one another by minimally small distance relations. Whilst Forrest's response may overcome the tile argument, it would seem to allow the existence of point-sized objects, and thus undercut the version of intermediate nihilism being currently proposed. See Dainton (2010), 297-8 for further discussion.

would be regardless of one's view about space. The idea is, I think, that there are lots of natural limits, so to speak, in the world. The speed of light, for instance, marks the absolute maximum velocity at which an object can travel. And in a sense, this maximum speed limit is entirely arbitrary, since it surely could have been different. But at the same time, there is nothing *intolerable* about having such a limit. So by extension, why not just think that there is, as a matter of fact, some absolute minimum size that an object can be. We may not know exactly what that size is (although presumably it will be very small indeed), but that is not to say that there isn't one at all.

I am unconvinced by such a response. Firstly, it should be noted that the speed of light was an *empirical discovery,* and, moreover, it was through an application of the laws of physics *plus* empirical data that we discovered that the speed of light represented a universal speed limit. The fact that the speed of light is what it is (approx 670 million mph) is in a sense arbitrary, but it has been corroborated by empirical data; it is just a standard contingent fact, and there is nothing at all objectionable about that. What *would* have been objectionable is if some philosophers, upon pondering the speeds at which objects could travel, had simply asserted that there *must* be a maximum limit, in the absence of any reason to so assert. That *would* be arbitrary and unjustified. But this, it seems to me, is what would be being done if an intermediate nihilist were to suppose that there were some absolute minimum size that an object could be. At the very least, the burden of proof would be on them to provide a reason as to why we should believe such a claim. But what could be the reason?

Think of the opposite end of the spectrum: *large* objects. There is, it seems, a perfectly good reason to suppose that there is a maximum size an object can be. Supposing that an object can only exist *at a world,* then the world at which it exists will set the limit for its size. The absolute maximum size for an object would be the same size as the world of which it is part. Moreover, this is an *a priori* reason; one can establish its truth from the comfort of one's armchair, since it is logically incoherent to suggest that an object could be larger than the world in which it exists. But applying the same reason to the very small would surely indicate that the smallest an object could possibly be would be point-sized. Because for any non-zero size stipulated, it is always possible to suppose that an object could have been smaller, so to suggest that there is some non-zero minimum size limit looks objectionably

arbitrary. Just as it would do if one were to suppose that there was some maximum size limit for an object that is nonetheless smaller than world-size.

At this point, one may draw on the findings of contemporary science, and insist that there *is* a minimum size for objects; it is whatever the size is of the smallest objects known to science (quarks, or neutrinos, or whatever). But this argument doesn't wash either. These tiny fundamental particles are merely the smallest particles *known to man;* they are not the smallest particles *possible.* There is no reason to suppose that in the future it won't be discovered that what we currently take to be fundamental particles are in fact made of smaller parts, perhaps even point-sized parts. Or alternatively, there is no reason to suppose that entirely *new* particles will not be discovered that are even smaller than the fundamental particles of today. So whatever size science tells us that the smallest particles *are* is quite irrelevant. To infer that this is the smallest size *possible* would be unjustified.

To conclude, then, intermediate nihilism should be rejected. Most formulations of it are intolerably arbitrary, and those that are not arbitrary suffer from other significant problems, such as a commitment to substantive and controversial views about the nature of space, or the ubiquitous co-location of objects and widespread causal overdetermination. I will consider intermediate nihilism, therefore, eliminated.

## *§5.2.2 Eliminating Punctal Nihilism*

If one puts intermediate nihilism out of the picture, and in light of the preceding comments it would seem prudent to do so, one has two options remaining: punctal nihilism or monism. In the extant literature, the former of these options is by far the most popular.[195] But there are some serious problems with punctal nihilism. In particular, there are some serious problems involved with the supposition that all matter is point-sized. I think that these problems are serious enough to suggest that, other considerations notwithstanding, punctal nihilism should be rejected. By process of elimination, then, monism is vindicated as the only viable form (or at least, the most viable form) of compositional nihilism. For clarity's sake, the argument could be presented formally as follows:

[195] Sider (forthcoming), Cameron (2010b), Dorr (2005), for instance, are all punctal nihilists. Horgan & Potrč (2000; 2008) are the only monistic nihilists I know of.

1. Only three varieties of compositional nihilism are available: punctal, intermediate, and monistic.
2. Punctal and intermediate nihilism should be rejected.
3. Therefore, if one is a compositional nihilist, one should be a monistic nihilist.

I have hopefully already made a decent case for the claim that intermediate nihilism should be rejected. So now let me begin my case against punctal nihilism.

Punctal nihilism says that there are only simples, and moreover, that these simples are minimally small. I take minimally small to mean point-sized, where point-sized means having zero-extent in any spatial dimension. I think this is a fair assumption to make, purely because if they were *not* point-sized, they would not be minimally small. I.e. they could be smaller. (This is intended in the same sense as saying that any number greater than zero is not minimally small, since it could be smaller). Any nihilism which posits its simples as being larger than point-sized should not be classed as a punctal nihilism. But the very idea of point-sized objects raises immediate problems. These problems are long-standing, and can be traced right back to Ancient Greece. Zeno, in particular, is famous for his interest in, and rejection of, point-sized particles, as is demonstrated by the following passage:

> If it [a point-sized thing] were added to another existent, it would not make it any larger. For if it is of no size but is added, there cannot be any increase at all in size. Thus what is added is therefore nothing. And if when it is subtracted the other thing is no smaller – and will not increase when it is added again – then clearly what was added and subtracted was nothing.[196]

This kind of thought should be familiar to anyone who has ever taken an undergraduate metaphysics course. The reasoning is simple yet powerful. To make things a little clearer, I think one could construct a more structured Zeno-esque argument as follows:

1. Extended material objects are ultimately composed of point-sized atoms. (Assumption for *reductio*).
2. Point-sized atoms have zero extension.
3. No matter how many atoms of zero extension you may have, the total extension of their sum will always equal zero. ($0 \times n \equiv 0$).

[196] Taken from Barnes (2001), 102

4. Therefore, extended objects could not be composed purely of point-sized atoms.
5. Thus we have a contradiction; 1 must be false.

It is clear how this argument is supposed to affect the punctal nihilist. The punctal nihilist says that there are only point-sized simples. All *illusions* of extended objects are to be explained by the existence of point-sized simples arranged in particular ways. The Zeno-esque objection, then, would be that there *could not be* any extended objects, or even the *illusion* of extended objects, if there are only point-sized simples. Non-extended objects have not the capacity to conjure *even the illusion* of extended ones.

Whilst this kind of reasoning can seem compelling, I think its attraction is superficial – for there are ways one could respond. Specifically, it should be noted that the Zeno-esque argument only works if extended matter is taken to be *continuously dense.* By that I mean that the simples which are taken to compose some extended object (or taken to be arranged object-wise) are taken to be directly contiguous, i.e. touching one another. If one makes this assumption, then the argument looks a threat, since the total extension of any number of *directly contiguous* non-extended simples will always be zero. But why should one make such an assumption? Why not just think that whilst there are only point-sized simples, they are spread out (as depicted in *fig. 3*)?

*Fig. 3.*[197]

We are already used to this kind of understanding of macroscopic objects. Modern science tells us that ordinary objects (and indeed the atoms that compose them) are mostly empty space. Thus it would be wrong to think of the atoms that compose your table (or that are arranged table-wise) as being packed together so tightly that they are directly contiguous. Rather, they are spaced out at intervals really quite huge relative to their microscopic size.

[197] The diagram is not a truly accurate representation of the idea being proposed, of course, because the punctal simples (represented by the blue spheres) would not have any extension (as they clearly do in the diagram), and there would be no actual *things* (i.e. the white bars) connecting them together, but just empty space. However, the diagram gives a good flavour of what is being put forward.

So it seems that this kind of view is also available to the punctal nihilist, *and* it would avoid the force of the Zeno-esque argument. To be precise, this view would undermine premise 3 of the argument, for whilst the simples themselves have no extension, the distance from one simple to another would not be zero so long as they are not directly contiguous. So it is simply not the case that for any amount, *n,* of point-sized simples that their total extension will be zero. If the *n* simples are arranged in a line, for instance, and are spaced with gaps of a non-zero magnitude, *x,* then their total extension will be $x(n - 1)$. Punctal nihilism, it seems, can overcome the threat from Zeno.

However, a more modern take on this Zeno-esque argument has been offered by Peter Simons, and *this* argument is not so easily overcome.[198] I will tweak his argument somewhat to target it more squarely in the direction of punctal nihilism, but the thrust of it will remain the same as the original:

1. If punctal nihilism is true, then the only physical things there are are point-sized simples.[199]
2. Therefore, whatever fundamental physical properties there are must be instantiated by those simples (since there are no other candidate objects to instantiate them).
3. Mass is a fundamental property.[200]
4. Therefore simples have mass.
5. But if simples have mass, then they must be infinitely dense, since they have a non-zero mass in a zero volume.
6. Physical objects cannot be infinitely dense.
7. Therefore, either 1, 2, or 3, must be false (or a combination all three).

[198] Simons (2004), 373

[199] I include the 'physical' qualification since nothing in the letter of punctal nihilism rules out the possibility of there being abstract objects.

[200] Peter Simons has suggested to me that mass may not be a fundamental property after all. Very roughly the idea is that if *processes* are taken to be more ontologically fundamental than *objects,* then fundamental properties would be those instantiated by processes rather than objects. On this view, mass *would* be instantiated by objects, and thus would be definable in terms of more fundamental properties of processes. Specifically, mass would be definable in terms of Planck seconds per square metre. Whilst this proposal is undoubtedly very interesting, it is most certainly controversial, and to discuss it further would take us too far afield, so I will not do so.

Premise 2 looks undeniable. Physical properties can only be instantiated by physical things, so if the only physical things are *x*s (whatever those *x*s might be), then *x*s are the only things that can instantiate physical properties. Premise 3 also looks unshakable, unless one wants to deny that the world is massive. So it is premise 1, then, which should be rejected. But of course, rejecting premise 1 is just to reject punctal nihilism. One should also note that the 'spacing out response' that I laid out above, which overcame the original Zeno-esque argument, doesn't look as though as it will be of any help here. For no matter how far apart material simples are taken to be, they are still the only candidate objects available to instantiate mass. Mass cannot be instantiated by 'gaps', unless we reify those gaps into material things, which would go against the central thesis of punctal nihilism. So the conclusion is undeniable; if punctal nihilism is true, then all matter must be infinitely dense.

Despite the clear force of these types of arguments that highlight the physical difficulties of admitting infinitely small, point-sized objects, they are summarily dismissed. It is routinely asserted that in the age of post-Cantorean mathematics, metaphysical problems stemming from considerations of the infinite are mere chimeras; they can be explained away. Sorensen confidently asserts that "all of Zeno's paradoxes were solved by Cantor a hundred years ago".[201] While Dean Zimmerman, somewhat more reservedly, tells us that "these paradoxes are significantly diffused by Cantor's discovery of the distinction between denumerably and non-denumerably infinite numbers"[202] Elsewhere, Sider nonchalantly asserts that "traditional arguments" against point-sized objects "are unconvincing since we now know that theories of point-sized things are mathematically coherent".[203] But I think this is far too quick. For even if one concedes that such theories are *mathematically* coherent, does it necessarily follow that they are *metaphysically* coherent? It is one thing to demonstrate that within a certain abstract, idealized model that the numbers somehow add up, but it is quite another to suggest that this is what the world is really like. The fact that point-sized objects must be infinitely dense is referred to by Simons, somewhat politely, as a "minor embarrassment" of any theories which posit them. But I wholeheartedly agree with him when he says that whilst this may be acceptable for the physicist, or mathematician, whom are dealing with idealized models or theories, these

[201] Sorensen (2003), 56-7
[202] Zimmerman (1996a), 1
[203] Sider (forthcoming), 29

"minor embarrassments" are simply not acceptable for the metaphysician. The metaphysician must take them "literally and seriously".[204] The difficulties with point-sized objects are very real. They cannot simply be 'mathematised' away.

In his 1951 paper, Max Black demonstrates how mathematical solutions to metaphysical problems are not always satisfactory.[205] That is, he shows that whilst they may be *mathematically* satisfactory, they will not always satisfy the metaphysician, since they won't always get to the metaphysical *heart of the matter.* Black's particular target is Zeno's paradox of Achilles and the Tortoise. Zeno supposed that Achilles could never catch up with the tortoise, since before he could do so, he would first have to reach the point at which the tortoise started, from which the tortoise would then have moved. But this sequence would go on infinitely, since although the distance the tortoise moves decreases in each case, it will always have a non-zero value. Thus the thrust of the argument is that Achilles will never catch up because in order to do so he will have to travel an infinite number of non-zero distances. Black's point is that the mathematical solution to this paradox, offered by Whitehead (among others), doesn't really *solve* anything at all.[206] He says of the solution: "It tells us, correctly, when and where Achilles and the Tortoise will meet, *if* they meet; but it fails to show that Zeno is wrong in claiming they *could not* meet".[207] It is as though the mathematics deals with only one side of the paradox and leaves the other side untouched. But therefore it *remains* a paradox; it is no good telling us *when* Achilles will catch the tortoise if at the same time it is supposed he can never catch the tortoise! It would be akin to claiming that one had solved Russell's paradox by merely asserting that the barber actually shaves himself!

Similarly in this case, whilst mathematics can provide us with coherent models that posit point-sized particles, they fail to get to the metaphysical heart of the matter. The mathematical models employed by physicists will often take fundamental particles to be point-sized, but this is only for theoretical convenience. Physicists do not take the particles to *actually* be point-sized, but merely maintain that they are small enough for the difference to be theoretically negligible. This is why consequences like infinite density only represent

[204] Simons (2004), 373
[205] Black, (1951)
[206] Whitehead (1929), 107
[207] Black (1951), 93

minor embarrassments for physicists since, as Simons points out, these consequences are merely a "product of idealization".[208] But the punctal nihilist does not claim simples to be point-sized purely for theoretical convenience. Rather, she claims that all matter *actually is* point-sized. Thus the punctal nihilist has to take the claim that point-sized simples must be infinitely dense at face value. But this, it seems to me, is no minor embarrassment; it is surely a much more serious consequence than that. But it is a consequence that the punctal nihilist must accept, all the same. Thus theories which posit only point-sized particles may well be, as Sider says, mathematically coherent. But that only tells one side of the story. Their *metaphysical* coherence is in serious doubt, since they postulate that all matter is infinitely dense.

At this point, one may reasonably ask what is so bad about infinite density. The punctal nihilist may well be happy to simply accept that point-sized simples, if massive, must be infinitely dense, and have done with it. That is, I suppose, a possible response. But it is surely a costly or counter-intuitive one. Admittedly, I cannot *demonstrate* that infinite density is impossible, but it is surely not something to be taken lightly. And remember that the punctal nihilist does not only have to accept that it is merely possible for an object to be infinitely dense, but rather, she has to accept the much more significant conclusion that *all* massive objects are infinitely dense. It is my view that such a conclusion is, at best, far-fetched, and at worst, bordering on the absurd.

The monist, by contrast, can sidestep the whole issue. There is no room for point-sized objects in the monist's ontology, and because of this, there is no room for the problems that come with them. Simons's concerns about infinite density certainly don't apply to the monist, for the monist will simply assert that it is the *world* that instantiates fundamental properties like mass (because the world is the only candidate object *available* to instantiate properties), and yes, the world will have a non-zero mass, but since it also has a non-zero extension, no problems arise. Zeno-type concerns that extended objects cannot be broken down entirely into non-extended objects will also not apply to the monist, for the obvious reason that the monist takes there to be only *one* extended object, which cannot be broken down at all.

[208] Simons (2004), 374

On the face of it, then, monism should be the default position for the compositional nihilist. Intermediate nihilism should be rejected since most versions of it are objectionably arbitrary, and those which are not entail significant and problematic views about the nature of space. Punctal nihilism should be rejected because it entails the undesirable consequence that all matter must be infinitely dense. Other considerations notwithstanding, then, monism should be the nihilism of choice.

## §5.3. Gunk

It is a fairly common assumption, among both philosophers and scientists alike, that matter must 'bottom out' at some point, in the form of fundamental, or elementary, particles. In other words, the assumption is that, at base, matter is constituted by some very small, fundamental minima which are themselves indivisible: we can divide macroscopic objects into smaller and smaller bits, but there is a point at which that division must stop. The standard model of modern physics, for instance, says that the ultimate constituents of all matter are quarks and leptons, and that they are simple – they have no parts. Leptons and quarks, we are told, are the ultimate building blocks of reality.

But what if this assumption is wrong? What if it turns out that these elementary particles *do* in fact have parts, and that moreover, those parts have parts, and so on *ad infinitum*? The question is: why should we assume that there is a fundamental layer, at which matter can no longer be divided, at all? The possibility of matter which is divisible 'all the way down' has been postulated by a number of philosophers throughout history. It has been said to go back even as far as Anaxagoras, who, according to one interpretation, thought that, "every stuff always contains a 'portion' or 'share', however small, of every other stuff".[209] But it was David Lewis who is largely responsible for the recent surge of interest in this intriguing conception of matter. He introduced the term "atomless gunk", to apply to any substance "whose parts all have further proper parts".[210] This terminology seems to have stuck, and interest in gunk has increased ever since.[211]

[209] Barnes (2001), 189
[210] Lewis (1991), 20.
[211] See Hudson (2007) for a useful introduction to some of the recent research into gunk.

Gunk presents a problem for the compositional nihilist. For nihilism says that everything is mereologically simple, but gunk is, by definition, mereologically complex. So if any gunk were ever to be discovered, it would instantly sound the death knell for any type of compositional nihilism. But I don't think many nihilists are particularly worried about that, because gunk is not really the type of thing that can be discovered. We saw in chapter 2 how empirical research is seemingly incapable of revealing the mereological structure of material objects. So if one cannot empirically demonstrate that matter is mereologically complex, it seems to follow that one cannot demonstrate it to be *infinitely* complex either. Despite this, however, the threat of gunk remains. The believer in gunk does not have to produce any *actual* gunk to undermine the nihilist's thesis, but rather, she has only to show that such stuff is *possible.* The threat of gunk hinges on the fact that nihilists usually take nihilism to be not only true, but *necessarily* true. With this modal assumption in place, the threat from the possibility of gunk can formalised in the following argument[212]:

1. If compositional nihilism (CN) is true, then it is necessarily true.
2. If CN is necessarily true, then all matter, at all possible worlds, is mereologically simple.
3. If gunk is possible, then there is a possible world where matter is gunky (and thus mereologically complex).
4. Therefore, if gunk is possible, CN is not necessarily true.
5. Gunk is possible.
6. Therefore CN is not necessarily true.
7. Therefore, from 1, CN is not true.

There are a number of ways in which a nihilist could respond to this argument, some of which show more promise than others. Importantly, though, the degree to which a nihilist response succeeds against this argument will depend, I claim, on the variety of nihilism one endorses. In particular, I think it can be shown that the responses available to the monist can be considered vastly more successful than those available to the punctal nihilist. The threat from the possibility of gunk can, somewhat ironically, strengthen the case for monism.

[212] An argument of this type was first given by Sider (1993a), although the precise formulation is slightly different.

The first way a nihilist could respond would be to renege on premise 1. In other words the nihilist could weaken her nihilism somewhat by claiming that whilst it is true, it is not *necessarily* true. On this view, the world would be seen as being a nihilistic one, compositionally speaking, but only as a matter of contingent fact; there would be nothing to preclude the fact that it might not have been so. This view would nullify the argument, for it is entirely compatible with the possibility of gunk. But its advantages come at a steep cost. The cost is that to endorse this view, one must also endorse the view that the facts about mereology are contingent. There is a vast array of conflicting views about mereology, as has been shown in the preceding chapters, but despite this, it is widely agreed that whatever the facts of mereology are, they hold with metaphysical necessity. And there is good reason for this. Mereology is a paradigmatically *metaphysical* subject matter. Its possibilities are to be explored with *a priori* reasoning, and thus the conclusions one reaches should be considered *a priori* truths. At least, this is the orthodox view on these matters. One or two contemporary thinkers are beginning to question this orthodoxy,[213] but that is not a path I wish to pursue here. Moreover, many of the arguments given earlier in support of compositional nihilism are *a priori* arguments, and it strikes me that the very force of these arguments would be severely undermined were one to maintain that their conclusions were only contingently true. Thankfully, however, and as I shall presently show, one does not have to endorse this controversial claim about the modal status of mereology in order to overcome the argument from the possibility of gunk.[214]

A much more promising response to the argument from gunk is to deny premise 5. That is, to deny that gunk is possible at all. This response too would nullify the argument, but of course, it will need some independent support. For lots of philosophers think gunk *is* possible.[215] There are, I suppose, reasons to think that gunk is possible. Inductive reasoning may be touted here, for instance, in that over the course of the last few centuries, science has continued to discover increasingly small particles of matter. It used to be thought that atoms were the ultimate constituents of matter (hence the name 'atom', presumably), until sub-

[213] See Cameron (2007). Peter Simons, too, suggests that the laws of mereology may not be knowable *a priori.* See Simons (2006), 597.

[214] I will revisit the issue of the contingency of mereology in the next chapter.

[215] See, for example, Zimmerman (1996a); Sider (1993); Forrest (2004); to name but a few contemporary thinkers. Going further back, Leibniz was an ardent defender of infinitely divisible matter. Van Cleve (2008) believes that gunk is not only possible, but that in actuality, "all is gunk" (p.9)

atomic structure was discovered, at which point it was presumed that sub-atomic particles (protons, and neutrons, and so on) were the bedrock. But then of course it was discovered that these too were made of smaller things – quarks and leptons. So the inductive argument appeals to this progressive sequence and suggests it will, or at least *could,* continue *ad infinitum,* such that for each smaller layer we uncover, there will be a further layer below at which objects are smaller still. But as has been noted elsewhere, this inductive argument should not be taken all that seriously.[216] After all, to infer an *infinite* sequence from such a small number of cases is philosophically naive at best, and at worst, methodologically irresponsible.

Perhaps the most common reason for thinking that gunk is possible is that it is *conceivable.* Some thinkers consider this to be simply evident, and that this evident conceivability is a sure sign of possibility. Jonathan Schaffer, for instance, thinks the view that matter is gunky is not only conceivable, but that it is "if anything the *more natural* doctrine [as opposed to the atomistic view of matter]" and that the existence of opposing views merely "evidences brainwashing".[217] Likewise, Ted Sider says: "*Surely,* there are both atomistic possible worlds *and* gunk worlds", before going on to assert: "I find the *possibility* of gunk so compelling that I am willing to reject any theory that rules it out".[218] And J.R.G. Williams unfalteringly states that "a world of infinite descent is conceivable [...] *all* parties will agree".[219]

I remain unconvinced. Firstly, I am not convinced that gunk *is* as easily conceivable as is often suggested, to the point that I sometimes genuinely wonder *exactly what it is* that is being conceived of when it is stated that one is conceiving of atomless gunk. The *definition* of gunk is easy to understand, sure, in terms of the semantic content of the terms and concepts employed. But that doesn't necessarily mean that gunk itself is easily conceivable. I have a good grasp of the terms 'five', 'side' and 'square', for instance, but that doesn't mean I can conceive of a five-sided square. So just because I can understand what you mean when you tell me that gunk is matter whose parts all have further parts, that does not necessarily mean I can conceive of such stuff existing, let alone conceive what it would actually be like.

[216] See Sider (forthcoming), 29-30
[217] Schaffer (2003), 502. For the record, I do not find Schaffer's arguments for this view at all convincing.
[218] Sider (1993a), 288. Interestingly, Sider later recanted on this, and is now a committed punctal nihilist. See Sider (forthcoming).
[219] Williams (2006), 503

Perhaps this is just a sign of a limited imagination on my part (in the gunk case, of course, not the five-sided square case; I refuse to believe anyone who says they can conceive of one of those) – that I am willing to accept.[220] But that is fortunately of little consequence, because importantly, even if it is conceded that gunk is perfectly conceivable, that does not entail that it should be considered a metaphysical possibility. Compositional nihilism rules out the possibility of lots of things that are easily conceivable. For if nihilism is a necessary truth, then it rules out the possibility of any composite objects whatsoever. But composite objects are not inconceivable! I can quite easily conceive that some simples arranged table-wise, say, may compose a table. But that does not pose any threat to my nihilism. It is simply a consequence of nihilism, which is argued for on entirely independent grounds, that there are no such things, *necessarily,* as tables. Tables and chairs and cats and dogs are ruled out as metaphysical impossibilities.[221] And that is *despite* their evident conceivability.

It is fine for metaphysical theories to rule out certain types of entity, or certain types of possibilities. All metaphysical theories will do that to some extent. Metaphysical theories of all stripes are often in tension with other doctrines; they often deny the existence of the posits of common sense, for example, or the posits of science, or whatever. The important thing is, however, that if these theories want to be taken seriously whilst maintaining that much of what common sense or science states is false, they must be able to offer some kind of *explanation* as to why these falsities are so commonly believed. They must be able to be *reconciled* with one another. Punctal nihilism, for instance, denies that there are chairs. But it does not deny that we can sit down! For it comes ready furnished with a plausible explanation of why we ordinarily think that there are chairs. There are no chairs, says the punctal nihilist, but there are simples arranged chair-wise. And the sentence "there is a chair" is not strictly true, but you are excused for thinking it is true because there is a relevant paraphrase of that sentence ("there are simples arranged chair-wise") that *is* true. The tension between common sense and nihilism is relieved, then, by the provision of these suitable explanations. The *illusion* of chairs is explained away by the existence, and arrangement, of simples.

[220] Although, in the next chapter I will say a little more about one should be suspicious about the conceivability of gunk.
[221] At least, if they are considered composites. There is nothing to rule out a mereologically simple cat.

But what about the illusion of gunk? Many philosophers, as we have seen, think gunk is a genuine possibility. Certain scientists and philosophers of science, too, are positing gunk, or theories of gunky space-time.[222] Schaffer even goes so far as to suggest that "the folk may be committed to atomless gunk".[223] So if nihilism is to maintain its plausibility in the face of such commitment to gunk, it must be able to *explain away the illusion of gunk.* Just as it denies tables and chairs but replaces them with simples arranged table-wise and chair-wise, so if it denies gunk, it must replace it with some suitable explanation. But it is hard to see what this explanation could be. The sentence "there could be some gunk" cannot be paraphrased with "there could be some simples arranged gunk-wise". It makes no sense to say that simples could be arranged gunk-wise, because to be arranged gunk-wise would involve, *by definition*, being arranged in such a way that there are no simples! The illusion of gunky objects cannot be explained away by the positing of punctal simples.

Monism, by contrast, does not have the same problems when dealing with the possibility of gunk. Obviously, the monist, just like any nihilist, must deny that gunk is *genuinely* possible, since it would involve the instantiation of parthood relations (and multiple objects) which the nihilist eschews. But the monist, it seems, can deal with the *illusion* of gunk perfectly well. That is, the monist can explain why it might seem like gunk is a genuine possibility, whilst the punctal nihilist is left floundering. The monist maintains that the world is a single, mereologically simple object. She suggests that all illusions that the world has parts are to be explained in terms of the world's properties. So it may seem like the world has table-like parts, but it doesn't. Rather, it has the property of table-ness distributed in those areas we posit the tables. Likewise, it might seem like those alleged tables themselves have parts: legs, tops, molecules, or whatever. But, of course, they don't. Rather, the world has various properties (table-leg-ness, and so on) that are responsible for the illusion of these things. The important thing is, that if this explanatory strategy is accepted, then it seems to explain the illusion of gunk just as well as it explains the illusion of more ordinary composites like tables and their legs. If the world can instantiate the property of table-ness, which explains the appearance of tables, then there is no reason why the world can't instantiate the property of gunkiness, which would explain the appearance

[222] See, for example, Artzenius (2008)
[223] Schaffer (2003), 502

of gunk. In more general terms, once it is recognised that the monist can explain the illusion that the world has parts *at all*, then there seems no reason why any *particular* parts should be more difficult to explain than others. If the world can instantiate properties that give rise to the illusion that it is composed of galaxies or molecules, then it could just as easily instantiate a property that gives rise to the illusion that it is composed of gunk.

The key difference between the punctal nihilist and the monistic nihilist, in terms of their explanatory strategies, is the *order* of explanation. The punctal nihilist starts at the minimally small and explains upwards; whereas the monist starts with the maximally large and explains downwards. And this is why the punctal nihilist has a problem with gunk, because even if you start with minimally small entities, to explain gunk is to explain downwards even further. But the punctal nihilist cannot go any further down! By stipulation, the minimally small simples represent the ultimate bottom layer, so there is nowhere left to go. Since all the monist's explanations of ordinary objects are already downwards, in the sense in which the term is currently being used, then there is no obstacle to letting those explanations go on infinitely. Gunk is perfectly explicable in monistic terms.

Gunk, then, presents a problem for the compositional nihilist. The nihilist cannot allow that gunk is a genuine possibility (at least, not if she wants to hold that her nihilism is a necessary truth). But the punctal nihilist cannot even explain the illusion of gunk – why it is that it at least *seems* that gunk is conceivable, and thus possible. The monist has no such problems. So if the compositional nihilist takes the illusion of gunk seriously, then she has good reason to prefer monistic nihilism over punctal nihilism.

## §5.4. A Very Brief Word on Junk

In the interests of balance, I feel that I should say a least a few words here on what has been called 'Junk'.[224] One can think of Junk as a kind of weird and wonderful alter ego of Gunk. For whilst a gunky world is one in which everything *has* proper parts, a junky world is one in which everything *is* a proper part. In a junky world, then, there is no maximal fusion (what one would ordinarily call 'The Universe') than contains everything else as proper

[224] The term is from Schaffer (2010c), 64.

parts. Rather, a junky world is a world consisting of "an *infinite plurality* of objects such that each thing in the plurality is a proper part of something else in the plurality".[225]

Junk has the potential to threaten Monism in the same way that Gunk threatens Punctal nihilism. Of course, all versions of nihilism will deny that Junk is a *genuine* possibility, because junk is, by definition, mereologically complex. However, if Junk is conceivable, then just as was the case with gunk, the nihilist needs to be able to explain the *illusion* of Junk. But this time around, it is the monist who is going to struggle. The punctal nihilist, by contrast, should have no trouble. For the punctal nihilist explains the illusion of complex objects from the bottom up, i.e. she starts off with minimally small simples, and explains how they can be arranged to give the illusion of larger, complex objects. If this strategy works at all, then there is no reason why it could not continue infinitely. I.e. if we can conceive of an infinite plurality of simples, then we should also be able to conceive how such a plurality could give rise to the illusion of an infinite hierarchy of ever increasing complex objects. Thus the illusion of junk is explained.

The monist, by contrast, starts from the top and explains downwards. That is, she begins with the world itself, a maximally large simple, and explains the illusion of its having parts in terms of its properties. But it is difficult to see how the world could instantiate a property that could give rise to the illusion that the world itself was a proper part of some other, bigger, object. That would be a strange property indeed. In fact, the whole top-down explanatory strategy that the monist employs looks doomed to failure when it comes to junk. Because if the world was junky there would be no top level at all – the monist's explanation could never even get off the ground! Junk, therefore, has the potential to undermine monism.

I reject the threat from junk solely because I reject the claim that it is even conceivable. I have already expressed doubts that *gunk* is genuinely conceivable, but in comparison to Junk, gunk is veritably *ordinary* fare. *For junk is plain weird.* Too weird for me, it must be said. And I am not alone in thinking this. Einar Bohn is the only contemporary philosopher I know of who endorses the possibility of Junk.[226] But, of course, weirdness is

[225] Bohn (2009a), 29

[226] See Bohn (2009a; 2009b). Bohn claims, however, that Junk is backed with a little authority from philosophical history, since he claims that Leibniz and Whitehead believed the actual world to be junky.

not enough to simply reject something out of hand.[227] To show why I think that junk is inconceivable, I invite you now to try to conceive of some. I think you will struggle. Firstly, one should note that you cannot conceive of a piece of junk that has a finite size. If that is what you were conceiving of (i.e. some medium-sized lump of junky stuff), then you were not, I'm afraid, conceiving of junk. To show why this is so, assume for *reductio* that you had a finite-sized piece of junk, *J*. Let's say that *J* is one metre cubed. This would mean that within that *J* there is a maximally sized part, i.e. a part that is one metre cubed. But this piece cannot be a proper part of anything else, because by stipulation, it is the largest bit of junk there is in *J*, and for *x* to be a *proper* part of something entails that *x* is *smaller* than that something. So therefore we have a contradiction. I.e. *J* has a part that is not a *proper* part, yet *J* is meant to be junky, such that *all* of its parts are proper parts. Thus we have to conclude that if a junky object is possible at all, it must be infinitely large.

But this is not the main problem. The main problem I have is that in talking of a junky object *at all,* one is already in danger of descending into incoherence. For in conceiving of an object, *any object,* one seems to be conceiving of the *whole* object. But I don't think it makes sense to talk about a *whole* piece of junk. Because talking of a whole implies that you are talking of *all* of the object. But if that's right then what you are talking of can't be a proper part of anything else, or you wouldn't actually have been talking about *all of it*. So in fact, you can't have been talking about junk at all. In short, then, I fail to see how one can conceive, and I mean *genuinely conceive,* of junk. Jonathan Schaffer seems to take a similar view. He says: "No world – providing that worlds are understood as possible concrete cosmoi – could contain worldless junk because a world that contained junk would be an entity not a proper part of another entity at that world. A world would top-off the junk".[228] It is for this reason that I reject the threat of gunk. I deny that it is either possible or conceivable. The monist cannot explain the illusion of junk, but neither can anyone else, because there is actually no illusion there to explain.

[227] For I am sure that there are a number of philosophers who would dismiss the thesis of monism that I am presently defending as *plain weird.*
[228] Schaffer (2010c), 65

## §5.5. The Problem of Emergence

Emergence is a term applied to events or properties, or sometimes just phenomena in general. For clarity, I will restrict my focus to emergent properties, although not a great deal should hang on this restriction. An emergent property is a property of an object or system that cannot be explained or accounted for solely in terms of the properties of that object's or system's parts. These properties *emerge* at a certain level, and cannot be reduced to properties instantiated by objects at a lower level. Emergent properties are, in this sense, *irreducible.*

Emergence is somewhat of a hot topic in contemporary discussion. Over the last few decades, interest in emergent phenomena has been rekindled thanks to the work of a number of notable philosophers.[229] It would be fair to say that over recent years, there has been a resurgence of emergence. The most common forum in which emergence is discussed is within the philosophy of mind. For it is often thought that there are certain aspects of consciousness that are emergent, i.e. cannot be reduced purely to properties of cerebral cells. The thought is that intentional properties or mental states (things like qualia) are so entirely distinct in character from the neurological properties that are instantiated by parts of the brain, that they cannot be explicable purely in terms of those properties. They may well be *caused* by activity in the brain, yet they emerge holistically as being far greater than the sum of their causal beginnings. Another quite distinct field in which emergence plays a prominent role is quantum mechanics. Very roughly, the thought is that certain composite quantum objects or systems (often referred to as 'entangled systems') can exhibit properties that are quite inexplicable in terms of the object's/system's sub-atomic constituents alone.

I will not be going into any great detail about either neuroscience or quantum mechanics in what follows, for I am an expert on neither and, moreover, it would lead us unnecessarily far astray from the point of the argument. But thankfully the argument can be made without having to plumb the depths of these densely technical areas of modern science. The basic thrust of the argument is that emergent properties seem to pose a serious problem for punctal nihilism. At base, the problem is that punctal nihilism is very much a reductive theory in spirit. It aims to reduce all purported macroscopic objects and events to

[229] For an excellent collection of essays on emergence see Beckerman, Flohr, & Kim, (eds.) (1992). See also, Clayton & Davies (eds.) (2006).

objects and events at the microscopic scale.[230] Indeed it says that the only things there are exist at the microscopic scale. But reductionism and emergence clash. This is because emergent properties resist reduction – that is part of the very nature of emergence. Emergent properties are properties that emerge only at specific levels of ontology (specifically, higher-levels than the microscopic level). But punctal nihilism doesn't admit different 'levels' of ontology. It posits a single level: there are punctal simples and nothing else. It seems that emergent properties and punctal nihilism are incompatible.

The problem that emergent properties pose for punctal nihilism is quite straightforward. Punctal nihilists explain the alleged existence of composite objects and their alleged properties, purely in terms of punctal simples. But if certain properties of an alleged composite object are taken to be emergent, then they *can't* be explained purely in terms of the simples that allegedly compose that object. So emergent properties have no place in a punctal nihilist ontology – for there are no candidate objects in that ontology to instantiate them.

The possibility of emergent properties, and the problem it causes for punctal nihilism, is analogous in many respects to the possibility of gunk, and the problem that follow from it. I.e. since punctal nihilism is (usually) taken to be *necessarily* true by its proponents, then the very possibility of emergent properties is enough to undermine it. Thus the punctal nihilist, it seems, will have to deny that emergent properties are even possible. Furthermore, just as in the gunk case, the monistic nihilist is not susceptible to the argument in the way the punctal nihilist is. Since the monist's explanations of worldly phenomena are always 'top-down', then an alleged *emergent* sub-world property is no more difficult to explain than *any* alleged sub-world property.

Consider an example. Take an ordinary human being – let's call him Roger. Roger has lots of ordinary human properties. He has brown hair and blue eyes; he is six feet tall; he weighs seventy-five kilos; and so on and so forth. Now both the punctal nihilist and the monist will take Roger's alleged existence, and his alleged properties, with a good pinch of salt. Neither will actually accept that he exists, as a macroscopic composite object in his own right. But both will have their own explanations of why it at least *appears* that there is a

[230] Indeed even this is generous. Strictly speaking, the nihilist does not *reduce* macroscopic entities to their microscopic constituents but, rather, she *eliminates* all macroscopic objects. But the argument is unaffected by this fact.

composite object – Roger – which has the properties described. The punctal nihilist will say that there are simples arranged Roger-wise, and that the arrangement of those simples, along with the fundamental properties they instantiate, make it such that it appears that the alleged object they compose – Roger – exists, and has the properties that are ascribed to him. In contrast, the monist will say that the world instantiates certain maximal properties that give rise to the appearance of a six foot tall, brown-haired, blue-eyed, seventy-five kilo man – Roger. The world is, in a certain spatio-temporally localised manner, *Rogerish.* These types of explanation should hopefully be quite familiar by now, and in this case are, so far at least, unproblematic.

But now let's suppose that, as any ordinary human being would, Roger instantiates certain mental properties. Specifically, let's say that Roger has the intentional property of *believing it will rain tomorrow.* Let's call this property *'X'.* Now let's suppose that the best current theories of neuroscience state that intentional mental states, like X are emergent. That is, they emerge when certain brain activity occurs, but crucially, they are not reducible to the properties and relations instantiated by the fundamental particles that (allegedly) compose the brain.

The monist can explain this situation perfectly well, whereas the punctal nihilist is going to struggle. The punctal nihilist, of course, can make reference only to the fundamental particles (punctal simples) and the properties and relations they instantiate, when providing an accurate and literally true explanation of this scenario. But *ex hypothesi,* describing the arrangement of the punctal simples, and the properties they instantiate *will not be enough* to describe the intentional mental state that is also present – since the arrangement of the punctal simples and their properties alone *do not fix* the intentional states that are also occurring. The monist, by contrast, has no problem. Let me demonstrate why this is so. Let's call the property the world instantiates that gives rise to the appearance of a composite object, *Roger,* 'world-property A' (WPA). Let WPA include all the purely physical properties we would ascribe to Roger from his hair colour and height, right down to the atomic structure and electrical activity of his brain. Now let's call the property that the world instantiates when Roger believes it will rain tomorrow 'WPB', and the property the world instantiates when Roger does not believe it rain tomorrow 'WPC'. It seems evident that there is no logical obstruction to the world simultaneously instantiating either WPA and

WPB, *or,* WPA and WPC.[231] And this shows that the monist's explanation is perfectly consistent with the thought that certain intentional states (such as *believing it will rain tomorrow*) are emergent, i.e. are not reducible to the microphysical activity occurring in the brain.

Overall, there is a strong conclusion and a weak conclusion that one can draw from these remarks. A strong conclusion would be that there *are* emergent properties, and thus since punctal nihilism is incompatible with such properties, it must be false. A weaker conclusion would be to say that emergent properties seem to be at the very least a genuine possibility. Since punctal nihilism is incompatible with the existence of such properties, it rules out even the possibility that there may be such things. This should count against it as a theory, particularly because it seems most plausible that emergent properties are *at the very least* possible. This is made even more apparent when contrasted with monism, which has no problems accommodating emergent properties. I think the weaker conclusion is probably the more prudent of the two, but either way, the moral seems clear. If you want to be a nihilist but you don't want to rule out emergent properties, then you should be a monist.

### *§5.5.1 A Potential Response*

There is, I think, a way in which the punctal nihilist could respond to the problem of emergence. Specifically, I think that the question the punctal nihilist should consider is this: why should the existence of emergent properties imply the existence of composite objects? Just because new *properties* emerge at certain levels, this doesn't seem to entail that new *objects* emerge also. Why can the nihilist not say that when simples are arranged in particular ways they *collectively instantiate* new properties which are not reducible to the properties each simple instantiates individually?[232] To return to our hypothetical character, Roger, for instance, the nihilist has no need to posit a new, composite entity (Roger's Brain, say) to instantiate the emergent property *X,* but rather, she can just state that the particles arranged brain-wise collectively instantiate *X.* There is no *logical* obstruction to this proposal, as far as

[231] There is, of course, a logical obstruction to the world instantiating WPB and WPC simultaneously, since even an illusory object cannot be both *p* and *not-p,* but this is by the by.

[232] This solution was mentioned to me in conversation by Ross Cameron. To my knowledge, it has not been explored (let alone defended) in print, apart from a few cursory remarks from Ted Sider. See Sider (forthcoming), 27

I can see, so it is worthy of consideration. But it is certainly in need of a little clarification. In particular, the notion of *collective instantiation* needs some explanation, since it is this notion that is doing all the work which the nihilist requires.

The idea is essentially that instantiation need not be a one-one relation.[233] That is to say, it needn't always hold between a single object and a single property. Rather, we should view instantiation as a many-one relation, in that many objects can, *collectively*, instantiate a single property. One may say, for instance, that the *x*s collectively instantiate *F-ness.* Now it is important to note that in saying this I do not mean that the collection, or aggregate, or set, (or any other collective term) of *x*s instantiates *F-ness.* And neither do I mean that each of the *x*s instantiates *F-ness* individually. Rather, the *x*s collectively instantiate *F-ness.* To help explain, there is a potential analogy to be drawn here with the thesis that *composition is identity.* As we explored in chapter 3, those who endorse CAI endorse a somewhat different relation of identity than that which is classically recognised. They take identity to be a many-one relation. So when one says that a composite object is identical to its parts, one is not saying that it is identical to the *sum* of its parts (for that is just a case of self-identity), nor that it is identical to *each* of its individual parts. Rather, one is saying that it is identical to its parts *collectively.* The same kind of idea is being employed here, but concerning the notion of instantiation, rather than identity.

One will also remember from chapter 3 that there are some significant concerns with the claim that identity can link many things to one thing. Most importantly, it was suggested that CAI was simply incoherent, since it would involve objects simultaneously instantiating incompatible properties. In contrast, however, many-one *instantiation* does not immediately appear incoherent. Intuitively, many-one instantiation is no way near as objectionable as many-one identity. To illustrate, consider an example. Suppose there was a large crowd of people, each of which were humming a single note at a fairly modest volume – a volume which on its own, no-one could reasonably call loud. Now if there were enough people engaged in this activity, it is reasonable to think that the resultant sound would be of a volume sufficient to be reasonably called loud. In such a situation one may think it appropriate to say that the crowd is *being loud*. Now of course, one may be a little suspicious

[233] In fact, I don't think instantiation is a relation at all. However, for present purposes I will call it a relation. It won't affect the argument.

of admitting a new entity – a crowd – to which the property *being loud* is being attributed. It is not a single entity, after all, which is making the noise – it is the individual people, each making their own noise. But as was originally stipulated, no individual is *being loud.* In such a case, then, it may seem appropriate to say that the people *collectively instantiate* the property *being loud,* even though no single individual is *being loud.*

A few qualifications are needed. Firstly, I am not for a minute suggesting that the crowd's *being loud* is an emergent property. For it is quite clear that the cumulative volume is reducible to the sum of the volumes of each individual hum. Neither am I suggesting that this is some kind of paradoxical situation which one can only avoid if one countenances many-one instantiation – it is clearly not. All I am aiming to show, however, is the much more modest conclusion that it is at least coherent (or perhaps even more timidly, it is at least not obviously *in*coherent) to suggest that many-one instantiation could be invoked to explain such a situation. The positing of many-one instantiation does not *jar* with us in the way that positing many-one identity does. The point is, then, that one should at least be open to accepting it as a possibility.

The nihilist, if equipped with a coherent notion of *collective instantiation,* has a clear route to solving the problem of emergent properties. Sure there are emergent properties at the macroscopic level, she will say, but that doesn't mean that there must be some macroscopic entities there to instantiate them. Rather, they are instantiated collectively by microscopic simples. So a nihilist could admit that intentional mental states, say, are emergent properties, but merely maintain that they arise when simples arranged in a particular way *collectively instantiate* consciousness. No extra entity is needed. Nihilism and emergence are quite compatible.

I think that this response is worth taking seriously. Indeed I take it seriously enough to concede that it weakens the original argument significantly. But I maintain that it does only weaken it, not destroy it. For there are certain costs and concerns associated with this response that one can avoid entirely if one opts for monism over punctal nihilism, and that in turn gives one reason to *favour* monism over punctal nihilism. In particular, one has to endorse this new and unfamiliar notion of *collective instantiation.* And that is *in addition* to the standard notion of singular instantiation, since the punctal nihilist will still have to make use of that. So whilst the punctal nihilist may be able to admit the possibility of emergent

properties by employing this response, in so doing she is adding some extra ideological baggage to her theory – baggage that the monist has no need to take along.

This extra baggage *does* represent a theoretical cost, although just how costly one takes it to be may vary. At one extreme, the proponent of collective instantiation may assert that she is not introducing a new relation at all, but merely using an existing relation (instantiation) in a new way (by relating many things to one property), in which case it does not look costly at all. But I don't think this is very convincing. Proponents of CAI have tried to make a similar move by asserting that they are employing one and the same relation of identity, but just using it to link many things to one thing. But the same response is available in both cases: by employing the relation in this way you are twisting my *very understanding* of what the relation is, to an extent that it no longer looks like the same relation at all. At the other extreme, one may reject collective instantiation as incoherent (in the same way that CAI is often rejected as incoherent) on the grounds that one simply cannot understand how many disparate things could instantiate *collectively* one and the same property. I don't think this view is very convincing either, however, as it just appears overly obtuse. For as I mentioned earlier, collective instantiation just doesn't seem to be as obviously incoherent as CAI. Overall, however, collective instantiation (or any other notion that is invoked to do the same work) surely represents some kind of theoretical cost.

I think that this notion of collective instantiation is a fairly revisionary one, and as such, should be considered a *significant* theoretical cost. To impress why this is, I want to draw a comparison between collective instantiation on the one hand, and collective (or plural) predication on the other. Collective predication is commonplace in ordinary language. We may say, for example, "the soldiers surrounded the building". The predicate *surrounded the building* is collective and non-distributive, for it applies to *all* the soldiers taken together, but doesn't apply to any of them individually. (A single soldier cannot surround a building). Because of this, then, the predicate is also *irreducible*, because it can't simply be reduced to a conjunction of each soldier satisfying the predicate (none of them satisfy the predicate!). But there is nothing at all mysterious about such predication; we can make perfect sense of it.[234] So a natural thought may be that collective instantiation is just the

[234] Although it does raise certain logical issues, in that the standard logical apparatus of quantifiers and predicates struggles to deal with it. See McKay (2006) for an interesting study into this topic.

same. Just as we have no problem in understanding what it is for some soldiers to (collectively) surround a building, we should have no problem in understanding how some particles (collectively) instantiate a property.

I think this analogy fails. Whilst one should accept that plural predication is irreducible in the sense espoused above, one should also note that plural predicates are satisfied *in virtue* of singular predicates being satisfied. That is, the reason the soldiers collectively satisfy the predicate *surround the building* is that each of the individual soldiers satisfies some normal singular predicate (e.g. is located here, is located there, etc.). So although the predicate is irreducibly plural, it's obtaining holds in virtue of singular predicates and states of affairs. The same cannot be said for collective instantiation, however, or at least, not when the collectively instantiated property is emergent. If some *x*s collectively instantiate a property *Y*, and *Y* is taken to be emergent, then that instantiation does not obtain in virtue of the singular properties that each of the *x*s instantiate. This is because, *by stipulation*, *Y* is emergent, so cannot be explained solely in terms of the properties instantiated by the individual *x*s. It is irreducibly collective, and it resists further explanation. What this shows, therefore, is a stark disanalogy between collective predication and collective instantiation. Collective instantiation is much more revisionary than it might initially seem, thus its acceptance should not be taken lightly. It may be able to hand the punctal nihilist an escape route from the problem of emergence, but it certainly doesn't come free of charge.

## §5.6. The Argument from Parsimony – Take I

Parsimony is often touted as a theoretical virtue in metaphysics. That is to say that when faced with a choice of theories, those which are more economical or simple in some regard are often preferred precisely because of that economy or simplicity. The number of entities a theory posits, or the number of *kinds* of entities it posits, or the number of primitive terms it employs, are all respects in which simplicity is often valued. Parsimony appeals to those philosophers who, like Quine, have a taste for desert landscapes. I should state from the outset that I am uncertain as to how much weight parsimony should be afforded when weighing up competing theories, and so I am somewhat sceptical of arguments which

appeal directly to it. This is not to say that I reject the principle of Occam's Razor outright, but merely that I think it can be significantly blunted through the consideration of other significant factors.

A fairly standard way of expressing Occam's Razor could be with the following conditional: if all other things are equal, then the simpler of the competing theories should be preferred. This is a perfectly reasonable principle but it is one that, in practice, will seldom come into play. The reason for this is that very rarely are all things equal between two theories other than the number of things they posit. That is, the caveat in the antecedent of the conditional is *crucial*, yet often overlooked. My point here is that if there are other significant factors in play when weighing up competing theories, it is not at all clear how much weight we should give to concerns of parsimony.

However, the fact remains that parsimony *is* often touted as a theoretical advantage. So it should be worthwhile for the monist to point out that *if* one has a taste for desert landscapes, one will most likely have a taste for monism too, for when it comes to parsimony, monism is right at the front of the pack in a number of different regards.

The most obvious way in which monism is parsimonious is down to the number of concrete objects it posits. The monist posits only a single concrete object – far fewer than most metaphysical theories. Indeed the only way one could claim to be more parsimonious in this respect would be to posit no concrete objects at all, and adopt some kind of complete nihilism when it comes to material things.[235] Of course, the vast majority of theories will posit more than a single concrete object, and indeed most of those will posit a *great deal* more.

To re-concentrate our focus on the issues currently at hand, one will see that compositional nihilism, in any of its guises, will be much more parsimonious with respect to concrete objects than universalism. For in addition to the simples posited by the nihilist, the universalist will posit many more composite objects which those simples are said to compose.[236] But of course these types of parsimony concerns can also be divisive *within*

[235] See Turner (2011) for an interesting investigation into the possibility of such a theory. Ultimately, Turner concludes that it fails. See also Hawthorne & Cortens (1995).

[236] Note that a universalist who doesn't countenance simples at all (i.e. who thinks matter is gunky) will posit an *infinite* number of concrete objects, making her theory score very poorly on the parsimony scale. At the other end of the scale, however, one *could* be both a monist *and* a universalist. I.e. accept the universalist principle that for any two or more objects there is a further object which they compose, but maintain that there is in fact only a

nihilism. The punctal nihilist, for instance, will (presumably) posit a vast quantity of point-sized simples, whereas the monist posits only one (admittedly, much larger) simple. So one may be tempted to appeal to this fact in order to argue for monism over punctal nihilism.

As it happens, Terry Horgan and Matjaž Potrč, argue for monism in exactly this fashion.[237] As a result of various theoretical considerations, Horgan & Potrč claim that there are only three viable theories which meet the criteria they set out for a satisfactory ontological view. Those three theories are punctal nihilism, monism and universalism.[238] Having narrowed down their options to three, they make the following claim:

> These three candidates can be ordered with respect to comparative ontological parsimony. The simplest is blobjectivism [monism]; it maximizes ontological parsimony by countenancing only one real concrete object, the blobject [the world].[239]

Horgan and Potrč do recognise that arguments resting on claims of ontological parsimony such as this one will not be entirely decisive, because parsimony represents an issue on which "reasonable people can reasonably differ".[240] However, they claim to be persuaded that concerns about ontological parsimony clearly show nihilism to be favourable to universalism, and among competing nihilisms, swing the balance in favour of monism as opposed to punctal nihilism.

Despite the fact that Horgan & Potrč's argument supports my own view, I think there are reasons to consider it dubious. First of all, one could question the value of the *type* of parsimony they invoke. Focusing only on reducing the number of *tokens* of a particular type of entity is to focus on what David Lewis would have called *quantitative parsimony.* This is to be contrasted with *qualitative parsimony,* which is concerned with reducing the number of *types* of entity posited by a theory.[241] Horgan and Potrč seem concerned solely with quantitative parsimony, since both monism and punctal nihilism posit the same type of objects (concrete simples), but the former posits fewer tokens (one) of that type than the

---

single object. This would make both theories equal with regard to parsimony, but it is hard to see why anyone would endorse such a view, so I will ignore it from here on.

[237] See Horgan & Potrč (2008), 182-9. Horgan & Potrč are the only contemporary philosophers of whom I am aware, that defend monism.

[238] Horgan & Potrč (2008) use different, and somewhat zany, terminology for these views (Snobjective noncompositionalism, Blobjectivism, and snobjective universalism, respectively), but their content is the same.

[239] Horgan & Potrč (2008), 183

[240] Horgan & Potrč (2008), 190

[241] Lewis (1973), 87

latter (many). According to Lewis, however, the value of quantitative parsimony is scant, and it is only qualitative parsimony that is to be considered a theoretical virtue. This view is shared by many contemporary metaphysicians, such that one could say it is the default view. If the default view is correct, then, Horgan & Potrč's argument is not going to carry much weight, since it appeals only to quantitative parsimony.

I do not subscribe even to the default view. I think there are cases where quantitative parsimony does seem to be of at least *some* value. These cases tend to be those in which two (or more) theories are seemingly identical in all respects *except* the number of objects they posit. Daniel Nolan demonstrates this by asking us to consider some examples. The most vivid example he offers concerns the empirical data, and subsequent reasoning, that led scientists to posit the neutrino.[242] To sum up very briefly, scientists noticed that during the process of beta-decay, whereby the nuclei of radioactive atoms emit electrons, the loss of energy of the nucleus was greater than the energy accounted for by the electrons emitted. It seemed as though some energy (albeit a truly miniscule amount) was literally evaporating into nothingness. The problem was, of course, that this would contravene the principle of the conservation of energy – the principle which states energy cannot be created or destroyed, but only transferred. The ultimate solution was to postulate an extra particle – the neutrino. The idea was that as well as emitting an electron during beta-decay, these radioactive nuclei must also emit a neutrino, and it is that which accounts for the discrepancy in energy. Thus even though it was never physically detected, the neutrino was borne.

Nolan's point is that whilst the theory postulates a *single* neutrino to account for the slight energy discrepancy for each electron emitted, it could just as well have postulated many more. Let us say that the discrepancy in energy for each electron emitted is $x$. There is nothing stopping the theory from stating that with each electron emitted there are *twenty* neutrinos emitted, each with energy of $x/20$. Or indeed any other number, $n$, of neutrinos, each with energy of $x/n$. But of course the theory does *not* state this – it postulates a single neutrino for each discrepancy of energy. Despite the fact that it would be perfectly coherent to posit 618 neutrinos for each electron emitted, it would be theoretically gross. And the same surely goes for any version of the theory other than the one which posits the minimum amount of neutrinos required to fill the explanatory gap – and that amount is one.

[242] Nolan (1997), 332

It is hard to know how to *argue* for such a claim, indeed I don't claim to able to do so. Yet it strikes me as patently obvious that it is true. The postulation of a new particle to explain the discrepancy in energy is an ingenious one, but to posit *more* than one would be gratuitous and unacceptable. What this suggests is that quantitative parsimony *is* in fact a theoretical virtue, and that Occam's Razor *is* in fact a sensible principle to adhere to. Nolan's example is a rare case in which all other things *are* equal, and the only way to distinguish the possible theories is by the number of entities (in this case neutrinos) they posit. In such a case it just seems obvious that the most quantitatively parsimonious theory is the one which should be accepted. This in turn suggests that quantitative parsimony is to be valued.

So I disagree with Lewis and others who claim that quantitative parsimony is of no value. However, I still find Horgan & Potrč's argument unconvincing. The reason for this is that the case over which Horgan & Potrč argue (the very same case that is the subject of this chapter) is markedly different from the case of the neutrino highlighted by Nolan. It is clearly not the case that punctal nihilism and monism are equal in all respects other than the number of simples they posit. The two theories are vastly different in a number of ways. An obvious way in which they differ lies in the *nature* of the simples each theory posits. The simples differ in size for a start, for on one account they are minimally small (point-sized) and on the other they are maximally large (world-sized). And they are surely going to differ in complexity. Since the monist claims there to be only a single concrete simple, it must surely exhibit enormous structural complexity if it is to give rise to the rich qualitative variation we experience in the world. The point-sized simples of the punctal nihilist, however, will presumably be bland and uncomplicated in comparison, or so it would seem reasonable to think. These differences alone suggest that there is a lot more to separate these two theories than merely the number of entities they posit. Indeed the clear difference in nature of the posited simples may even suggest that we are not even counting tokens of the same type of thing at all. Yes they are both subsumed under the categories of *concrete object* and *mereological simple*, but their evident differences may suggest that a more fine-grained method of distinction may be required to accurately categorise them. The simples posited by the two nihilisms look like very different beasts. In any case, it seems that when focusing on plain quantitative parsimony, Occam's Razor just cannot be applied to separate punctal

nihilism and monism, because its crucial caveat is clearly not observed; all other things are *not* equal.

### *§5.6.1 A Brief Digression on Quantitative Parsimony*

To digress slightly from the issue at hand, I think that there is a little more that could be said here about the value of quantitative parsimony *in general.* I previously claimed that certain examples, like that given by Nolan about the neutrino, can demonstrate that quantitative parsimony *can* be considered a theoretical virtue, albeit one that is easily outweighed by other factors. But this is not the only way to interpret the case. One could maintain (and indeed, on reflection, I think this is probably the more sensible option) that quantitative parsimony is *not* in fact a theoretical virtue in the sense that it has any independent, intrinsic value, but rather, that it can act as a tie-breaker in otherwise inseparable cases. In the example of the neutrino, as presented by Nolan, there was quite literally *no difference at all* between a theory which posited one neutrino and a theory that posited more, other than the number of neutrinos posited. And in such a case it seems entirely reasonable to invoke quantitative parsimony in favour of the former theory. But is this parsimony of any *real theoretical value,* or is it just an intuitive way of breaking the stalemate? I think perhaps it is the latter which is correct.

To illustrate, imagine an example very similar to Nolan's, in which the only difference is that the energy discrepancies noted during beta-decay varied in magnitude from time to time. (Let's also say, just to keep the example simple that the discrepancies were always precise multiples of $x$, where $x$ was the minimum energy discrepancy noted). Now there are various ways in which one could explain this example. On one hand, one could say that in each case the nucleus emits a single neutrino, but maintain that neutrinos must vary in energy levels from case to case. On the other hand, one could say that neutrinos have fixed energy levels (i.e. $x$), but different numbers of them are emitted from case to case.[243] It's not obvious which of these explanations would be better here, although my intuition pulls me towards the latter, but what does seem clear is that parsimony doesn't seem of much

[243] Alternatively, a third explanation would be that there are different *types* of particle being emitted from case to case, each of which have varying energy levels, but for simplicity, I will consider just the two possibilities.

importance. It would seem strange to opt for the former explanation purely because it posits fewer neutrinos than the latter, for it just doesn't seem to hold any real advantage. And this suggests that quantitative parsimony is perhaps not of any intrinsic theoretical value in itself. Furthermore, on the face of it at least, neither theory looks obviously more likely to be true than the other, suggesting that quantitative parsimony provides no guide to truth. Indeed this is even the case in the original example by Nolan. For even if positing two neutrinos rather than one seems somewhat gratuitous, we have no reason to believe (from the data given) that it is less likely to be true.

But I digress too much. These remarks hopefully emphasise some interesting issues that surround the notion of quantitative parsimony and what value, if any, it should be afforded. But I will not press the issue any further here. What remains the case though is that whichever way you look at it – i.e. whether you take QP to be a genuine theoretical virtue, yet one which is easily outweighed, or whether you deny it any intrinsic value at all and maintain its maximum application should be as some kind of theoretical tie-breaker – quantitative parsimony should have little effect on the debate between punctal nihilism and monism. For if you take the former view, you will see that there are all sorts of other factors in-play that will outweigh its importance, and if you take the latter view you will see that there is no tie which requires breaking. Quantitative parsimony, then, is of little use here.

## §5.7. The Argument from Parsimony – Take II

Quantitative parsimony is not the only type of parsimony, of course. It has already been noted, for instance, how David Lewis distinguished between quantitative parsimony and qualitative parsimony; between reducing the number of *tokens* of things on the one hand, and the number of *types* of thing on the other. More recently, attention has been turned to what has been called *ideological* parsimony: the aim of reducing the number of primitive terms one's theory employs.[244]

In what follows I will outline how parsimony can be invoked to support monism over any other rival version of compositional nihilism. I think that the type of parsimony I

[244] See Sider (forthcoming), & Cowling (forthcoming).

will appeal to could fall under the head of ideological *or* qualitative parsimony, but either way, its results will be of far greater success than the appeals to quantitative parsimony considered in the previous section. Very broadly, I will argue that the monist can do away with all *relations.* Relations will not appear in a monistic ontology, and they will play no part in a monistic explanation of reality. This disposal of relations represents a significant theoretical saving, I claim, and thus makes monism a much more attractive view. As to whether this is an appeal to qualitative or ideological parsimony, I repeat that I am unsure. For relations may well be viewed as a *type* of entity, or perhaps a particular type of property, in which case dispensing with them would seem to represent a qualitative saving. But it also seems that theories which invoke relations will have to make use of irreducibly relational language in their fundamental descriptions of reality, which suggests that dispensing with them represents an ideological saving. Perhaps it would involve a saving of both types? And perhaps that, in turn, would increase the force of the argument? Once more, I am unsure. But what *does* seem clear is that however one wants to categorise it, dispensing with relations represents a significant and valuable theoretical saving.

My argument will proceed in two parts. In one part I will argue that on a monistic view, all and any relations there may be are *internal.* In the other, I will argue that internal relations are not real, since they represent nothing over and above the existence of their relata. As such, I will conclude that there are no (real) relations at all. But I will present the parts in what may seem like reverse order, firstly explaining why out of external and internal relations, only the former should be considered real things, and secondly by explaining why on a monistic picture there are no external relations.

Firstly, however, I should probably note, as the reader may well have done already, that there are many similarities that could be drawn between the argument I will be presenting here, and a famous argument (or perhaps more precisely, a famous *family* of arguments) presented by F. H. Bradley.[245] Bradley articulated a number of problems concerning relations, most famously, the eponymously named *Bradley's Regress,* to argue for his thesis of monism. However, whilst it would certainly be an oversight on my part were I not to mention these similarities at all, I would like to distance myself somewhat from Bradley. The main reason for this is that it is not *precisely* clear as to how Bradley's

[245] Bradley (1930)

arguments are meant to work, nor indeed what his thesis of monism actually consists in. His writing is notoriously obscure, and interpretations of his work vary widely. In fact it has been said of Bradley that "obscurity came to him naturally, and it is hard to resist the suspicion that he also exploited it as a weapon in argument".[246]

It is hard to be certain, for instance, whether the distinction he employs between internal and external relations is of exactly the same nature as that distinction which is commonly accepted.[247] Moreover, it is not certain whether he was trying to claim that all relations are internal, as some interpreters suggest,[248] or whether there are in fact no relations at all.[249] Perhaps most significantly, it is not even clear whether the type of monism Bradley was arguing for was the same monistic thesis that I am arguing for here. It has recently been argued, and in my view quite plausibly so, that Bradley's monism is not a mereologically nihilistic, Parmenidean monism, but rather, one that admits a multiplicity of things, but merely maintains that none of them are entirely distinct or separated from one another or from the whole in which they are said to inhere. If this interpretation is right, then Bradley's monism is more in line with the thesis put forward by Jonathan Schaffer – *Priority Monism* – a thesis quite distinct from the one being proposed here.[250] So whilst I recognise the Bradlean roots from which my argument stems, I will make no further mention of his work in what follows. This is for the sake of clarity. After all, the present goal is to argue for my thesis of monistic nihilism, not to become entrenched in a project of historical or philosophical interpretation. So let me get back to the argument.

To begin, it should be worthwhile to get clear on the precise nature of the distinction between internal and external relations.[251] Following the likes of G. E. Moore, I take the distinction to be best characterised as being *modal* in nature.[252] That is, if *R* is an internal relation that relates *a* to *b,* then the very existence of *a* and *b necessitates* the fact that they stand in *R*. Or, to put it another way, it would not be possible for *a* and *b* to exist and *not* stand in *R,* if *R* is internal. To make this somewhat clearer, consider the following examples:

---

[246] Wollheim (1959), 11

[247] Some have suggested that he didn't even recognise that there was such a distinction at all. See Vallicella (2002), 5

[248] See, for instance, van Inwagen (2002), 34

[249] See Vallicella (2002), 5

[250] See Schaffer (2010a; 2010b; 2010c)

[251] For an excellent recent discussion on relations, which includes a section dedicated to the internal-external distinction, see Heil (2009).

[252] See Moore (1919),

1) 3 is greater than 2

2) Brian is sitting next to Freddie

According to the modal characterisation of the distinction, then the relation referred to in 1 should be considered *internal,* and the relation referred to in 2 should be considered *external.* The reason for this is that the very existence of 3 and 2 *necessitates* the fact that the former is greater than the latter. There is no possible world in which 3 is *not* greater than 2 (or so I will assume). Conversely, however, the mere existence of Brian and Freddie is not enough to necessitate the fact that they will be sitting next to one another. For it is perfectly possible that both Brian and Freddie could exist yet were *not* sitting next to one another. There may be someone sitting between them, for instance, or they may not even be sitting down at all. So this is what I take the distinction between internal and external relations to consist in. An internal relation is one whose obtaining depends *solely* on the existence of the entities it relates; if the relata exist, then the relation *must* obtain.

Having clarified the nature of the distinction, I should now explain why I will assume that out of internal and external relations (if indeed there are any such things at all), it is only external relations that should be considered real. I think the best way to proceed here is to frame the explanation in terms of *truthmakers*. There is a popular school of thought, largely Quinean in spirit, which claims that the only entities we need accept as existent are the truthmakers required to make true sentences true. For instance, take the following sentence

(1) 'Tables exist'

On this view, for sentence (1) to be true, there needs to be some existent thing (presumably a table)[253] whose existence *makes it true.* And whatever it *is* that makes the sentence true, one is *ontologically committed* to, if it is true. But another common addition to this view is that one is only committed to the *minimal* truthmakers required to make a sentence true. For instance, say universalism was true, and there was some object composed of a table, a lion, and the

[253] Although not necessarily. See Cameron (2010b). Cameron claims that sentences like 'there are tables' can be strictly and literally true even if there are, in fact, no tables. The truthmakers for sentences such as these, he claims, need not be *direct* referents of the nouns employed. If there are simples arranged table-wise, for instance, then this may be enough to make true the sentence 'there are tables'.

top half of the moon. Now the existence of that object would be *sufficient* to make (1) true, since it includes the existence of a table. But the truth of (1) in no way commits one to the existence of such an object (that would be absurd!), for it is not the *minimal* truthmaker required to ensure (1)'s truth. This qualification on the requirement for truthmakers stops one from generating a gratuitously profligate ontology just from analysing what sentences are true. To see why, consider:

(2) 'The average mother has 2.4 children'

Despite the somewhat unusual nouns that (2) employs, its truth does not commit one to the existence of some object to which those nouns (attempt to) refer, i.e. an average mother, or four tenths of a child. Rather, for the sentence to be true, it is required only that there are a certain number of mothers and a certain number of children, and that the latter number divided by the former number gives 2.4. These mothers and children alone are the minimal truthmakers required for (2) to be true, and therefore, they are the only entities that one is committed to believing in *because* (2) is true.

Getting back to internal and external relations, consider the following two sentences:

(3) '2 is greater than 1'

(4) 'Brian is sitting next to Freddie'

Let's take (3) first. As I have already shown, it refers to an *internal* relation. (At least, it does so according to the modal characterisation of the distinction that I am endorsing). Moreover, it certainly expresses a truth, or so I will assume. But what *makes* it true? What are the minimal truthmakers required for such a sentence to be true? The obvious answer appears to be the numbers 1 and 2 themselves. As I have already stated, the mere existence of the two numbers alone is enough to guarantee the truth of (3). It would not be possible for them to exist and *not* be so related; if you have the numbers, you have the relation. Thus there is no need to posit any relation or relational property, or any other such thing, to make (3) true; the numbers alone do all the required work. So whilst it is correct to say that 2 is greater than 1, we are not committed to the existence of any relational entity in so doing. And it would appear that the same holds for any putative internal relation. Any true sentence that

purports to refer to an internal relation does not commit one to the existence of such a relation, but in fact, it commits one only to the existence of the relata.

By contrast, consider (4). For this sentence to be true, it would seem that we require more than the mere existence of Brian and Freddie. For as was noted earlier, they could exist yet have someone sitting between them. Indeed they may not be sitting down at all. Brian and Freddie alone are not sufficient to constitute the truthmakers for (4). So for (4) to be true, there needs to be something else, something *external* to Brian and Freddie, which will *guarantee* its truth. There are a number of candidates for what this thing could be. It could simply be the relation *sitting next to,* which obtains between the two men. It could be a fact or a state of affairs, i.e. the fact that Brian is sitting next to Freddie, or the state of affairs of Brian's sitting next to Freddie.[254] It could perhaps be a relational 'moment' or 'trope'.[255] But whatever candidate one plumps for, one will be ontologically committed to it. Whatever explanatory route one takes, if (4) is true then its ontological commitments are greater than just Brian and Freddie. And once again, it seems that this conclusion will apply to any putative external relation. For the very nature of what an external relation is entails that the existence of the relata alone is not enough to guarantee the existence of the relation. So external relations, if there are any such things, represent a very real ontological commitment. Therefore, any metaphysical theories in which external relations play a part will have to accept this commitment.

From this analysis we can conclude that out of internal and external relations, it is only external relations that can be considered real entities. Internal relations are, in the words of D. M. Armstrong, "not ontologically additional to their terms".[256] It is solely because of the existence of 1 and 2 that the latter is greater than the former, there is no extra entity, no extra thing involved in this being the case. We can *talk* of internal relations obtaining or not, but there are no *things* to which we refer when we so talk; there are but the things which we talk of them relating. External relations, by contrast, *are* ontologically additional to their terms. It is not just the existence of Brian and Freddie that makes it the case that one is sitting next to the other; some extra entity is needed (a relational entity of some sort) to make it the case. So if there are any relations – *any real relations* – they must be

[254] See Armstrong (2004) for a defence of this response.
[255] See Mulligan, Simons and Smith (1984) for a defence of this response.
[256] Armstrong (1997), 12.

external relations. And from this we can conclude that if it were the case that there were *no* external relations, then it would follow that there were no relations at all.[257] In what follows, I will show why pluralists are committed to the existence of external relations, whereas monists are not. The monist can do away with relations entirely.

Before moving on, I would like to quickly pre-empt and dispense with a potential objection to the above reasoning. In the preceding comments I readily accepted the Armstrongian thesis that internal relations are 'nothing over and above' their relata. Yet the reader may remember that in Chapter 3, I rejected Armstrong's claim that mereological fusions are 'nothing over and above' their parts. One may take this to create some tension in my views, so I would like to say a few words aimed at alleviating that tension. There is a crucial difference between the connection between internal relations and their relata on the one hand, and the connection between mereological fusions and their parts on the other. The difference is that the former are logically inseparable whereas the latter are not. What I mean by that can be summed up in the following two statements:

1. If the numbers 1 and 2 exist, then the sentence '2 is greater than 1' is true of *logical necessity*.
2. If three simples (*x*, *y*, and *z*) exist, then the sentence 'there is an object which *x*, *y*, and *z* compose' is not true of logical necessity.

Internal relations simply, and logically, 'fall out' from the existence of their relata; you cannot have the latter and not have the former. Mereological fusions do not simply 'fall out' from the existence of some particles. One needs to believe that there are such things as mereological fusions to begin with. That is why it is consistent for one to accept that internal relations are no addition to one's ontology, yet reject the same claim about mereological fusions. So on with the argument.

I claim that positing multiple concrete objects commits one to the existence of external relations. To demonstrate this, consider a very simple world consisting of just three simple particles, *a*, *b*, and *c*. Let's suppose that they are arranged in a line as depicted below:

[257] Providing, of course, that the internal-external distinction is exhaustive, i.e. there can be no relations which are neither internal nor external. I will assume this to be true.

*Fig. 4.*

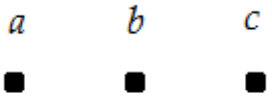


Now in a world like this, where there are multiple concrete objects, there will always be spatial relations that hold between those objects. In this world, for instance, there will be distance relations between the particles. Furthermore, the distance *ac* is twice as great as the distance *ab*. But this is surely an external relation![258] The mere existence of *a*, *b*, and *c* is not enough to guarantee that this relation will obtain. After all, *b* and *c* could have been further apart (or closer together). External spatial relations like this will occur, I claim, in *all* worlds where there are multiple concrete objects. The only way around it I can see would be to suggest that all concrete objects at a world have their particular spatial locations *of necessity*. But that would be a quite bizarre view to hold – I cannot think of anything that would possibly motivate it. It is for this reason that I claim ontological pluralism, in any form, is committed to the existence of external relations.

The more difficult part of the argument is to show that the monist is *not* committed to external relations. In fact, I'm not sure that it can be shown, *conclusively*, to be true. However, my claim that the monist is not committed to external relations will take the form of a challenge rather than a demonstration; a challenge to any dissenters to provide an example of an external relation that the monist *is* committed to. I am not sure that such a challenge can be met – I cannot think of any suitable examples. I will say a little more in what follows to show why I don't take the monist to be committed to external relations, but I will leave the challenge open for the reader to tackle at his or her leisure.

My argument begins with the following, and in my view, plausible, premise: for any external relation to obtain, there must be at the very minimum, two distinct relata.[259] Since

[258] Of course, there is a sense in which one may think it is an *internal* relation. Let's say for instance that the distance between *a* and *b* and between *b* and *c* is 1 metre, thus the distance between *a* and *c* will of course be 2 metres. Now I am not saying that the relation of *2 being twice as great as 1* is an external relation, for it is clearly not. What I am saying, however, is that it the spatial positioning of each of the particles is surely a contingent feature of this world. They could have been in different positions, and thus different spatial relations could have obtained between them. It is in this sense that I mean to say these relations are external. The mere existence of the three particles does not entail the fact that they will be equally spaced.

[259] Relations can, of course, obtain reflexively, i.e. between an object and itself, but all these reflexive relations will surely be internal. Examples of such relations might be: *being identical to*, or *being the same shape as*, or *being in the same place as*, and so on. These are all internal relations.

the monist only posits a single *concrete* object, it follows from this that there can, on a monistic view, be no external relations that obtain between purely concrete relata. Now presuming that the abstract-concrete distinction is exhaustive of all objects (and I will assume that such a presumption is true), then on a monistic view there remain only two combinations of objects which an external relation could relate: concrete object (i.e. the world) to abstract object *or* abstract object to abstract object. Both these combinations include abstracta, and thus the question I wish to pose is this: are there any external relations that could obtain in which at least one of the relata was abstract?

Initially, it looks plausible that the answer to this question is 'no'. Consider, for instance, some paradigm examples of external relations. Spatial relations, for instance, will normally be external, since things tend not to have their spatial locations *of necessity.* But, of course, these are not the types of relation that can obtain between *abstracta.* There may not be a universally accepted and complete definition of what it takes for an object to be abstract, but it is commonly accepted that abstract objects are not spatially located. And this would appear to exclude them from participating in spatial relations.

There are, of course, other sorts of relations which could be external. There is, for instance, a whole host of what one could call 'social' or 'human' relations. '*... is married to ...*' or '*... loves ...*' are examples of social relations. Both are non-spatial and non-causal, yet both are certainly external. Now under normal circumstances, these sorts of relations would obtain between only concrete objects (e.g. Brian is married to Anita, or Freddie loves his piano), but, I suppose, there is nothing stopping them from relating abstracta, or at least, from relating a concrete object to an abstract object. Brian may love the number three, for instance, and, if the laws governing matrimony were changed in the relevant ways, he may even be permitted to marry it. But I don't think any of this should worry the monist. For one should only be committed to the existence of external relations if they are indispensible to one's description of *fundamental reality.* And I doubt that social relations of the sorts mentioned will feature in *any* plausible description of fundamental reality – pluralistic, monistic, or otherwise.

I think it plausible to claim that spatial relations will be indispensible to *any* pluralist theory's description of fundamental reality. Thus pluralism of any stripe will be committed

to spatial, thus *external*, relations.[260] Monism, by contrast, will not be so committed. The monist's description of fundamental reality will be given in terms of only a single concrete object (the world) which instantiates various monadic properties.[261] Spatial relations will not be indispensible to the monist's explanations, for they will simply not figure in those explanations, and thus she will not be committed to them. This represents a significant theoretical advantage for the monist.

## §5.8. Concluding Remarks

I hope to have shown that there are a number of good reasons to believe that monism is the most plausible form of compositional nihilism. If one believes only in mereological atoms, then it is a single, maximally large atom that one should endorse, rather than a plurality of microscopica. We have seen how there are *prima facie* problems that afflict both punctal and intermediate nihilisms, which in turn, give *prima facie* reason to suppose that monism is to be preferred. We have seen that where the possibility of gunk spells trouble for the punctal nihilist, the monist emerges unscathed. And we have seen that where the monist can easily accommodate the possibility of emergent properties, the pluralist nihilist can do so only at the expense of inflating her ideology, and endorsing notions that seem, at the very best, controversial. Finally it was shown how the monist can boast the most parsimonious form of nihilism, whether it be quantitative, qualitative, or ideological parsimony that one is concerned with. So in the battle of the nihilisms, monism wins the day. In one corner we have the Parmenidean nihilist, promoting the existence of a single world-object, and in the other we have the Democritean nihilist, proposing a plurality of atoms in the void. In light of the preceding comments I hope there remains little doubt; Parmenides has emerged victorious.

[260] I should note that there have been some recent attempts to defend the idea that even in a pluralist ontology, all relations are internal. See, for example, Mulligan (1998) or Wachter (1998). Needless to say, I think these attempts are unsuccessful.

[261] I recognise that at this stage, this seems quite vague. Much more detail will be provided in Chapters 6 & 7.

# §6. Extended Simples

If monism is true, the world is simple: it has no proper parts. So given that the world is spatially extended (which I will take to be true), monism commits one to the existence of spatially extended simples. More precisely, monism commits one to the view that all there is, is one, large, spatially extended simple: the world itself. But to many, extended simples are complete anathema. So for a defence of monism to be thorough and robust, which I sincerely hope this defence to be, it should include some independent support for the claim that extended simples are (at the very least) possible. In what follows, I will provide that support.[262] I will first identify two arguments that can be, and have been, presented against the possibility of there being extended simples.[263] I will consider each argument in turn, and show it to be wanting, and thus conclude that there is no compelling evidence to suggest that extended simples are impossible. Finally, I shall go on to consider what views of matter remain for those who reject extended simples, and show that those views are subject to a number of difficulties. As a result, I shall conclude that extended simples should not be considered anathema after all.

## §6.1. The Cartesian Argument

I believe that this argument, or at least some related form of it, provides the most common and widespread source of suspicion towards the possibility of extended simples. It is the most common, I suspect, because it is (or at least it *seems*) the most intuitive. Whilst

[262] My use of the term '*independent* support' perhaps requires a little qualification. It will be independent in the sense that I will not be invoking monism, or any theses committed to monism, in order to back up my arguments for the possibility of extended simples. Such a strategy would beg the question. However, I will be keeping the monistic project in mind throughout. As such, there are various lines of response to certain arguments that I will not be considering because they either conflict with, undermine, or are simply not relevant to, monism.

[263] These are not the only two arguments that have been given against extended simples but they are, in my opinion, the most significant. They are certainly the most prominent in the literature. Other arguments against extended simples can be found in Hudson (2006), 108-121 and Zimmerman (1996b). Space restrictions prevent me from addressing these arguments here, but needless to say, I think that they can also be overcome.

endorsed by many philosophers (although often only implicitly), the argument is perhaps most clearly and explicitly stated by Descartes, hence my naming of it. He says:

> There cannot exist any atoms or parts of matter that are of their own nature indivisible. For however small we suppose these parts to be, because they are necessarily extended, we are always able in thought to divide any one of them into two or more smaller parts, and may accordingly admit their divisibility. For there is nothing we can divide in thought which we do not thereby recognise to be divisible.[264]

An uncharitable way to respond to this argument would be to reject it outright on the grounds that it is objectionably question-begging, in that it seems to simply assume the very thing that is up for debate: that all extended objects are divisible, and thus have proper parts. But it is generally good practice in philosophy to interpret the arguments of one's opponents in the most charitable manner possible, in order to avoid accusations of merely knocking down straw men. Moreover, if one can defeat the most charitable interpretation of an argument, one should be able to defeat *any* interpretation of it. (Ironically, one can strengthen one's own position this way, by first strengthening the arguments of one's opponents). So that is what I shall do. I think the most charitable reconstruction of Descartes's argument would be as follows:

1) If an object is spatially extended, then one can always conceive that it may be divided into at least two smaller parts.
2) If it is *conceivable* for an object to be so divided, then it is *possible* for that object to be so divided.
3) It is not possible for any *simple* object to be divided.
4) Therefore it is not possible for any simple object to be extended.

[264] Descartes, R. (1934), 209. Interestingly, Descartes is not just arguing against the possibility of *extended* simples here, but against the possibility of material simples *altogether.* Thus he is arguing for the controversial view that all matter is gunky. This is a consequence of his view that all matter is necessarily extended; indeed that extension is the fundamental nature of matter. Since, for Descartes, all matter is necessarily extended, and since all extended objects are necessarily divisible, then it follows that all matter is, of necessity, divisible. The Cartesian world (or at least the Cartesian *material* world) is one made entirely of atomless gunk.

5) Therefore, there can be no extended simples.[265]

There is a certain intuitive pull about the Cartesian Argument, but I think it can, and should, be resisted. I think there are a number of different ways one can go about responding to the argument, some of which are more plausible than others. Below I outline three possible responses, leaving my preferred response until last.

## *§6.1.1 Response One: Simplicity is Contingent*

Central to the Cartesian Argument is a substantive modal claim about mereological simplicity. The claim is that simplicity is *necessary.* That is to say that *if* there are any mereological simples at all, then they are *necessarily* mereologically simple. This claim is crucial to the argument; without it, it would not go through. To illustrate why this is so, consider the following toy example. Firstly, let 'groundpig' mean any pig that is incapable of flight. (This is analogous to the assumption implicit in the Cartesian Argument that 'Simple' means any object that is incapable of being divided). Now consider the following argument:

1. If an object is spatially extended, then one can always conceive that that object could fly.[266]
2. If it is *conceivable* that an object may fly, then it is *possible* that object may fly.
3. It is impossible for a groundpig to fly.
4. Therefore, it is impossible for a groundpig to be extended.
5. Therefore, there can be no extended groundpigs.

The argument is clearly fallacious, for as we well know, there *are* such things as extended groundpigs. Indeed *all* actual pigs are extended groundpigs. So where has it gone wrong? It has gone wrong, I would suggest, at premise 3. By stating that it is *impossible* for a groundpig

[265] A similar argument to this is considered, and rejected, in both McDaniel (2007), 137, and Markosian (1998a), 223-224.

[266] Perhaps it is impossible to conceive that the *entire* world could fly? If one thought that what it took to fly was to move through the air relative to the ground, then it certainly seems impossible that the world could fly. For those gripped by this peculiar thought, however, just substitute the word "sub-world object", or perhaps "earthly object" for "object" in premise 1. The argument will be unaffected.

to fly, one is stating not only that groundpigs are incapable of flight, but actually stating the much stronger claim that they are *necessarily* incapable of flight. Now as it happens, there is no good reason whatsoever to believe that this claim of necessity is true. On the contrary, there is good reason to believe it is false. It seems fairly obvious that the evident lack of porcine aviation is a purely contingent feature of the world. It surely *could* have been the case that the air was thick with flocks of pigs on the wing, it just so happens that it isn't. But if this is so, premise 3 is false. And if premise 3 is false, then the argument is unsound. So it is here that we have identified the source of confusion in our clearly fallacious argument.

So according to this line of response, the thought is that perhaps a similar thing has gone wrong in the Cartesian Argument. Premise 3 of the argument assumes that mereological simplicity holds of necessity, and this leads to the conclusion that simple things, if there are any at all, cannot be extended. But perhaps this assumption is mistaken? Perhaps there are simple things that could have been complex? Perhaps there are complex things that could have been simple? More generally, perhaps the mereological facts that obtain are purely contingent? Here, then, we have the beginnings of a possible response to the Cartesian Argument.

First of all I should clarify exactly what *sort* of necessity we are talking about here. For types of necessity differ, and the plausibility of one's claims will depend significantly on the type of necessity one invokes. For instance, if one were concerned only with *de dicto* necessity, then premise 3 (in both the Cartesian argument *and* the toy argument) would be true. Indeed, if we take the term 'simple thing' to mean 'indivisible thing', which is a plausible assumption to make, then on the *de dicto* reading, *of course* premise 3 is true; it is an *analytic* truth. And the same goes for premise 3 in the toy argument if we take 'groundpig' to mean 'pig incapable of flight'. But of course it is not the *de dicto* reading that is of our current concern here. This is not a dispute about what certain mereological terms *mean*; it is a dispute about what certain things are *really like.* Thus it is the *de re* reading that should be the subject of our focus. In the toy argument, premise 3 makes a *de re* modal claim about groundpigs, to the effect that it could not possibly have been the case that they could fly. This claim is false (or so it would plausibly seem), and it is the falsity of this claim that spoils the argument. Similarly, premise 3 in the Cartesian argument makes a *de re* modal claim about simples, to the effect that they are necessarily simple, and thus that it could not

possibly have been the case that they were not simple. So the line of response that is currently under consideration would have to reject this claim, and instead insist that if an object is mereologically simple, then it is only *contingently* so.

That this response overcomes the argument is undeniable. For if the mereological structure at a world is purely contingent, then premise 3 of the Cartesian Argument would be false, and thus the argument would be unsound. But the success of the response comes at a significant cost: admitting that mereological facts *are* contingent. This is a very controversial claim, for it is widely agreed amongst philosophers that the laws of mereology, whatever they may be, are necessary. Indeed, this is *so* widely agreed that it is very much the default view.[267] To claim anything to the contrary is to go very much against the grain. But, of course, we shouldn't care too much about going against the grain if we have good reason to do so. And as it happens, a small number of philosophers currently believe that, in the present case, we *do* have good reason for straying from the norm.[268] Ross Cameron, for instance, argues that there is no reason to believe that mereological facts hold with necessity, despite the fact that they are *a priori* justifiable.[269] I won't go into Cameron's arguments here, but I should point out that if they are successful, then they would affect much more than just claims about mereology. All manner of metaphysical truths would be at risk of having their modal status downgraded from necessary to merely contingent.

I must admit that Cameron's arguments tend to leave me feeling somewhat disconcerted. For I find the arguments themselves really quite compelling, yet I find the conclusion they lead to utterly abhorrent. Fortunately, however, for present purposes I think it is safe to push the whole issue to one side. It merely remains to be said that *if* one is sympathetic to the position that mereology is contingent, then one has an easy way of overcoming the Cartesian argument. Quite simply: if the mereological facts are not necessary, then the argument fails. Fortunately, however, one does not *have* to endorse this view to overcome the argument, for there are other responses available, as I shall presently show. So for those who want to protect the modal status of their metaphysics, please read on.

---

[267] Sider (1993a; 2001), Markosian (1998b; 2008), and Armstrong (1997), are just a few of the many thinkers who take mereological facts to be necessary.
[268] Specifically, see Cameron (2007). But also see Nolan (2005).
[269] Cameron (2007)

### *§6.1.2 Response Two: Physical Divisibility Does Not Entail Mereological Complexity*

There is a certain sense in which the Cartesian argument is undeniable. *Of course* one can conceive that any extended object could be divided into at least two smaller parts. Provided one can conceive of a knife sharp and strong enough, and a technician skilled and precise enough, there should be no conceptual barrier to stop one conceiving that any extended object could be physically bisected. But I think we ought to remember that the type of division we are conceiving of here is exactly that: *physical.* So the question I think we should consider is: why should this tell us anything about mereology? In other words, why should we suppose that just because something is capable of being physically divided, that it is therefore not mereologically simple?

At first glance this may seem an outlandish proposal. The thought would be, I suppose, that if you physically divide an object into two separate halves, then *of course* that object has proper parts. It has the two halves you have just divided it into at the very least! For what else are the two distinct objects that remain after your act of division, other than proper parts of the original object? On this view, it may be seen to be an analytic truth that something which has been divided (or indeed something that *could* be divided) cannot be mereologically simple. However, I think that a lot hinges on what one takes the term 'physical division' to mean. If one takes it to mean 'physical division *into parts*' (where 'part' is taken to mean 'mereological part'), then it *is* an analytic truth that anything physically divisible cannot be mereologically simple. But this is not what I mean when I talk of physical divisibility. Perhaps my choice of terminology is a little misleading, so let me try to explain what I *do* mean without employing potentially ambiguous terms like 'division'.

To begin, imagine an extended simple, call it *S*, that is shaped like a rectangle.[270] *S* would look something like the object depicted in *fig. 5* below:

*Fig.5*

[270] You can think of *S* as a two- or three-dimensional object, it won't affect the argument.

Now imagine that one were to apply a certain force to *S* with a sharp instrument of some kind, like a knife. Let's call this act 'slicing'. Now suppose that after our act of slicing, *S* were to look something like the object depicted in *fig. 6* below, let's call it *S**:

*Fig. 6*

It would appear that after our initial slicing, *S* has certainly undergone some change. It has, for instance, changed shape; it was initially rectangular, whereas now it is not. It is also likely to have different properties after its slicing; presumably it would, for instance, have a higher density after the slicing than it had before. It may well have different aerodynamic properties were it to move through the air, and so on and so forth. However, there is no obvious reason to suppose that its *mereological structure* has changed in any way. That is to say that provided one accepted *S* as being mereologically simple initially, there is no reason that I can see that should cause one to think it is no longer mereologically simple after it has been sliced and becomes *S**. It may have changed shape, or density, or whatever, but it seems that this should not have any bearing on its simplicity.

With that in mind, then, now imagine that we sliced *S**. Suppose that the force and direction of this slicing was such that afterwards, *S** looked like the object depicted in *fig. 7* below, call it *S***:

*Fig. 7*

In other words, *S*** has been *sliced right through*. Now the question that I want the reader to consider is whether *S*** is a mereologically complex object or not. As we have seen above, when *S* was first sliced, and became *S**, there seemed no apparent reason to suggest that *S** was no longer a simple. The slicing it had undergone had merely changed its shape, along

with certain other of its properties, such as its density and so on. So perhaps we could say the same of *S***? Perhaps we should conclude that *S*** is still a mereological simple, but it has simply changed shape from when it was *S**? Now *if* this were the case, then the Cartesian argument against extended simples would be severely undermined. For when it is claimed that for *any* extended object, it is possible that it may be divided, I think it must be something like this (i.e. being sliced right through) that the Cartesian has in mind. In fact I can't think of any other reasonable way to interpret what is meant. It can't be supposed to mean that for any extended object it is possible that it is mereologically complex, for that would simply beg the question. Thus the only reasonable interpretation of premise one is that it claims that any extended object could conceivably be physically bisected, i.e. sliced right through. But, of course, if it transpires that an object could be sliced right through yet remain mereologically simple, then the mere fact that an extended object could be sliced right through would not warrant the inference that it must be mereologically complex.

To properly analyse this line of response, we need to investigate whether it is at all reasonable – indeed whether it is even coherent – to suggest that an object such as *S*** could be mereologically simple. The *prima facie* objection would be to simply assert that of course *S*** is not simple, *for it is in two bits* (read: "two parts"). But this is not to *argue* against the response being considered, but rather, it is to flat-footedly deny it. The advocate of the current view is unlikely to be worried about this response, for she will simply reaffirm her view, and say that *S***, contrary to appearances, is not in two parts, but merely has an unusual shape. It is a single, simple object, which happens to occupy two disconnected regions of space. It is, one might say, a *scattered simple.*

I am willing to concede that scattered simples are a little odd. They may even be considered *radically* odd. So for this response to be considered even as a possibility, the burden lies with me to provide reasons why we should entertain it at all. So let me try. Firstly, when considering mereology in metaphysics, we are considering the *fundamental structure* of physical objects. And there is a fairly obvious sense in which this should go beyond mere considerations of how we can physically manipulate those objects. Just as in most areas of metaphysical enquiry, when trying to uncover the *real* truth of the matter into which we are enquiring, we must remember that superficial considerations of how we perceive the world, and how we physically interact with it, may sometimes mislead us; they

may be just that – *superficial*. That is not to say, of course, that such considerations should be ignored, but merely that they should be treated with caution. And with that in mind, then, it may well seem a little naive to think we can uncover any real truths of mereology just by thinking about whether we could physically cut an object or not. Disputes over mereological complexity are not to be decided with a scalpel and a steady hand, I would suggest, but instead, through a process of rational and logical thought.

To reinforce the point, I invite you to reconsider the earlier discussion of the Special Composition Question, and in particular, those potential solutions to it which van Inwagen called the "simple bonding answers". Van Inwagen showed quite comprehensively the futility of attempting to answer the SCQ by considering different methods of physical bonding. Being thorough and methodical as he is, van Inwagen presented philosophical arguments to show that the simple bonding answers were not satisfactory. But I think it is evident that he never held out any genuine hope that they would be. By tackling them he was instead trying to show that those approaches to the SCQ which may immediately present themselves through common sense may not, in fact, be correct. And he is surely right! When trying to answer the *metaphysical* question of what it takes for two material objects to compose a further material object, it seems positively childlike to think the answer is to be found by considering various different methods of physical adhesion ("hmm" said one metaphysician to the other, "shall we fetch the spot-welder, or do you think a staple-gun will do the trick?"). It is akin to thinking we could answer metaphysical questions about the passage of time by finding a suitably accurate stopwatch.

But if this is the case when considering mereological *composition*, then it seems quite reasonable to suppose it is also the case when it comes to *decomposition*. That is to say that if we shouldn't let our approach to the SCQ be clouded by considerations of how objects can be physically bonded, then perhaps we shouldn't let our views on extended simples be obscured by considerations of how objects can be physically cut. Consider the following quote from Hud Hudson: "It may be the physicist's job, for example, to tell us whether the fundamental entities that physics appeals to are physically indivisible one-dimensional strings, but it is the job of the metaphysician to tell us whether those uncuttable things are composite".[271] I think Hudson is right to remind us that there is a clear distinction between

[271] Hudson (2006), 107. Note that despite this quote, Hudson does not believe extended simples are possible.

the physical and the metaphysical.[272] And it is on these grounds that the metaphysician could resist the Cartesian argument. Sure, we can conceive that any extended object could be *sliced right through*. But for that to have any bearing on the debate over extended simples, we would need the further premise that slicing entails mereological complexity. But in light of the above comments, there is no obvious reason to believe that such an entailment holds.

Endorsing scattered simples, therefore, offers one line of response against the Cartesian Argument, although it is, admittedly, a fairly controversial view. For instance, there could, on this view, be a single, simple object located at some point in this very room in which I am writing, and also simultaneously located at some point in a distant galaxy, yet not located at any point or region connecting the two. Furthermore, this object's 'dual-location', if I may so call it, would not obtain in virtue of it having distinct parts - one in this room and one in the distant galaxy – for by definition it has no proper parts. It would be reasonable to consider such an object as a little peculiar. But for now, despite their peculiarity, I shall consider their existence as an open possibility, and thus consider this line of response as a live option (albeit a controversial one) for the anti-Cartesian to take.

## *§6.1.3 From Conceivability to Possibility*

The Cartesian Argument is a classic example of a conceivability-to-possibility argument, in that Descartes argues from the fact that he can conceive of $x$ being the case to the fact that $x$ really could be the case. Of course, Descartes is no stranger to this method of argument; his argument for the distinctness between mind and body is one of the most famous examples of a conceivability-to-possibility argument in the entire history of philosophy.[273] However, regardless of the success (or otherwise) of that particular argument, I remain unconvinced that in *this* case, the mere ability to conceive of any extended object being divided, is enough to justify the much stronger conclusion that all extended objects are of necessity divisible.

I think that the best way to respond to the Cartesian Argument is to simply deny that, in this case, conceivability entails possibility. That is, one can assent to the first premise

[272] It may be more accurate to say it is clear that there *is* a distinction between the physical and the metaphysical, rather than that the distinction itself is clear.
[273] Descartes (1986), *second meditation*

in the argument, whilst denying the second. The proponent of extended simples can agree that there is nothing conceptually incoherent in supposing that an extended object could be divided. There is no logical contradiction involved in such a supposition. And that is because it is logically possible. But that is *all* it is: a *mere logical possibility.* But just because something is merely logically possible does not mean that it must be accepted as a genuine *metaphysical* possibility – Kripke taught us that.[274] I mean, a nihilist can surely accept the logical possibility of there being composite objects like tables and chairs – she can conceive of such things – but it is simply a consequence of her theory that such things are genuinely metaphysically impossible. Conversely, the universalist can accept the logical possibility of two material objects failing to compose, but this does not undermine the necessity of her thesis. One could not argue against the universalist by merely saying: "but I can imagine a scenario where some objects fail to compose!" Well, one *could* argue in that way, of course, but it would be a poor argument indeed.

To press the point home, consider something that I pointed out in an earlier footnote: that Descartes's reasoning about mereological complexity entails that the material world is, of necessity, gunky. For if all matter is, of necessity, extended, and extended things are, of necessity, divisible, then all matter is necessarily divisible, i.e. gunky. Let's call this view 'The Gunk Hypothesis'. The Gunk Hypothesis is controversial but it is not nonsensical; it is a view that deserves to be taken seriously at the very least. James van Cleve, for instance, endorses the very same view.[275] But imagine if an opponent of the view were to present the following argument against it:

1. It is conceivable that there are simples.
2. Therefore, it is possible that there are simples.
3. If the Gunk Hypothesis is true, it is impossible that there are simples.
4. Therefore, the Gunk Hypothesis is not true.

This argument is of exactly the same type as the Cartesian Argument. And whilst it is valid, it is hardly convincing! It is certainly not going to convince the proponent of the Gunk

[274] Kripke (1980)
[275] Van Cleve (2008). Note that van Cleve's reasons for endorsing this view are different from those of Descartes. His conclusion is the same, however: "all is gunk".

Hypothesis that their view is false. For they will simply deny premise 2. And they are well within their rights to do so; for it is simply a consequence of the Gunk Hypothesis that simples are impossible. A simple object may well be conceptually coherent. There may be no logical or epistemic barrier to conceiving of a thing with no proper parts. But the proponent of the Gunk Hypothesis can accept that. He will merely maintain that simples are *metaphysically* impossible. I think that the proponent of extended simples should respond to the Cartesian Argument in exactly the same way. The mere fact that one may be capable of imagining that any extended object could be divided is not enough to undermine the possibility of extended simples.

## §6.2. The Problem of Spatial Intrinsics

Anyone who is at all familiar with contemporary metaphysics will, more likely than not, be familiar with what has come to be known as *the problem of temporary intrinsics*.[276] Roughly, the problem arises when trying to explain the seemingly evident fact that an object can change its intrinsic properties over time yet still remain the same object. Over time, a banana will change from *being yellow-all-over* to *being brown-all-over*, yet still remain the same banana. But then nothing can be both yellow-all-over *and* brown-all-over, so what gives?[277]

It has been suggested that extended simples are at threat from a similar problem.[278] The problem is not how they can have different intrinsic properties at different *times,* but rather, how they can have different intrinsic properties at different *spatial locations.* Suppose, for example, you had an extended simple in the shape of a cube – call it 'Cube'. And now suppose that each of Cube's faces was a different colour. The question is: how are we to explain this qualitative variation across Cube? Since Cube is a single mereologically simple object, any properties it may have must be instantiated by it, the whole of it, and nothing other than it. But if that's the case, then it seems that one and the same object – Cube – is instantiating multiple and incompatible properties (e.g. redness and blueness) at the same

[276] This term was coined by Lewis (1986a), 202, but the problem itself dates back much further, probably all the way back to Parmenides.
[277] Of course, colour properties are unlikely to be considered *intrinsic,* but they serve to illustrate the problem nicely.
[278] See McDaniel (2003), 272

time. And that's impossible. So the argument is that the *only* way we can explain the qualitative variation of Cube is to insist that it is made of parts. That way we can say that it has a part which is blue, a part which is red, and so on and so forth, and there is no resulting contradiction. If Cube is mereologically simple, we have no such explanation available.

Since this problem is the spatial analogue of that highlighted by Lewis, it has been called "the problem of *spatial* intrinsics".[279] The first thing to note is that it is not an argument against the possibility of extended simples *per se.* Rather it is an argument against the possibility of extended simples that exhibit any qualitative variation. Thus the argument would have us conclude that if there are any extended simple things at all, then they must be qualitatively homogenous. Not all friends of extended simples will feel the need to overcome this argument – they may just accept that extended simples *are* all qualitatively homogenous and have done with it.[280] From a monist's perspective, however, this doesn't look like a live possibility. That the world is awash with qualitative variation is a fact so evident as to be indubitable. And since the monist claims the world to be a single mereological simple, the conjunction of these two claims entails the fact that extended simples *are* capable of exhibiting qualitative heterogeneity. So with the monistic project in mind, the problem of spatial intrinsics must be overcome. In what follows I will press this problem a little further and consider some possible responses to it. I will argue that most of these possible responses, whilst independently quite promising, should be considered quite unfit for purpose by the monist. I will conclude by presenting a rough outline of how I think the problem is best overcome. The outline will only be very rough at this stage, since I will be revisiting the problem in much more detail in the next chapter.

## *§6.2.1 Response by Analogy*

The problem of *temporary* intrinsics has received a great deal of attention from metaphysicians since Lewis brought it into focus a few decades ago, and a number of interesting and innovative responses have been offered in reply. So those responses present a natural place to start looking when thinking about the analogous problem of *spatial*

[279] McDaniel (2003), 272.
[280] McDaniel seems to have endorsed this view for a period, although he has now rejected it. See McDaniel (2009), 325

intrinsics. If any responses are considered successful in the temporal case, then there may well be an analogous response in the spatial case that also proves successful. So let me spend a little time considering those responses, and seeing how they fare in the spatial case.

Lewis's own solution to the problem of temporary intrinsics was to take the view that temporally persisting objects are composed of instantaneous temporal parts. Going back to our earlier example of the banana, then, Lewis would say that the banana has a temporal part that is yellow and a distinct temporal part that is brown. Since these temporal parts are distinct from one another, there is no contradiction involved in the fact that they instantiate incompatible properties. Lewis's solution has enjoyed significant success, at least if success is to be gauged by popularity and endorsement among other philosophers.[281] But it offers no help at all in the analogous spatial case. The spatial analogue of Lewis's solution would be the view that all extended objects are made up of non-extended parts and that qualitative variation across extended objects was to be explained by appealing to the properties instantiated by those objects' parts and the relations in which they stand. For the friend of extended simples, then, Lewis's response is a non-starter.

Another response that can be quickly dismissed is derived from the Presentist.[282] According to Presentism, only the present time exists, and thus only present properties exist.[283] Hence our banana is saved from contradiction since the only existing properties it instantiates are the ones it instantiates *now.* And it doesn't instantiate *being brown-all-over* and *being yellow-all-over* right now, therefore there is no problem. Once again, however, it seems fairly clear that this response is not going to help in the spatial case. The spatial analogue of presentism would be the view that only the present *location* exists. This view seems plainly false - that is, if one can even make sense of it at all - so I will not consider it any further.

Perhaps the most common way of responding to the problem without resorting to temporal parts, is to take a *relational* view of properties or property instantiation. There are various ways in which this can be done. On one view, which I shall call 'Property Relationism', the idea is that properties that may *seem* to be monadic, such as *being yellow,*

[281] Notable proponents of temporal parts, other than Lewis himself, include Armstrong (1980), Heller (1984; 1990), Jubien (1993), Sider (2001), to name but a few.
[282] See Lewis (1986a), 203, for a discussion of the Presentist response.
[283] Defenders of Presentism include Bigelow (1996), Hinchliff (2000), Markosian (2004), and Zimmerman (2008).

are in fact dyadic relations linking an object to a time, such as *being yellow at time t.* Going back to our banana, then, the problem is overcome because the banana stands in the *being yellow at* relation to *t1,* and the *being brown at* relation to *t2.*[284] On another view, which I shall call 'Instantiation Relationism', the idea is that there is a three-place relation of *instantiation,* which holds between objects, (monadic) properties, and times. Thus on this view our banana stands in the instantiation relation to the property *being yellow* and the time *t1,* and also stands in the same relation to the property *being brown* and the time *t2.* The final view I shall consider is known as 'Adverbialism'. The idea here is similar to Instantiation Relationism, in that it is not the properties that are seen as relational, but the way in which those properties are instantiated. Thus one way of expressing the fact that a banana is yellow at a particular time, on this view, would be to say the banana instantiates-at-time-*t* the monadic property of *being yellow.* A more common, if somewhat less grammatical, way to express this claim is to say the banana instantiates *being yellow t-ly*.

All three of these views can be modified in such a way that they respond to the analogous problem of *spatial* intrinsics. Very roughly, the difference is that rather than relativising properties, or instantiation, to *times,* they relativise them to *locations* (i.e. points or regions of space). This can be fleshed out in a little more detail as follows. Firstly, recall our hypothetical extended simple, Cube, with its different coloured faces. The problem was how to explain the qualitative variation across Cube's extension without resorting to positing its proper parts. Below is an explanation of how our three views, when suitably modified, would respond to that problem:

**A. Property Relationism**

On this view, putatively monadic properties, such as *being red,* are taken to be dyadic relational properties that link objects to locations, such as *being red at region r.* Thus on this view, Cube's qualitative variation is to be explained by saying it instantiates the property *being red at* $r_1$, and *being blue at* $r_2$, and so on, where $r_1$ and $r_2$ are those regions which correspond to the different faces of Cube.

[284] See Lewis (1986a), 202, for a discussion, and rejection, of this view.

**B. Instantiation Relationism**

On this view, instantiation is a triadic relation that links objects, properties, and locations. Thus Cube's qualitative variation is to be explained by saying that it stands in the instantiation relation to the property *being red* and the region $r_1$, and also stands in the instantiation relation to *being blue* and region $r_2$, and so on and so forth, where regions $r_1$ to $r_6$ correspond to the different faces of Cube.

**C. Adverbialism**

On this view, the problem of Cube's qualitative variation is answered by pointing to the way in which Cube instantiates its properties. Just as in the temporal case, it is difficult here to express the view grammatically, since there are no existing adverbs fit for the precise purpose. One could follow the lead of Jonathan Schaffer, and say that an object can instantiate a property (such as *being red*) *here-ishly* or *there-ishly*.[285] One could instead try to be a little more precise, as Horgan & Potrč have, and say Cube, for example, instantiates the property *being red* $r_1$*-ly* and the property *being blue* $r_2$*-ly*, and so on, where $r_1$ and $r_2$ refer to the appropriate regions.[286] But whatever locution one chooses, the difficulties in actually expressing the view should hopefully not obscure its content – that it is the fact that a single object can instantiate different (and incompatible) properties in *different ways* that is invoked to solve the problem of spatial intrinsics.

In the temporal case, I think that a decent case can be made for all three views as alternatives to an endorsement of temporal parts. And this is in spite of the fact that all three have been subject to criticism in the extant literature.[287] However, things are less straightforward in the analogous spatial case. When considering the possibility of qualitatively heterogeneous extended simples *per se*, I think that the three views above may well offer genuine promise. For I can see no principled reason why putatively monadic properties should not actually be dyadic, or why the instantiation relation should be limited to just two places rather than

[285] See Schaffer (2007), 179n

[286] Horgan & Potrč (2008), 177-178

[287] For a critique of Property Relationism, see Lewis (1986a), 202. For a critique of Adverbialism, and a more sympathetic view of Property Relationism, see Hawley (2001), 16-24. For a defence of Adverbialism, see Johnston (1987) or Haslanger (1989).

three. But things are very different when considering monism, and there are two main reasons for this. The first reason is that all three views seem unavoidably committed to a robust realism about substantival space. The first two views require either regions or points (or both) to act as relata for the relations they posit, and the third view requires regions or points in order that the adverbs it posits (such as *r1-ly*) have a discernible source of semantic content. (By that I mean that if one were an anti-realist about space, then it is not clear to me what one could possibly mean when saying that an object instantiates a property *r1-ly*).[288] I cannot see any way in which any of the three views could be held, coherently, without also committing one to the view that spatial regions and/or points *really* exist.

The problem with this is that the monist should, in my view, reject Substantivalism about space. The reason for this should be fairly obvious. Monism states that there is only a single concrete object in existence – the world – whereas Substantivalism seems committed to the existence of multiple concrete objects – the many points and regions of substantival space. So by adopting a Substantivalist view in response to the problem of spatial intrinsics, the monist gives with one hand whilst taking away with the other. She gets to explain the qualitative variation exhibited by the world, but only at the expense of the very integrity of her monism – by admitting multiple concrete objects.[289]

The second reason I think these responses fail is that they all seem unavoidably committed to external relations. For instance, if one adopted Instantiation Relationism, then the instantiation relations one posits are surely external. Cube may stand in the instantiation relation to *being red* and *r1,* for instance. But it surely *could* have stood in that relation to *being blue* and *r1.* The existence of the properties alone do not entail that the relation in question will obtain, thus it must be an external relation. But in the last chapter, you may remember that I touted it as a major advantage of monism that one can do away with external relations entirely. That advantage would be lost, and indeed that argument would be nullified, if one

[288] Interestingly, Horgan & Potrč (2008), 177-178, endorse Adverbialism in this exact way, i.e. endorse Monism, Adverbialism, and reject the existence of space-time points or regions. They say, for instance, that Mass could be instantiated *R-ishly* by the world, and even that relations can be instantiated (R1-R2)-ishly by the world, whilst they deny the reality of any spatial regions. I'm just not sure I can make sense of what *R-ishly* or *(R1-R2)-ishly* could possibly mean, if there are no regions to which R refers. It is because of this that I reject Adverbialism as a viable option for the monist.

[289] It may be objected that space is not really a concrete object at all, or at least, that it is of a different ontological category from the concrete objects that reside in it. Maybe such a response holds promise? I will just say that I think, from a monist's point of view, it looks like a bit of a cheat. A region of space has physical properties. It has a size, a shape, and so on. So regardless of what you call it, to me, it will always be a physical object.

adopted one of the above solutions to the problem of spatial intrinsics. I conclude, therefore, that the monist should reject them all.

### *§6.2.2 The Preferred Solution: Maximal Irreducible Properties*

Fortunately, the monist does not have to endorse any of the above responses in order to overcome the problem of spatial intrinsics. Instead, she needs only to endorse what I will call *Maximal Irreducible Properties.* This is a rather broad term, and how one characterises these properties could vary from case to case, but very roughly the idea is as follows. The properties are *maximal* in that they are instantiated by a *whole* object. (That is, if we are talking about extended simples *per se,* then it would be an entire extended simple that instantiated the property, and if we are talking specifically about monism, then it would be the entire world that instantiated the property). The properties are *irreducible* in that they cannot be reduced to more basic properties instantiated by an object's parts.

Consider once more our hypothetical extended simple, Cube. The idea is that the colouration across Cube's extension is to be accounted for by its instantiating a maximal irreducible property. Let's call it the property of *being multi-coloured.* Roughly speaking, then, Cube, as a whole, instantiates the single irreducible property of *being multi-coloured* and it is this that gives rise to its appearance. This line of thought would overcome the problem of spatial intrinsics, since it would not be the case that Cube simultaneously instantiates incompatible properties, like *blueness* and *redness.* Cube does not, in fact, instantiate *blueness* or *redness at all.* Rather, Cube instantiates a maximal and irreducible colour property; Cube is irreducibly *multi-coloured.*

Once the idea is grasped, it should be relatively easy to see how such properties could be invoked to explain all manner of qualitative variation that an extended simple could exhibit. An object could be irreducibly *textured,* for instance, in that it could be, say, *ridged,* or *perforated,* or even *perfectly smooth*. But, importantly, its texture would be explained by its instantiating a maximal irreducible *textural property.* Alternatively, an object could have an irreducible *heat property.* It could be a perfectly regular heat property, such that the object was the same temperature throughout, or it could be an irregular heat property, in that it was hotter in some places and cooler in others. Crucially, however, it would be a

maximal irreducible heat property that was responsible for this temperature variation, not the individual temperature values of the object's parts. Of course, the same idea could be applied to all sorts of qualities that an object could exhibit.

I hope, at this point, that the reader will at least agree that such properties are *possible.* For they are surely conceptually coherent, indeed, the idea of such properties is surely very easy to grasp. After all, if one can grasp what it is for an object to be *blue,* for instance, then one could surely grasp what it is for an object to be, say, *perforated* or *multicoloured.* Moreover, I hope that the reader will agree that if there *are* such properties, then they would overcome the problem of spatial intrinsics. For if an object can be irreducibly *perforated* or *multi-coloured,* say, then there is no need to appeal to the object's parts in order to explain its variable texture or colour. Despite this, however, I envisage that the reader may still have their suspicions. In particular, I have not given any good reasons for believing that there *are* such properties in the first place. They may, at this stage, appear little more than a desperately *post-hoc* postulation serving merely to overcome the problem of spatial intrinsics. A natural thought, for instance, may be to question *why* we should accept that such properties are *irreducible.* Sure, we can accept that a cube could be multi-coloured, the response may be, but the instantiation of that property only holds *in virtue* of the object's having different parts which are different colours. Indeed, it may be claimed that *all it is* for an object to be multi-coloured is for it to have distinct parts that are different colours.

Such a response is entirely understandable. In fact, I would respond the very same way were there no good *independent* reasons to believe in this particular type of maximal irreducible property. Fortunately, however, there *are* some good independent reasons. In the next chapter, I will spell out those reasons, and make what I hope is a compelling case for the existence of this type of property. Moreover, I will go into much more depth as to how these properties are best characterised. But that is for the next chapter. For now, I hope to have shown only how maximal irreducible properties, if there *were* such things, would overcome the problem of spatial intrinsics.

## §6.3. Alternatives to Extended Simples

As a final note, it should be worth considering what alternative views remain if one rejects the possibility of extended simples. For when considering what the ultimate nature of matter may be, there is a limit to the number of options available. Firstly, one must determine whether matter is, at its ultimate base, mereologically simple or not. Of course, all of the matter that makes up the actual world could be a combination of both types, simple and complex, but the fact remains that the two types are themselves exhaustive. For any particular piece of matter, it either has proper parts or it does not. There is no other way.

But then, as we have just been discussing, if one opts for the simple view of matter, there is still the small issue to be settled as to whether simple things can be extended or not. Once more, the simples that there are (if any at all) may be a combination of both types, extended *and* point-sized, but again, the two types should surely be regarded as exhaustive. For any particular material simple, it must either have extension or not. There is no other way.

Conversely, one could reject the idea of material simples altogether. On this view, then, all matter is mereologically complex. For any piece of matter you choose, it will have proper parts, and moreover, those parts will have further proper parts and so on and so forth. It is matter such as this that David Lewis labelled "atomless gunk".[290] Of course, the matter that makes up the actual world could be a combination of all of the three types just considered, but needless to say, those three types themselves should be considered exhaustive. At base, matter must bottom out into either extended simples or point-sized simples, or not bottom out at all. There is no other way.

So we have an exhaustive triumvirate of candidate theories for the ultimate nature of matter. Matter could consist of spatially extended simples. Alternatively, it could consist of non-extended, point-sized simples. Finally, matter could be complex *all the way down,* that is, consist of atomless gunk. Therefore, if one rejects the possibility of extended simples, one has to plump for one of the other alternatives: points or gunk. But both of these alternatives come replete with problems. I have already highlighted some of these problems in the previous chapter, so I will not labour the point here, but it remains to be said they are very

[290] Lewis (1991), 20

*real* problems. Gunk, for instance, is just weird. It is often said that it is perfectly possible that there is such stuff, since it is perfectly *conceivable.* But as I intimated previously, I am not convinced that this is true. Try as I might, I just don't think I can conceive of what a hunk of gunk would be like. To help us in this conceptual task, we are often provided with certain hypothetical scenarios that are claimed to contain gunk. Jonathan Schaffer, for instance, tries to spur us on with the following rhyme:

> Great fleas have little fleas upon their backs to bite 'em,
> And little fleas have lesser fleas and so *ad infinitum.*[291]

Charming as the rhyme may be, it provides no help in conceiving of gunk. It helps merely with conceiving of a series of objects ever decreasing in size (indeed it doesn't even suggest that these objects are *parts* of one another). But conceiving of *that* is not conceiving of *gunk.* Indeed, as I have stated before, I wonder what it is, *precisely,* that people *are* conceiving of when they say they are conceiving of gunk. It is my suspicion that when one claims to conceive of gunk, one conceives of something along the following lines. First one thinks of an ordinary object of some kind – it doesn't matter what – and then one thinks of cutting it in two. Then one thinks of cutting one of the remaining halves in two again, and then repeating this process over and over again. There is nothing troublesome about that; I have no difficulty in conceiving of such a scenario whatsoever. But I refuse to accept that it gets one any closer to conceiving of gunk. There are two reasons for this. Firstly, even if one imagines this cutting process to go on for a really long time, one can't imagine it going on *infinitely.* To imagine an infinite number of cuts would take an infinite amount of time. It would be simply beyond us. Secondly, and more importantly, however, even if one *could* imagine such a process going on infinitely, it would not demonstrate that the object being cut was gunky. Such an object could still be composed, ultimately, of point-sized particles. For no matter how many times you cut an extended object in half, no matter how small it gets, you will never get to a stage where the remaining halves are point-sized. Half of an extended object will *always* be an extended object. So bisecting an object an infinite amount

[291] Schaffer (2010c), 61. (Schaffer attributes the rhyme to Bohm (1957))

of times does not show that object to be gunky; it does not rule out the possibility that it bottoms out in points.[292]

Now if my suspicion is correct, and this *is* what people are imagining when they say they are conceiving of gunk, then it turns out that they are not conceiving of gunk after all. And if they insist that this is *not* what they are imagining, then it needs to be explained what they really *are* imagining. It will not suffice to merely say the words, "I am conceiving of an object whose parts all have further parts", for whilst I understand the words, they don't help me get a grip on what such an object would really be like. Until such an explanation is forthcoming, I am unwilling to accept that gunk really *is* conceivable. And if it's not even conceivable, then there is little reason to suppose that it is possible. It is for this reason, then, that I refuse to accept that matter could really be gunky.

So if gunk is out of the picture, and in light of what has just been said I will assume that it is, then he who rejects extended simples must accept that matter bottoms out in point-sized simples, for that is the only available view that remains. But this view is not without its problems either. In fact, in the last chapter, I demonstrated that if all matter is ultimately composed of point-sized objects, then it follows that all matter must be infinitely dense. But this is surely a troubling consequence; even if one claims it is not *impossible*, it still smacks of the absurd. So the punctal view of matter hardly represents some kind of metaphysical safe haven that offers shelter from controversy and counter-intuitive consequences; far from it.

So what does all this show? It shows, I think, that matter, however one conceives of it, is a lot more peculiar than we may ordinarily imagine. But this fact significantly diminishes the force of some of the standard objections to extended simples. For extended simples are often objected to on the grounds that they are too weird; too counter-intuitive. Consider, for example, these remarks from Dean Zimmerman:

> Perhaps a miraculous substance or a universal may be partless yet spread throughout an extended region; but it is hard to imagine an extended solid [...] behaving in any way which would lead us to describe it as partless.[293]
>
> When one sees the top half of the thing [an extended simple], one thereby sees the bottom half as well; when one touches the left side, one touches the

[292] I thank Mike Bench-Capon for pointing this out to me, and for much helpful discussion on this topic.
[293] Zimmerman (1996b), 154

> right also. Perhaps this bizarre mode of space-filling is not absolutely impossible; but surely the more sensible kinds of extended objects ought not participate in it.[294]

It is not uncommon to encounter this type of intuitive distaste for extended simples. And I suppose it is not that hard to see why; for I concede, extended simples *are* a bit odd. But what is seldom recognised, however, is that the alternatives to extended simples (i.e. gunk or points) are at least *equally odd.* Gunk is surely the most peculiar of the three. Indeed, as I have already made clear, I doubt that it is even possible; I certainly can't imagine what it would be like. And the punctal view of matter is peculiar too. In addition to the *intuitive* problems with it (i.e. if an object has literally no size, then in what sense is it there at all?), it entails the even more damaging fact that all matter must be infinitely dense. But once these facts are recognised, it makes no sense to rule one option out on the grounds that it is counter-intuitive and weird. They are all counter-intuitive and weird! In fact, what I hope to have impressed on the reader over the course of the present chapter is that, contrary to Zimmerman-esque views, extended simples actually represent the *least* weird conception of matter of the three.

## §6.4. Concluding Remarks

I hope to have shown that the widely held distaste for spatially extended simples is little more than prejudice. The arguments that are often wielded against them are unconvincing, for there are all manner of ways in which they can be responded to and overcome. But perhaps even more importantly than that, I hope to have shown that contrary to popular belief, extended simples actually represent the least problematic metaphysical conception of matter. Yes, they are a little odd in certain respects. But these oddities pale into insignificance in comparison with those of their rivals: points and gunk. It is for these reasons combined that we should endorse the conclusion that extended simples are a very real possibility.

---

[294] Zimmerman (1996b), 154. I actually think that Zimmerman's remarks are in danger of begging the question against the proponent of extended simples. For he considers what would happen when one sees the top 'half' of an extended simple. *But you never see the top half of an extended simple, because extended simples don't have halves!*

# §7. A World of Properties

So far, I have been quite vague in my explanations of what the monistic world is actually like. I have pushed to one side questions of how to explain this or that feature of the world, with little more than a wave of the hand. I have repeatedly suggested that monistic explanations will appeal to the world's *properties* rather than the world's *parts,* but I have refrained from going into any great detail about what those properties must be like, or how they can fulfil the explanatory role that the monist requires of them. This is because, up to this point, I have been arguing *for* monism; providing reasons and arguments to entertain it as a possibility in the first place. But now those arguments have been given, and hopefully the reader will have found them compelling, the time has come for the monist to do a bit of *explaining.*

The central task for the monist is to be able to sufficiently explain the appearances. In other words, the monist has to explain the appearance of plurality, given that there is, in fact, only singularity. For it is beyond doubt that it certainly *seems* as though there are many things. The world as it appears is populated by a vast number of distinct material objects: trees, rocks, cats, dogs, and so on and so forth. The monist needs to give us a story to explain how it is that it seems that there are such things, when in fact there are none. This is what I shall do in the present chapter. The problem is essentially the same as the problem of spatial intrinsics mentioned in the previous chapter, but on the grandest of scales. How is it that a single, maximally large yet mereologically simple object, can exhibit the kind of dazzlingly rich qualitative variation that we encounter in the world? The answer, as I have already alluded to, is through its instantiation of maximal irreducible properties. According to the monist, we live in a world of properties, not in a world of things.

## §7.1. Distributional Properties I: What They Are

I believe that the monist should endorse *distributional properties.* A distributional property is a property which not only ascribes a particular quality to an object (e.g. *redness*), but one

which also specifies the way in which that quality is distributed across the extension of that object (e.g. *red all over* or *red in the corners* etc.). The idea of distributional properties was first introduced by Josh Parsons, who tells us that "a distributional property is like a way of painting, or filling in, a spatially extended object with some property".[295] Distributional properties are most intuitive – we talk as though there are such things all the time. Consider, for example, the following statements:

1. The zebra is black and white striped
2. The surface of the lake is rippled
3. The Swiss cheese is holey

These are all perfectly ordinary, everyday kind of statements, yet taken at face value, they all seem to refer to distributional properties, i.e. the properties of *being black and white striped; being rippled; being holey.* Now I do not expect the reader to be immediately convinced, on the back of these example statements, that there *are* distributional properties. Arguments will follow which aim to do that. However, these statements illustrate nicely what distributional properties are supposed to be like. The idea of a distributional property is, I would suggest, very easy to grasp.

However, I envisage that there will be much disagreement over how seriously we should take distributional properties. Another way of saying that may be that there are starkly different ontological conclusions one can draw from analysing statements such as 1 – 3. One possible and fairly extreme view would be to reject distributional properties altogether. On this view, one would deny that there are any properties to which the predicates in 1 – 3 refer. The motivation for such a view would be, I suspect, some kind of appeal to ontological economy. For instance, one may think that there is only a sparse stock of objects and properties (presumably microphysical objects and properties), and that all features of the world can be explained fully and perspicuously in terms of these alone. On this view, there would simply be no need to posit objects like lakes or zebras, nor the distributional properties that they purportedly instantiate. I will call this view, the *Eliminativist View.*

[295] Parsons (2004), 173

A somewhat milder stance could be called the *Reductive View.* On this view, one would accept that there *are* distributional properties, but maintain that they are reducible to combinations or conjunctions of more fundamental, non-distributional properties. The reductionist, then, would deny distributional properties any ontological seriousness. The property of *being black and white striped,* say, would be considered on a par, ontologically speaking, with the property of *being king* or *being tasty,* for example. There *are* such properties, but they are ontologically lightweight; they would have no place in a complete and perspicuous description of what the world is really like. On the reductive view, then, to say that an object has the property of *being black and white striped,* for example, is really just to say that that object has some parts which are black and some parts which are white, and that those parts stand in the appropriate spatial relations to give rise to stripy-ness.[296] The distributional properties *reduce* to non-distributional ones.

There is nothing particularly unusual or revisionary about either of these views. Eliminativism and reductionism are commonplace in many areas of philosophy. Both, however, are incompatible with monism. The reason for this should, I think, be fairly clear. It is that both eliminitavism and reductionism about distributional properties seem necessarily committed to a pluralistic ontology, i.e. an ontology of multiple concrete objects. To see why this is so, consider an example. Let's continue with the example of the zebra having the property of *being black and white striped.* It was suggested that the reductionist about distributional properties would say that for an object to have the property of *being black and white striped* is nothing more than its having *parts* which are black and *parts* which are white, and those parts standing in the appropriate spatial relations required for stripy-ness. The eliminativist takes a similar view but simply rejects that there is any distributional property there at all; rather, *all there is* are the parts, and their properties, suitably arranged. In both cases, however, the appeal to parts looks *essential.* I fail to see how any reductionist or eliminativist strategy could work here *without* appealing to sub-portions (i.e. parts) of the extended objects purported to have distributional property. With regard to our hypothetical zebra, for instance, I just don't see how the reductionist strategy would work *without* appealing to the zebra's parts? It would have to reduce the property of *being black and white*

[296] Of course, as Parsons himself notes, it is unlikely that any reductionist would appeal to colour properties *simpliciter.* More likely, they would appeal to the fundamental properties of the microphysical particles, or something like that. Colour properties serve merely as an illustrative heuristic device.

*striped* to properties that are *non*-distributional, yet that are nonetheless instantiated by the entire zebra. It is hard to see what these properties could possibly be.[297]

The monist, then, must take distributional properties to be *irreducible.* This is for the simple reason that if they *were* reducible, then they could not be reduced to properties of the world's parts (because the world doesn't have parts), so they would have to be reduced to other maximal, yet non-distributional properties of the world. But it is hard to see what these properties could be, and, moreover, if there were any suitable candidates for such properties, then their existence would surely undermine the requirement for positing distributional properties in the first place. In other words, why posit distributional properties at all, if there are other candidate properties that do the explanatory work just as well without them?

So the monist requires an ontology that contains fundamental, irreducible distributional properties. It is the instantiation of these properties that gives rise to the rich and diverse qualitative variation of the world. Of course, when pressed, I think it is unlikely that the monist will want to posit properties like *being black and white striped* in any description of fundamental reality – these properties were merely invoked for illustrative purposes. I think that the monist can, and should, allow that there *are* such properties, *and* that they are distributional properties, but that she should maintain that they are somehow derivative on more fundamental distributional properties of the world.[298] When it comes to the world's *fundamental* properties, the monist may want to piggyback on the theories of modern science, for instance, and suggest that the world's fundamental distributional properties are things like *mass distributions, charge distributions, spin distributions,* and so on. More will be said about that later.

[297] I should note here that Kris McDaniel (2008) has attempted a reduction of exactly that sort. He suggests that distributional properties are reducible to short-lived, localised tropes. Whilst McDaniel's view is certainly interesting, I think the monist should reject it. For one thing, he takes distributional properties to be mereological sums of tropes. This directly conflicts with mereological nihilism, a central motivation for endorsing monism in the first place. I won't be considering McDaniel's view in what follows.

[298] Perhaps these types of 'lightweight' distributional properties *should,* then, be taken as being reducible. Importantly, however, they would only be reducible to other *distributional* properties.

# §7.2. Distributional Properties II: Why believe in them?

The monist could, of course, just *assert* that there are distributional properties and that they must be taken as irreducible; that this is a non-negotiable feature of monism. But it would make her position a whole lot stronger if she could argue for this claim on independent grounds. For if one were at all suspicious about distributional properties, one may use the monist's reliance on them as a reason to object against monism itself. If the monist could provide independent arguments to support the existence of irreducible distributional properties, then she could undercut this line of objection. Fortunately, I think there are some good independent reasons to believe in irreducible distributional properties, which I will set out below.

## *§7.2.1 Parsons's Argument*

Parsons has provided a most compelling argument for the irreducibility of distributional properties.[299] The argument aims to show that the reductionist view entails the claim that all extended objects must, *of necessity*, be composed of non-extended points. The reductionist view, then, rules out the possibility of both gunk and extended simples, or so Parsons would have us believe. This leads him to reject the reductionist view since gunk and/or extended simples should be thought of as, at the very least, possible. I think that Parsons's argument works convincingly, and, moreover, that it can be put even more forcefully than he puts it. I will set out the argument below.

The argument works best when set out in terms of an example, so let's return to our hypothetical zebra. The reductionist strategy would be to claim that the zebra's distributional property of *being black and white striped* is reducible to the conjunction of various non-distributional properties instantiated by the zebra's parts. But which parts? The obvious candidates are those parts which we can call the zebra's 'stripes'. Thus the claim may be that the zebra has some parts (i.e. stripes) that are black *simpliciter* and some which are white *simpliciter*, and that those parts stand in certain spatial relations appropriate for

[299] Parsons (2004)

stripy-ness.[300] But then the question is: what is it for an extended object to be black *simpliciter*? A natural response to this question would be that for an object to be black *simpliciter* is just for it to be plain black, i.e. uniformly black all over (I mean, how else *could* one respond?). But being uniformly black all over looks suspiciously like a distributional property, in that it describes not only the colour of the object, but the way in which the colour is *distributed* across the object (i.e. uniformly). But if this is right, then it looks like the parts that the reductionist must appeal in order for her reduction to work *cannot* be extended. For saying of *any* extended object that it is black *simpliciter* is merely to say that it has the uniform distributional property of *being black all over*. From this, then, it seems legitimate to infer that the only possible candidate parts that the reductionist can appeal to are non-extended parts, i.e. points. So the argument suggests that one can only reduce distributional properties to non-distributional ones *if* those non-distributional properties are instantiated by point-sized objects. Thus if there are no *irreducible* distributional properties at all, then *all* objects must be composed, ultimately, of point-sized particles.

The argument is compelling, particularly when it is set out with Parsons's usual flair. But it requires further examination; for there are a number of key assumptions upon which the argument rests that could be brought into question. Drawing out these assumptions explicitly will serve to make the argument more formal, and will enable us to subject it to a more rigorous scrutiny. So that is what I will do. First, however, it will be necessary to highlight an important distinction. The distinction is between *distributional* properties on the one hand, and *distributable* properties on the other.[301] Distributable properties are the qualities that distributional properties distribute. For example, the distributional property of *being black all over* ascribes a distribution of the distributable property *blackness*. Importantly, however, distributable properties like *blackness* are *not* distributional properties; they are non-distributional properties. What is also important, however, is that one should not take the terms 'distributable property' and 'non-distributional property' to be co-extensional. The reason for this being that there are, quite plausibly I think, non-distributional properties which are also non-distributable. To give a few examples, the property of *being a brother*, or

[300] I am supposing, for the sake of argument, that colour properties are fundamental here. Again, this is purely because colour properties illustrate the argument vividly. However, one could substitute any distributional property they like. The argument is designed to show the irreducibility of distributional properties *in principle.*
[301] Parsons also notes this distinction, and the term 'distributable property' is his.

perhaps, the property of *being located at*, are almost certainly non-distributional properties. But they also appear to be non-distributable properties, or so I will assume. The upshot of this is that the distributable properties are merely a subset of the non-distributional properties.

Now this distinction has been made clear, I will return to the key assumptions upon which Parsons's argument rests. Specifically, I think that three assumptions are required in order for the argument to go through. They are:

a) For each and every distributional property, there will be at least one corresponding distributable property that it distributes.[302]
b) An extended object cannot instantiate a distributable property without thereby also instantiating a distribution *of* that property.
c) All extended objects must, of necessity, instantiate at least one distributional property.

These assumptions are not made entirely explicit in Parsons's formulation of the argument, but I think they are required for the argument to go through. In due course I will say a little bit about why I think these assumptions should be accepted, but first, let me show why they are required for the anti-reductionist argument to work. With these assumptions in place, I think the argument can be set out as follows:

1. If reductionism is true, then all distributional properties are, *of necessity*, reducible to non-distributional properties. Specifically, they are reducible to their corresponding distributable properties (or conjunctions thereof). (E.g. the property of *being black and white striped* is reducible to the conjunction of instantiations of *blackness* and *whiteness*).

[302] Simple distributional properties, like *being black all over* will have only a single corresponding distributable property, i.e. *blackness.* More complex distributional properties, like *being black and white striped,* will have more than one distributable property, i.e. *blackness* and *whiteness.* The more complex the distributional property, the more distributable properties will be associated with it.

2. Therefore, if reductionism is true, an object's having a distributional property is necessarily reducible to its parts having only certain distributable properties (and standing in certain spatial relations).
3. From assumption *b*, it follows that these parts cannot be extended.
4. Therefore, if reductionism is true, all extended objects that instantiate at least one distributional property must be composed of point-sized parts.
5. From assumption *c*, it follows that if reductionism is true, then *all* extended objects are, *necessarily*, composed of point-sized parts.
6. From 5, it follows that if reductionism is true, then extended simples are impossible.[303]
7. Extended simples are possible.
8. Therefore, reductionism is false.

The argument is valid, but in order for it to be convincing some of its premises and assumptions are in need of some independent support. Specifically, I think that it is the three initial assumptions (*a*, *b*, and *c*), and premise 1, that require a defence. Let me provide that defence now.

Assumption *a* looks fairly uncontroversial. The very definition of what a distributional property is seems to entail that there must be at least one corresponding distributable property for each distributional property. To revisit Parsons's description: "a distributional property is like a way of painting, or filling in, a spatially extended object with some property".[304] What could be meant here by 'some property' other than a distributable property? If there were no distributable property corresponding to a particular distributional property, then what, exactly, would that property distribute? Nothing, seems to be the only answer, which suggests that there could not be any such distributional properties at all. Assumption *a*, then, looks fine.

Assumption *b* is also quite plausible, or so I would like to claim. For if *b* were false, then its falsity would entail the possibility of a scenario whereby an object instantiated a

[303] Parsons also claims that if reductionism is true, then it rules out the possibility of gunk, as well as extended simples. This is not strictly true. For that to be the case, it would have to be impossible for a gunky object to be composed of point-sized parts. But this needn't be impossible; at least, not if you accept that point-sized objects could have point-sized proper parts. Such a postulation would be strange, I think, but not impossible.
[304] Parsons (2004), 173

certain distributable property, yet at the same time, did not instantiate any distribution *of* that property. The problem is that I'm not convinced this is even a coherent possibility. How could an extended object instantiate *blackness*, for example, or *heat*, say, without having it distributed in some way or other across its extension? I just don't see that it could. I'm not sure I could even understand what was meant if someone said something like, "this poker is hot, but the heat is not distributed across it in any way, it is simply hot". They may mean, of course, that it is equally hot all over, but of course, this is *still* a distribution of heat; it is merely a uniform distribution. To suggest that something could be hot, but have no heat distribution at all is just nonsense, or so it seems to me. It is for this reason that I suggest assumption *b* is also true.[305]

Assumption *c* might initially strike one as a particularly bold claim. After all, I have only just introduced the very idea of distributional properties, yet I am now claiming that *all* extended objects have at least one distributional property *of necessity*. Perhaps it is a bold claim, but I think it is also a claim that has a great deal of plausibility. To see why, just think of an extended object – any extended object you like – and I am sure that you will find that it has at least *some* kind of property distributed across its extension. For surely, no matter how amorphous, bland or unremarkable an object may be, it will always have *some* kind of distributional property. Even a blank slate, for instance, will have a uniform distribution of *blankness* across its surface. Intuitively, then, it seems very difficult to conceive of an extended object that *didn't* have some kind of property distributed across its extension. Intuition, therefore, points to the truth of *c*.

But in fact, I think an even stronger case can be made for assumption *c*. For it is plausible to suggest that any extended object will always have a mass (even the fundamental particles of physics, tiny as they may be, are taken to have mass). But *mass* is surely a distributable property; for any extended object that has mass must have that mass distributed across it in some way or other. An object may have its mass uniformly distributed, in which case we would say it is uniformly dense, or alternatively it may have its mass non-uniformly distributed (e.g. have its mass concentrated in one corner, say), in which case we would say it has a non-uniform density. But either way, it seems that mass

[305] However, I could imagine an opponent remaining unconvinced. In fact, I think that assumption *b* remains the one potential weak spot of the argument. I will revisit this assumption, and its potential weaknesses, in §7.2.3, but for now, I will accept the assumption as true.

must always be distributed *in some way or other* across the extension of massive objects. It may appear that I am begging the question somewhat here, but I just can't see what the alternative would be? That an extended object could have a mass, yet not have that mass distributed across its extension in any way? That, to me at least, doesn't even sound coherent. For it seems reasonable to suppose that if an extended object has mass, then it is a perfectly reasonable question to ask whether it is uniformly or non-uniformly dense. But these two options are exhaustive – it has to be one or the other. (It would make no sense to say 'it is neither uniformly *nor* non-uniformly dense'). But to be either uniformly or non-uniformly dense *just is* to have a mass distribution of some kind. So having mass, and having extension, *entails* having a mass distribution. If these thoughts about mass are true, then assumption *c* is true. But even if you have reservations about the mass example - even if you disagree with me that mass is a distributable property - assumption *c* surely still remains extremely plausible. It would be a challenge, I think, to come up with an example of an extended object that did not have *any* property distributed across its extension. I don't think such a challenge could be met, and that is why I accept assumption *c.* And for anyone who *does* want to reject *c,* I think the onus is on them to meet the challenge set.

In addition to the three assumptions I listed at the outset, premise 1 might seem at least not *obviously* true, so I will say a little in defence of that too. The thought behind premise 1 is that one can only reduce distributional properties in terms of the distributable properties they distribute. For instance, if something has the distributional property of *being black all over,* then whatever it reduces to must be expressible in terms of the distributable property *blackness,* and no other distributable property. This claim has a certain intuitive pull about it. After all, it would be quite strange to suppose that one could reduce the property of *being black all over* in terms of any distributable properties other than *blackness*, say, *whiteness* or *redness.* However, I can imagine the following objection being raised: but colour properties *do* reduce to non-colour properties! The property of *being red,* for instance, is explainable in terms of much more fundamental properties like *wavelength reflectance properties,* say. But these are not colour properties! So why should I accept that distributional properties can only be reduced in terms of the distributable properties they distribute? Why not believe that distributional properties could be reduced to other types of property *entirely?* The objection makes a fair point, but it only arises due to the specific nature of the

example. I have been using colour properties to make my points because they provide particularly vivid examples with which to work. But I don't, when it comes down to it, actually believe colour properties to be fundamental, thus I don't expect them to be irreducible either. However, what I do maintain is that whatever properties colour distributional properties *do* reduce to must *also* be distributional properties. They may well reduce to *wavelength reflectance distributions* for instance.[306] The import of premise 1, however, is that whatever *fundamental* distributional properties an extended object has could only reasonably be thought to reduce to instantiations of their corresponding fundamental distributable properties. If you take *wavelength reflectance* to be fundamental, for instance, then you could not reasonably expect to explain a *wavelength reflectance distributional property* in terms of anything other than instantiations of *wavelength reflectance.* If you thought *mass* was a fundamental property, then you could not explain a *mass distributional property* in terms of anything other than instantiations of *mass.* This is what is meant to be asserted by premise 1.

In conclusion, then, I think that the argument presents a convincing case for the irreducibility of distributional properties. If you accept the initial assumptions and the premises, you should also accept the conclusion that distributional properties do not reduce to non-distributional ones. And in light of the previous remarks, I hope the reader will now agree that there are good reasons to take those assumptions and premises to be most plausible.

### *7.2.2 A Stronger Version of the Argument*

I think that Parsons's argument is fairly compelling as it stands. But I think it can be strengthened further if it is combined with some observations I made in the last two chapters. In chapter 5 I proposed that point-sized particles could not have mass, since if they did, they would have to be infinitely dense. I think that this claim, if true, causes irreparable damage to the reductionist view of distributional properties. The argument could be set out as follows:

[306] Cf. Parsons (2004), 177

1. If reductionism is true, then all distributional properties are reducible to conjunctions of non-distributional properties instantiated by point-sized objects.
2. Fundamental properties are irreducible.
3. Therefore, from 1 and 2, there can be no fundamental distributional properties.
4. Mass is a fundamental property
5. Therefore mass can only be instantiated by point-sized objects.
6. Mass *cannot* be instantiated by point-sized objects, since it would result in them being infinitely dense.
7. Thus we have a contradiction.

Let us assume, for the sake of argument, that premise 1 (i.e. reductionism) is true. If none of the following premises can be reasonably taken to be false, then we have an argument against reductionism. Premise 2 looks undeniable. Indeed it may be natural to suppose that the terms 'irreducible' and 'fundamental' are perfectly interchangeable.[307] There are, of course, different interpretations of the precise meaning of that metaphysical buzzword 'fundamental', but I think there is a *generally agreed* understanding of the term that will suffice for present purposes. Fundamentality is usually defined in terms of *dependence*, such that those things which are fundamental are those things which don't *depend* on any other things in order to exist.[308] A fundamental property, then, is one whose existence does not depend on the existence of any other properties. But if this is right, it is hard to see how one could maintain that a fundamental property could be reducible to other properties. The whole point of reduction is normally to reduce to a *more fundamental level.* Consider some often touted examples of reduction:

- Mental properties reduce to physical properties
- Colour properties of objects reduce to microphysical properties

[307] One may, of course, argue that there are no fundamental properties whatsoever. See, for instance, McKenzie (2011) or Morganti (2009). But that is a different project. Perhaps premise 2 should be better thought of as a conditional, i.e. *if* there are fundamental properties, then those properties are irreducible.

[308] This is certainly the way Schaffer defines it, see Schaffer (2003; 2009). Also, see Cameron (2008b) or Fine (2001) for endorsements of fundamentality being cashed out in terms of dependence. Although see Barnes (2012) for reasons to think that the two notions come apart.

In these kinds of case, it seems clear that what the reducible things are taken to reduce to, are assumed to be more fundamental. Moreover, it seems that the reducible things will always *depend for their existence* on the things they are taken to reduce to. If mental property *X* reduces to physical property *Y*, then *X* could not exist *without Y*. Reducible entities *depend* on their reductive grounds. On this understanding, then, it would not make sense to claim that a fundamental property could be reducible. For if it is reducible, there must be some other property (or set of properties) on which it depends. But if it depends on *anything*, then it is not, by this definition at least, fundamental. Thus on what I take to be the common understanding of fundamentality, premise 2 is true.

So we are left with premises 4 and 6. On the basis of what was said on the matter in the previous two chapters, I will assume that 6 is true. It would take a brave (and perhaps foolhardy) philosopher to reject the argument by rejecting 6. For that would result in the consequence that *all* matter was *infinitely* dense; a consequence that is immensely implausible at the very best, and physically impossible at worst.[309] So what about premise 4? Now I do not claim to be able to argue for the fact that mass *is* a fundamental property, but I will repeat what I stated in chapter 5: that if there are any fundamental properties at all, then mass is surely a prime candidate. It strikes me that it would be a far more radical move to accept reductionism at the cost of mass's fundamentality, than to hold on to the latter and reject reductionism. I will assume that I have the weight of the informed majority behind me here, in that I am assuming that the majority would ascribe far more certainty to the claim that mass is fundamental, than they would to the claim that distributional properties are reducible to non-distributional ones. The weight of the majority may not count for all that much, of course, but what it does do, I think, is put the burden of proof on the reductionist to explain why we should believe that mass is not fundamental. And it would not be enough to simply point to the above argument and say 'that's why'. What I mean is that the reductionist needs to explain *what mass is*, in more fundamental terms. Until that can be done, one should reject reductionism in favour of holding on to the fundamentality of mass.

[309] Simons (2004) certainly claims it to be physically impossible.

### *§7.2.3 A Potential Response*

Before moving on, there is a potential response to the argument that I have presented above, that needs to be pre-empted and dealt with. I gestured at this possible response above (see footnote 306), but will now expand upon it. The response would be to reject assumption *b*, the claim that extended objects cannot instantiate distributable properties without thereby instantiating specific distributions of those properties. As was shown above, this assumption is essential if the argument is to go through. But I can envisage a way in which it may be questioned, which I will elucidate below.

Let's consider a particular distributable property: *blackness.* It may be claimed that, contrary to assumption *b*, an extended object can just instantiate *blackness*, plainly and simply, without having to instantiate a particular *distribution* of *blackness.* This would make the object wholly black. Now I may call that a uniform distributional property, but if the objector is not already sympathetic to distributional properties, she may not be convinced that my calling it so actually makes it so. She may just think that what I call "being uniformly black" is just "being black", and that's all there is to it. In essence, then, this line of response would deny that what I have called uniform distributional properties *are* distributional properties at all. If this view is correct, then the whole anti-reductionist argument looks at threat. For the reductionist can simply say that non-uniform distributional properties (e.g. *being black and white striped*) are reducible to uniform ones (e.g. extended parts *being uniformly black* and parts *being uniformly white*). But since the reductionist does not take uniform distributional properties to actually be distributional, it looks like non-uniform distributional properties can be reduced to non-distributional properties after all.

In fact, our hypothetical objector could press her response even further and concede that if an extended object instantiates some distributable property, then *of course* that property will be distributed across the object in some way or other. But she may deny that any distributional property is required in order to account for this fact. Yes, it is true that the object has a uniform distribution of *F*, but no distributional property is required to make this true; it is made true merely by the object instantiating *F.* Instantiating *F simpliciter*, on this view, results in a uniform distribution of *F*, but this doesn't mean we need to recognise uniform distributional properties; we need only recognise *F.*

I don't think there is any knock-down response to this reductionist objection.[310] In fact, I think that, at base, it is likely to come down to a clash of intuitions about whether uniform distributional properties are really distributional properties or not. However, I think there are a number of reasons to suppose that the intuition that they *are* distributional properties is much more reasonable than the reductionist intuition that they are not. The first reason is that if the reductionist is right, then the very possibility of there being qualitatively heterogeneous extended simples is ruled out. To see why this is so, just imagine a spatially extended object, *O*, that *is* qualitatively heterogeneous; let's say it is black and white striped. We can say, therefore, *O* instantiates a non-uniform distributional property: the property of *being black and white striped.* But the reductionist's view *just is* that properties like this reduce to conjunctions of non-distributional properties. But the only way to perform such a reduction is to say that *O* has parts which are plain black and parts which are plain white.[311] More generally, then, the reductionist has to accept that her view entails the claim that all extended objects which exhibit qualitative heterogeneity must be composed of parts. Qualitatively heterogeneous extended simples are, on this view, impossible.

The reductionist may be happy to accept this consequence, but she must accept that it *is* a consequence of her view. The upshot is, then, that this reductionist response does not completely overcome the argument given above, but it does *weaken* it. For on this view, it does not follow from reductionism that all extended objects must be composed of points (i.e. that extended simples are impossible), as the initial argument suggests, but it *does* follow that extended simple objects cannot, of necessity, exhibit qualitative variation. And it is likely to be thought that the latter conclusion doesn't sound nearly as objectionable as the former. However, the anti-reductionist should remind her opponent here that her view doesn't entail *either* consequence. The view that non-uniform distributional properties are irreducible has no implications for the mereological structure of the objects that instantiate them; *any* extended object, regardless of its mereological structure, can in principle

[310] Parsons (2004), 179, does actually provide a response that he believes sufficient to render the objection false. I don't think his response works, however, so I won't go into it here.

[311] Providing, of course, that we are taking *blackness* and *whiteness* to be fundamental properties, which we are, for the sake of argument.

instantiate a non-uniform distributional property, if they are taken to be irreducible. This fact must surely be regarded as a cost of reductionism.

But I think that there is an even stronger reason to resist the reductionist claim that uniform distributions of *F* are not distributional properties. That reason can be made explicit through a consideration of what the initial motivation might have been for taking a reductionist view of distributional properties in the first place. The only reason I can think of that might motivate a reductionist view of distributional properties is some kind of desire for simplicity. More specifically, irreducible distributional properties might at first appear to be a bit unusual; some strange and entirely new type of property. Admitting that there are such properties, therefore, would constitute a significant addition to one's ontology. Thus the thought might be that if one can reduce distributional properties to mere conjunctions of more 'normal' properties, then one will have explained the phenomena, without having to inflate one's ontology with a new and mysterious type of property. If one can reduce properties like *being uniformly black* or *being black and white striped,* for instance, to uncontroversial properties like *blackness* or *whiteness,* then one will have undertaken a successful piece of ontological cost-cutting, or so the thought may be.

I am entirely sympathetic to this kind of reasoning, for I, just like most metaphysicians, am always partial to an ontological saving here or there. But my worry is that this parsimony-driven motivation may evaporate entirely if the reductionist also adopts the response that is currently under consideration. For the response is that uniform distributional properties *just are* 'normal' properties normally instantiated. So an object's *being uniformly black,* for instance, just is for that object to instantiate the property of *blackness.* The instantiation of this property results in the bearer being wholly black (or being black all over, or being uniformly black, or however you want to phrase it), but there is no distributional property involved at all. But my worry is, then, that this 'normal' property of *blackness* seems, after all, to be a very similar type of property to the distributional property *being uniformly black.* Both properties seem to play the same role, so to speak; both properties would give rise to the object being black all over, or wholly black, or whatever. But if that's the case, then in what sense is *blackness* any more 'normal' than the distributional property of *being uniformly black*? Or conversely, what is it about distributional properties, like *being uniformly black,* that the reductionist found so mysterious in the first place? If the reductionist

is happy to accept that for an extended object, *O*, to instantiate *F simpliciter* entails that *O* has a uniform distribution of *F*, then I can't see why she wouldn't also be happy to simply accept the existence of the distributional property *being uniformly F*. On what grounds could someone reasonably reject the latter as unusual, mysterious, or metaphysically excessive, while accepting the former as quite unexceptional? This is not at all clear.

In light of all this, then, I think there is good reason to reject this reductionist response, and maintain the success of the original anti-reductionist argument. For the reductionist response only works on the proviso that standard (i.e. non-distributional) properties, instantiated *simpliciter*, can take the place of uniform distributional properties. But that proviso seems to undermine the very motivation for taking the reductionist view in the first place, because if non-distributional properties do all the same work that uniform distributional properties are taken to do, then there seems to be no gain in simplicity in endorsing the one and rejecting the latter. Moreover, because the reductionist view involves a cost that anti-reductionist view does not (i.e. it rules out the possibility of qualitatively heterogeneous extended simples), then in the absence of any motivation for endorsing it, one *shouldn't* endorse it. The response, then, should be rejected, and irreducible distributional properties endorsed.

## §7.3. Distributional Properties III: Two Conceptions

I have, hopefully, painted a reasonably rich picture of what distributional properties are, what they are like, and most importantly, why we should believe in them. But I think that there is actually scope for a lot more detail to be added to this picture. In particular, I think that there are two quite distinct ways in which one can conceive of distributional properties. I will call them *The Structured Conception* and *The Unified Conception*. (Whether these names are truly representative of what the views entail, I will leave for the reader to decide, but they are merely names after all). The two views trade on the distinction that has been made earlier, between distributional properties and distributable properties. Essentially, the views differ as to how substantive this distinction should be taken to be. In very rough outline, the structured conception takes the distinction to be metaphysically substantive, and as such,

takes distributional properties to be structured properties composed out of distributable properties and distribution patterns.[312] In contrast, the unified conception denies any metaphysical weight to the distinction. Rather it takes distributional properties to be unified (i.e. *un*structured) and that distributable properties, if we admit that there are such things at all, are mere abstractions from their distributional counterparts. These are, of course, *very* rough outlines; I will go into more detail in what follows.

To begin, however, I should make it clear that I take distributional properties to be universals, as opposed to tropes. And that is on either of the conceptions that I briefly outlined above. (In the one case they are taken to be structured universals, and in the other they are taken to be unstructured).[313] The main reason for this is that I subscribe to the fairly common view that universals give us a better explanation of facts about resemblance, as opposed to tropes, whose proponents must accept that resemblance facts are primitive. Moreover, I think that universals are a whole lot less problematic when taken in conjunction with a monistic ontology than a pluralistic ontology. To see why this is, consider the objection that is most commonly presented against universals, and that is probably the main reason behind most trope theorist's acceptance of tropes: the problem of their multiple location. How can two distinct objects (two pillar boxes, say) share one and the same universal (redness, say)? More generally, how can a single universal inhere in multiple and distinct objects? This presents a problem that some philosophers have considered insurmountable. However, one can see that it only really rears its head when taken against the backdrop of a pluralistic ontology: an ontology of multiple and distinct objects. For the monist, the problem never arises – because there are not many things! There is only a single concrete particular – the world – and it is that and that alone which instantiates universals. There is nothing at all peculiar or problematic in suggesting that a single object can instantiate one and the same universal.

These comments need a bit of qualification. For one might reasonably think that for the monist there is actually no substantive distinction to be made between universals and tropes at all. The central distinction between universals and tropes is that the latter are taken

[312] The word 'compose' may be misleading here. I don't mean compose in a mereological sense. More will be said about this in due course.

[313] When I say 'structured universals', here, I have a different kind of thing in mind than the structural universals that can be found in the extant literature that have been argued for by the likes of Armstrong (1978) 69-71 & 120-127; (1986); (1997) and argued against by the likes of Lewis (1986b). Again, I will say more on this later.

to be particularised (i.e. each trope is unique and non-repeatable), whereas the former are taken to be repeatable (i.e. a single universal can have multiple instances; it can inhere in numerous distinct objects). But for the monist, there is only one object to instantiate properties, so neither tropes *nor* universals would be repeatable. Both could only have a single instance, by being instantiated by the world. Moreover, if the world is the only object there is, then there will be no facts about resemblance that need explaining, since there would be no other objects for the world to resemble, and thus that alleged advantage of universals would be lost. If this view is right, then the distinction between tropes and universals fades away, and it would seem a little pointless for the monist to choose between them.

I can see why one might adopt such a view, but I think it is false. The reason for this is that even if you're a monist, universals are still *in principle* repeatable, and this fact can help explain similarities between *ways the world could have been*. Consider, for example, a monistic world that is green and spherical. Presumably, this world *could* have been green and cubical. These two possible states of the world, although distinct, seem to bear some sort of resemblance to one another; they are both green after all. If the monist takes properties to be universals, she can explain this resemblance: both states share a universal – *green-ness.* If the monist endorses tropes, this explanation is not available.[314] This is why I think that the monist should maintain that there is a significant distinction between universals and tropes, and why she should endorse the former rather than the latter. With that said, let me now go on to elaborate on the two distinct conceptions of distributional properties of which I gave only the roughest of outlines earlier.

## *§7.3.1 The Structured Conception*

To begin, I should make it clear that when I talk of 'structured universals' here, I am not referring to the kind of structural universals that are perhaps best associated with David Armstrong.[315] This is because Armstrong's structural universals can only be instantiated by mereologically complex objects. As Armstrong states: "A property [universal], *S*, is

[314] More will be said about this in §7.3.3
[315] See Armstrong (1978) pp.69-71 & 120-127; (1986); (1997)

structural if and only if proper parts of particulars having *S* have some property or properties, *T* ... not identical with *S*, and this state of affairs is, in part at least, constitutive of *S*".[316] This implies that structural universals can *only* be instantiated by mereologically complex objects, thus for obvious reasons, the monist will not want to appeal to properties of this sort.

When I talk of distributional properties being structured, all I mean is that they have some sort of ontologically significant, describable structure. That structure is describable in terms of two distinct components: distributable properties and distribution patterns. The structure is ontologically significant because its components are taken to be very real, yet distinct, entities. In fact, I think that distributable properties and distribution patterns should be taken as universals themselves. Thus distributional properties are structured universals since they consist of further universals: distributable properties and distribution patterns. To give an example, the distributional property of *being black and white striped* is, on this view, taken to be a structured universal consisting of the distributable universals *blackness* and *whiteness* and the distribution pattern *stripey-ness.*

One may think that this is a kind of reductionist view, in that distributional properties are being reduced to other types of property (i.e. distributable properties and distribution patterns). Moreover, one may think it a little strange that I am even entertaining such a view, considering that I spent the previous section arguing for the irreducibility of distributional properties. But I think this would be to misconstrue the thought being proposed. I am not attempting to *reduce* distributional properties, but rather, I am giving a more detailed description of *what they are.* A distributional property is what the world instantiates when it instantiates a particular distributable property in a particular pattern. But there is no sense in which any of these properties should be seen as being any more or less real than the others. In fact, I think that they should all be seen as being mutually dependent. That is, I am supposing, *a la* Armstrong, that universals cannot exist uninstantiated.[317] Moreover, I claim that a distributable property cannot be instantiated without being instantiated *in some pattern or other,* and that a distribution pattern cannot be instantiated without there being some distributable property that it distributes. And since

[316] Armstrong (1978), 69
[317] Armstrong (1989), 75-79. He claims that this is the Aristotelian view of properties.

distributional properties *just are* conjunctions of distributable properties and distribution patterns, all three types of property are mutually ontologically dependent. The world cannot have a distributable property without it also having a distribution pattern, and it cannot have a distribution pattern without it having a distributable property to distribute. And having the two of these *entails* having a distributional property. So if you have *one,* you must have all three. So there is no reduction going on here at all. Just a more detailed picture of what distributional properties are.

But, once more, one may think why bother with the distributional properties at all? If a distributional property *just is* the conjunction of a distributable property and a distribution pattern, then why need the distributional property, in any real or non-reducible sense, at all? The thought is, I suppose, that this *would* be a reductionist view of distributional properties, but it would not be a view that is at all objectionable to the monist, since it does not involve an appeal to *parts* of the objects which are taken to instantiate the distributional properties. It is a reductive strategy that sits perfectly well with monism. This is true, but it is not a view I am willing to endorse, because there is another *independent* reason why it is problematic. This reason can be exposed by the following example. Suppose you had an extended, mereologically simple object which was black and white striped *and* had pink spots. This object, then, has the distributional property of *being black and white striped with pink spots.* Let's call this DPx. On the reductionist view of structured distributional properties that is currently under consideration, then, *all it is* for this object to have DPx is for it to co-instantiate the distributable properties of *blackness, whiteness,* and *pinkness,* and the distribution patterns of *stripy-ness* and *spotty-ness.* But herein lies the problem: what is it about these properties that make it the case that the *blackness* and the *whiteness* come together (so to speak) with the *stripy-ness,* and that *pinkness* comes together with *spotty-ness?* In other words, why isn't the object *black and pink striped with white spots* (or, indeed, any other possible combination)? The only possible answer to this question that I can see is that there is something that is *binding together* the distributable properties with the relevant distribution patterns. And whatever this something is, it looks suspiciously like an *external* relation. For it is surely a contingent fact that the object is *black and white striped with pink spots*. It surely *could* have been *black and pink striped with white spots*. But this represents a problem. For if you recall, I argued in chapter 5, that an advantage of monism is that it has

no need to posit external relations at all. So if it now turns out that I need to posit external relations to bind the world's distributable properties with the relevant distribution patterns, then that advantage will have been lost, and that argument will have been rendered useless.

Fortunately, however, the monist does *not* have to posit these external relations. But this is only the case if the monist rejects the reductionist view of structured distributional properties that was being considered above. Instead, the monist should take the distributional properties themselves to be fundamental. Thus if an object instantiates the property of *being black and white striped with pink spots*, then that property is fundamental and irreducible. But instantiating that property *necessarily involves* the world's instantiating of the distributable properties *blackness*, *whiteness*, and *pinkness*, and the distribution patterns of *stripy-ness* and *spotty-ness*, and also *necessarily involves* that those properties are bound together in the appropriate ways. That's just *what it is* for the world to instantiate that distributional property. Of course, one may insist that there is still *some* kind of relation holding the distributable properties together with the relevant distribution patterns, but it is only an internal relation, for the existence of the distributional property necessarily involves that those particular constituent universals are instantiated *and* that they are related in that particular way. So on this view, distributional properties are structured out of other universals, but there is no sense in which they are posterior to, dependent upon, or reducible to, those structural components of which they consist.

A fairly natural question at this point may be: *why* endorse the structural conception of distributional properties? For they seem like quite complex things. What is the need for introducing this complexity? There are reasons why I think that the structural conception is attractive, but I can't really make these reasons explicit until I consider the alternative view – the *unified conception* – in the next section. Very roughly, I think that the structural conception of distributional properties will help us answer some potentially troubling questions about resemblance between properties and between worlds. I also think that it gives us a better understanding of what distributional properties are. However, more will be said about this in §7.3.3. First, we must consider the other alternative: the unified conception.

## *§7.3.2 The Unified Conception*

In contrast to the structural conception of distributional properties, there is what I have called the *unified conception.* On this view, there is no underlying structure to distributional properties. They are, for want of a better word, *simple.* (I hesitate to use the term 'simple', since it has mereological connotations which might suggest that the structured properties I considered in the previous section are somehow mereologically complex. This is not the sense in which I intend the term to be taken). On this view, then, if something has the distributional property of, say, *being black and white striped,* then there is little more that can be said about it; there is little more that can be said by way of explanation about what it is to have such a property. Of course one could *talk* about properties such as *blackness*, *whiteness,* and *stripy-ness,* but these would be mere abstractions from the distributional property itself. There are no such properties as these, or at least, not in any real or fundamental sense. They would certainly not appear in any description of fundamental reality. So on this view, the distinction between distributional properties and distributable properties is eschewed. There are no distributable properties, and there are no distribution patterns. There are merely irreducible distributional properties. They are what they are and they resist any further explanation.

## *§7.3.3 How to Choose*

Both conceptions have their attractions. The defender of the unified conception (UC from here on) may appeal to economy for a start. After all, it is clear that the character of distributional properties is quite easy to grasp. If we are told that something has the distributional property of *being black and white striped,* say, then we know what that thing is going to look like. So why complicate things by introducing distinct *types* of property (i.e. distributable properties and distribution patterns), when a simple distributional property does the required work on its own? Moreover, one might question exactly what a *distribution pattern* is. For the defender of the structured conception (SC from here on) will presumably have to take distribution patterns as being fundamental, i.e. irreducible. They certainly won't be able to explain them in terms of an object's parts/regions and the spatial relations in which they stand. So the defender of UC might reasonably think that if we can understand

what a brute *stripy-ness* is, for example, with no further explanation, then we can just as reasonably understand what a brute *black and white stripy-ness* is too. Thus the separation of distributional properties into distributable properties and distribution patterns is, on this view, an unnecessary piece of metaphysical excess. The simplicity of the UC should surely count in its favour.

The SC has advantages of its own, however. In particular, I think that the defender of the SC may be able to provide better explanations of certain features of the world. Specifically, she will be able to explain certain facts about resemblance between properties. To illustrate this, consider an example. Imagine three quite simple two-dimensional worlds that are represented by the pictures in *fig.* .

*Fig. 8*

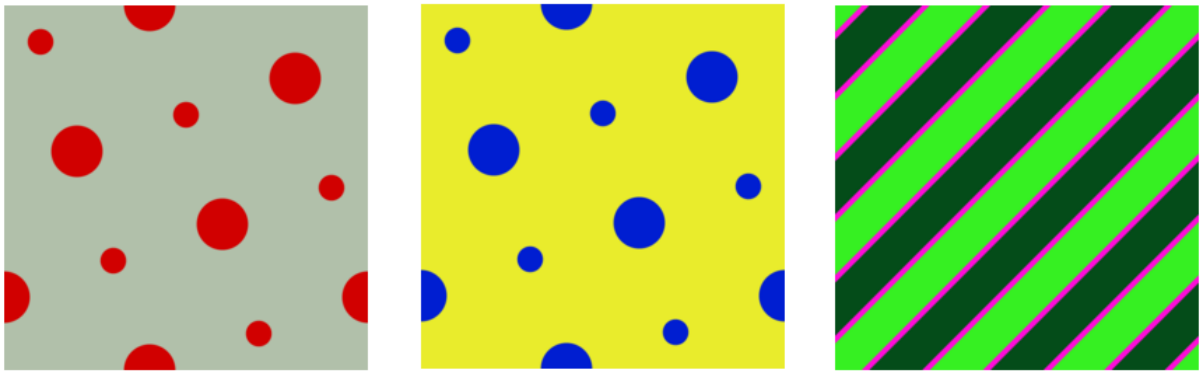


Let's further suppose that these worlds are monistic, i.e. they have no proper parts, and that they exhibit their qualitative variation by instantiating distributional properties. So we could say that *w1* instantiates the property of *being grey with red spots* (let's call it distributional property X, or DPX for short); *w2* instantiates the property of *being yellow with blue spots* (let's call it DPY); and *w3* instantiates the property of *being pink and green striped* (let's call it DPZ).

There seems to be an evident phenomenological fact that *w1* and *w2* are objectively more similar to one another than they are to *w3.* This fact is the kind of thing that needs explaining, or so it is reasonable to think. So how would our two conceptions of distributional properties go about explaining this fact of resemblance? The defender of SC would have a fairly obvious route. For her, the distributional properties instantiated by the three worlds are structured. Thus DPX, for instance, consists of the distributable properties

*greyness* and *redness,* and the distribution pattern *spottiness.* DPY consists of the same distribution pattern, but different distributable properties: *yellowness* and *blueness.* DPZ consists of the distributable properties *light greenness, dark greenness,* and *pinkness,* and the distribution pattern *stripy-ness.* (If one is concerned here that properties like *spottiness* or *dark greenness* look somewhat sloppy, then one should remember that they are only *names.* What we call these properties is fairly unimportant, it is what these names refer to that is of our genuine concern). So for the defender of SC, the reason why *w1* resembles *w2* but not *w3,* or more precisely, why DPX resembles DPY but not DPZ, is that the structured properties DPX and DPY share a constituent, the universal *spottiness.* DPZ shares no constituent universals with the other two distributional properties, and it is for that reason that it resembles neither of them.

For the defender of UC, the task looks a little more difficult. For her, there are no distributable properties or distribution patterns to appeal to; there are simply the three distributional properties. So on this view it looks as though the similarity between DPX and DPY will have to be admitted as a brute fact, which may be seen as a costly consequence of the view. One may be tempted at this point to remind us that DPX is the property of *being grey with red spots,* and DPY is the property of *being yellow with blue spots,* whilst DPZ is the property of *being pink and green striped.* Thus there *is* a similarity! DPX and DPY are both 'spotted' properties, whilst DPZ is a 'striped' property. This, of course, would be a blatant cheat. The term '*being grey with red spots*' is merely a name we have given to the property. We can't just assign names to properties and then note similarities between the names to explain similarities between the properties! (Imagine, for instance, the absurdity of saying a square was more like a triangle than a circle, because 'square' and 'triangle' both contain an 'a', whereas 'circle' doesn't). So I think that the defender of UC has no option but to accept these resemblance facts as brute. Why do DPX and DPY resemble one another? Well, they just do.

I think that these observations about resemblance should definitely count in favour of the SC. But I am unsure as to how much. For it strikes me that there are many claims about resemblance which will ultimately bottom out in brute facts. For instance, why is a red patch more similar to an orange patch than it is to a blue patch, even though none of them share the same universal? Perhaps they just are? Alternatively, perhaps the wavelength

reflectance profile of the red patch is closer to that of the orange patch than it is to that of the blue patch? But then, why are wavelength reflectance profiles more similar just because they are closer? Perhaps they just are? It seems that in many cases of resemblance, there will have to be a point at which one has to admit some kind of brute resemblance fact. According to the UC, we may reach this point sooner than if we adopt the SC, but it is not immediately clear how much of a problem that represents.

All in all, I take these considerations to point towards the SC as being the preferable view and, therefore, it is the view of distributional properties that I will endorse. However, I still see the UC as being a perfectly viable option. There may well be good reasons that I have not thought of to consider the UC as a much *better* option than the SC. But for now I am willing to leave that as an open possibility. In what follows, I will endorse the structured conception of distributional properties.

## §7.4. Distributional Properties IV: Saving the Appearances

Having established a suitably detailed conception of distributional properties, and having argued for why one should believe that there are such things, the final task for the monist is to explain, *precisely*, how they are best put to use. What I mean by that is that the monist has to explain how an ontology consisting of just a single concrete object (the world) and some irreducible structured distributional properties can sufficiently account for the phenomenological facts. The monist has to employ her metaphysical resources to give a satisfactory account of why the world appears as it does.

Ted Sider has summed up this task nicely by saying that the monist needs to be able to ground, or explain, the "as-if-facts".[318] Monism, just like a number of revisionary metaphysical theses, clashes with many facts of common sense and science; facts like, *there are tables* and *there are electrons.* However, whilst the monist denies that these are genuine facts – she denies that there really are any tables or electrons – she must accept that the world is *as if* they are facts. It is *as if* there are tables, and it is *as if* there are electrons. These 'as-if-facts' are not going to be considered fundamental facts, but they need to be grounded

[318] Sider (2008), 132

in, and explainable in terms of, the worlds fundamental facts. Essentially, the monist must explain how it is that a monistic ontology gives rise to the fact that it is *as if* pluralism is true.

Consider an analogy. If I hold up a fifty-pence piece to the night sky, from my perspective it may look *as if* the fifty-pence piece is bigger than the moon. This is an as-if-fact. But it is not a *fundamental* fact. If you were to give a complete and accurate description of reality, it would surely not include the fact that it is as if a fifty-pence piece is larger than the moon when observed in the right conditions. That is because this as-if-fact is grounded in more fundamental facts, like the actual size of the coin and the moon, their relative distances from the observer, and so on and so forth. The fundamental facts are the ones that give a true and perspicuous account of the phenomena. The as-if-facts simply *fall out* from the fundamental facts.

First of all, one may wonder why the monist should feel obliged to ground, or explain, the as-if-facts at all. Well, in truth, she doesn't *have* to. But it would make her thesis a whole lot more plausible if she could. Because if the as-if-facts were ungrounded, or inexplicable, then it would seem to follow that the world of appearances is nothing more than a mere illusion that bears no relation to how the world really is. One would simply have to accept that the world is *as if p*, when in fact it is the case that *not-p*, with no explanation of why this is. It is not incoherent, of course, to suggest that the world *is* like this; it is perfectly possible that the world of appearances *is* an illusion. But it is not a particularly appealing possibility. But more importantly, the monist does not want it to be a consequence of her thesis that the world of appearances *is* an illusion, for that would severely diminish the plausibility of her position. Recall, if you will, the common sense objection to nihilism that I considered in Chapter Four. Essentially, the objection was that nihilism must be false because it clashes with common sense. E.g. it says there are no tables when it is a common sense truism that there are tables. In response, I claimed that the common sense objection would lose all of its bite if the nihilist could provide a satisfactory explanation of why it *seems* as though there are tables, even though there are actually none. But I also conceded that if there was no such explanation available, that the common sense objection should be taken seriously. The situation is exactly the same here. To allow the as-if-facts to go ungrounded would be the same as if one were to reject the common sense

objection *without* giving any explanation of why it *seems* as though there are tables. It is for this reason that the monist must explain the as-if-facts.

Most metaphysicians (and probably most scientists too) would probably agree that common sense facts are not fundamental facts. Facts such as *there are tables* or *grass is green*, would have no place in a fundamental description of reality. A natural reason for thinking this is that these common sense facts are grounded in more basic scientific facts. So facts about tables or the colour of grass are going to be grounded in more fundamental facts about sub-atomic particles, fundamental properties, and wavelength reflectance profiles, or something along those lines. (The precise nature of these scientific facts is not overly important for present purposes as long as one recognises that whatever the scientific facts *are*, they will not involve entities like tables or properties like green-ness). It is *as if* this table in front of me is solid, and it is *as if* it is brown, but really it is mostly empty space, sparsely populated with tiny colourless particles.

But what about the scientific facts themselves? What are we to say about these? Well, one could simply say that the fundamental scientific facts *just are* the fundamental facts. So if the fundamental scientific facts involve entities like quarks and leptons having properties like mass and charge, say, then the world will be taken to be, fundamentally, one consisting of quarks and leptons having properties like mass and charge. Or if the scientific facts involve points of spacetime having certain field values, then the world will be taken to consist, fundamentally, of points of spacetime having certain field values. But, as Sider rightly points out, this view is only available to the pluralist. For whatever the ultimate scientific facts are, it is pretty certain that they will involve multiple concrete objects (whether those objects are sub-atomic particles, points of spacetime, or whatever).[319] The monist, then, cannot accept that the ultimate scientific facts are fundamental.

But that much is pretty obvious. *Of course* the monist will not accept that the ultimate scientific facts are fundamental – for monism conflicts with those very facts! For the monist, the scientific facts are just another set of as-if-facts, just like the facts of common sense. Sure, the world behaves *as if* our best scientific theories are true; that is, after all, why they are our best theories. But they are not true. And this is because the world is monistic, and the

[319] As Sider also notes, however, not all pluralists can *directly* endorse science like this. Punctal nihilists, for instance, may have trouble accepting certain facts about entangled systems whose properties do not reduce to properties of their constituent parts.

entities that our best scientific theories posit do not exist. So the task facing the monist is now fairly clear. She has to satisfactorily ground the ultimate facts of science.

To show how the monist should undertake this task, let's begin with a very rough outline. The monist will ground the as-if-facts of science by appealing to the world's fundamental properties. She will say that the world instantiates certain distributional properties that make it such that it is *as if* the scientific facts obtain. The world's distributional properties give rise to the empirical data that scientists collect, measure and analyse. But as it stands, this is very rough – it is not particularly informative. For whilst it tells us that the scientific as-if-facts *are* grounded in facts about the world and its properties, it doesn't tell us *how* they are so grounded. To leave it here would be to provide only what Ted Sider has called a "sketchy grounding story".[320] To make the explanation satisfactory, we will need to provide a little more detail.

In order to provide that detail, let me start with an analogy. Think back to the hypothetical extended simple I introduced in the last chapter: Cube. Cube is a mereologically simple object that is cube-shaped and multi-coloured. Specifically, each of Cube's faces is a different colour. Given the way that Cube appears, it would be fairly natural to describe it as having different parts (i.e. faces) that have different properties (i.e. colours). Of course, this description is false, strictly speaking, because Cube is, *ex hypothesi,* a mereological simple; it has no proper parts at all. But we must accept that it is *as if* Cube has parts – this is an evident as-if-fact. So how can the monist explain this as-if-fact? Well, armed with our new understanding of structured distributional properties, she has a fairly straightforward way of doing so. She will simply say that Cube instantiates six irreducible structured distributional properties. These properties give rise to, and thus ground, Cube's appearance. So, let's say that it is as if one of Cube's faces is blue. The monistic explanation of this fact is that Cube instantiates an irreducible structured distributional property – let's call it DP*x* – consisting of a distributable property, *blue-ness,* and a distribution pattern. The distribution pattern will be such that it appears as if Cube has a single face that is blue i.e. the *blueness* is distributed across the region which corresponds to Cube's alleged blue face. *Cube's* five other distributional properties will be similar, although they will consist of different distributable properties (i.e. different colour properties), and slightly different

[320] Sider (2008), 132

distribution patterns (i.e. they will distribute the properties to those areas which correspond to each of Cube's other faces).[321]

Now is this explanation still sketchy? I think not, or at least, not objectionably so. But I can envisage an objector thinking otherwise. In particular, I can imagine that one may be a little suspicious about the *distribution patterns* that I am invoking. For it may be thought that I have not really said much about what these distribution patterns *are like.* Rather, I have simply said that each distributional property involves a certain distribution pattern that distributes the appropriate properties in the appropriate ways. In order to allay such suspicions, my response would be that these distribution patterns are merely hard to describe, not hard to understand. It *is* hard for us to describe Cube's purported distribution patterns. Or, at least, it is hard for us to describe them in terminology that is not overtly pluralistic. We may need to talk of Cube's 'faces', for instance, or the regions of space that correspond with those faces. But this is, I repeat, merely a problem with description. For it is not difficult to *conceive* of an object like Cube, and it is not difficult to understand what the distribution patterns making up Cube's distributional properties must be like in order for it to look as it does. The mere fact that we can't accurately describe Cube's distributional properties in English without a misleadingly pluralistic slant should not cloud our understanding of what these properties are like; and neither should it force us into believing that Cube is mereologically complex. I am inclined to believe that the difficulties involved in describing these distributional properties should serve only to remind us of the descriptive limitations of our language, rather than tell us anything about what the object being described is really like. Indeed, if the descriptive powers of language are taken to put such heavy restrictions on our ontological conclusions, then monism would never have even got off the ground.

If one accepts the monistic explanation of the as-if-facts concerning Cube, then it should only take a short leap of the imagination to see how the same explanatory method could be employed on a much grander scale: to explain the *scientific* as-if-facts of the actual world. Let us suppose, for instance, that the ultimate scientific facts (i.e. those that constitute our best current scientific theories) involve sub-atomic particles (like leptons and quarks)

[321] In fact, one could say instead, that Cube instantiates only a single structured distributional property, consisting of the six distributable properties (i.e. the six different colours of Cube's faces) and the six different distribution patterns. Not much will hang on this point, however.

instantiating certain fundamental properties (like mass and charge). Just as she did with Cube, the monist can then translate these facts into facts about the world and its distributional properties. Firstly, she will take the fundamental properties posited by science (i.e. mass, charge, etc.) as the world's fundamental distributable properties. Then she will insist that rather than being instantiated by sub-world objects (like sub-atomic particles) they are instantiated by the entire world, in certain distribution patterns. These distribution patterns will correspond to the distribution of the alleged sub-world objects. Just like in the Cube case, this distribution pattern will be difficult, if not impossible, to describe in non-pluralistic terms. But just like in the Cube case, this should be seen only as a linguistic difficulty, not an ontological one.

Perhaps the ultimate scientific facts will be cast in very different terms than those I have suggested? One thing is for sure, they are likely to be very complex. But I do not pretend for a minute to know what the ultimate facts of science are. I am not a scientist after all. But I don't think that is too important, for it is the *principle* of the strategy that is the key thing. The monist has a *general strategy* for translating facts cast in pluralist terms into monist-friendly facts cast in terms of the world and its structured distributional properties. Take any fact that involves an alleged sub-world object having some property or other. Such a fact will have the general form, '*a* is *F*'. The monist will say that this is only an as-if-fact (because, of course, there are no sub-world objects), and it is more accurately expressed in terms of the world instantiating a structured distributional property. The *distributable* property component will be *F-ness*, and the distribution pattern will be such that it corresponds with the alleged location of *a.* I.e. the distribution of *F-ness* will overlap that, and only that, particular location. Once it is recognised that the monist has this strategy available, it should also be recognised that it doesn't matter what the alleged sub-world objects are purported to be (i.e. faces of a cube, sub-atomic particles, or whatever), or what the properties ascribed to them are (colours, mass, charge, or whatever).[322] It is the *principle* that is important. And if the principle is accepted, and I see no reason why it should not be, then it provides the monist with a satisfying way of grounding the as-if-facts. The world and

[322] Although perhaps one should insist that they must be qualitative, intrinsic properties. But this isn't too much of a problem, because I am assuming that facts involving *extrinsic* properties are reducible to facts involving only intrinsic properties.

its structured distributional properties are quite sufficient to explain the appearance of plurality.

## §7.5. How Did We Go So Wrong?

Before concluding, there is a fairly natural response to the thesis of monism that I am proposing that needs to be anticipated and dealt with. The response I am envisaging would run something along the following lines:

> *If monism is true and the world is, in fact, a single, part-less object, then how did we manage to go so catastrophically wrong? How could it possibly have come to be that we simply take it for granted that the world is populated by a multitude of objects, when in fact it is not? Sure, it would be foolish to think that we never get* anything *wrong about the way the world is, but the sheer magnitude of* this *error is surely* so *great as to render it entirely implausible.*

I sympathise with this response, but I also think it is flawed. One fairly obvious way to respond to this view is to appeal to the fact that we have got things drastically wrong about the nature of the world in the past, so there is good reason to suppose that there are a great many other features of the world that we are still wrong about today. To give an often used example, it was once taken to be obvious that the Earth is flat.[323] Presumably, this was taken to be *so* obvious, that to suggest that the Earth was anything *but* flat, would have been met with ridicule and scorn. But this was, of course, not only wrong, but *radically wrong*. Yet despite the sheer magnitude of the error, it is quite easy to see why people believed it. From our perspective the Earth *looks* flat, for a start. But likewise, from our perspective, the world *looks* like it is populated by a multitude of distinct physical objects. So using the same line of thought, it is easy to see why we think that it *actually is* populated by such objects.

There is, however, what could be an important difference between these two examples. The postulation that the Earth is flat is empirically testable; it may look flat to the

[323] There is, however, a debate amongst historians as to which peoples did and didn't believe that the earth was flat. The often told story that it was only when Columbus sailed the Atlantic that people first discovered the Earth's true shape is evidently a myth. Indeed it is generally thought that from about 300BC onwards the standard view in the west was that the Earth was spherical. Before this, however, it seems that a flat-earth view was the norm. For those interested, see Russell (1991) or Garwood (2007).

naked eye, but there is physical evidence, if one knows where to look, that shows that it must be curved.[324] The postulation of the monist, by contrast, does not seem to be empirically testable in this way. For it is hard to know what empirical data could ever be sufficient to demonstrate, *or to even indicate,* the mereological structure of the world. In other words, it seems like there is no way to tell, *just by looking,* whether monism or pluralism is true. This is just another case of underdetermination of theory by evidence. Science is afflicted by this phenomenon all the time, and metaphysics even more so. In fact, one may think that this makes it all the more understandable that we've got it wrong. If there's no way of telling, just by looking, whether the world is monistic or pluralistic, then how could we ever have even *expected* to get it right?

But this reasoning might be a little naive; or, worse, it might just be the hopeful musings of an optimistic monist. For another line of thought might be that for any dispute in which one of two incompatible theses must be true, then if the empirical evidence fails to point towards the truth of one thesis over the other, it would be reasonable to expect there to be a fairly even split in opinion about which thesis actually *is* true. Just think of the existence of God. No empirical evidence can prove or disprove His existence, and as such there is the resultant split in opinion about His existence that one may expect. But the same is not the case when it comes to the monism-pluralism debate. There is no widespread disagreement over whether there are many things or only one! For most ordinary people the very thought that there might only be one object in existence would surely have never even crossed their minds. And even among philosophers – those thinkers who are far more amenable than most to the acceptance of unconventional ideas – the unerring conviction in pluralism is almost universally held.[325] So perhaps this gives further weight to the initial objection. For if monism is correct, it means that our error-laden pluralistic tendencies are almost completely pervasive. But it is just not plausible that such a gross error could be so widespread, or so the thought may be.

I am not convinced. I think it is perfectly plausible that we could be grossly wrong about the fundamental nature of the world of which we are part. Indeed I think it suggests a

[324] Famously, Eratosthenes exploited this evidence over two millennia ago, to not only show that the earth was spherical, but to actually calculate its circumference.

[325] I say 'almost' universally held, because not *all* philosophers are pluralists. The only living philosophers I know of who defend monism in print, however, are Horgan and Potrč (2000; 2008). This represents an overwhelming majority of pluralists.

rather naive and misplaced confidence in our epistemic capabilities to be so unreceptive to the possibility that we could be so wrong about things. But at the same time, if we *are* radically mistaken about the nature of the world, it would be nice if we could at least trace the *source* of this mistaken view; if we could give some kind of explanation of how we have gone so far astray from the truth. I think that such an explanation can be given. Indeed, when trying to discern the source of the error of our pluralistic ways, I think that there is an obvious candidate: *language.*

Quine once said that "our language shows a tiresome bias in its treatment of time".[326] His point was that because our language contains all manner of tensed variants of verbs, it often leads us to presume that reality really is tensed when in fact (as far as Quine was concerned at least), it is not. I will remain neutral here about whether Quine was right or not about the reality of tense, but I think he *was* right in making the more general point that certain central features of our language often lead us to presume that there are corresponding features in reality, when there may not, in fact, be any such features at all. In this respect I think our language shows a number of tiresome biases in its treatment of all manner of things. But it shows no greater bias, I claim, than in its treatment of multiplicity.

Our language is *inherently pluralistic.*[327] It comes replete with a vast stock of nouns, each of which purports to pick out some *thing;* some object among many. Sentences are, in the vast majority of cases, structured in a subject-predicate form, whereby the subject is taken to be some object or other, and the predicate taken to refer to some characteristic of that object. Indeed the predication of such subjects often serves precisely to distinguish them from *other* subjects. But implicit in all this, of course, is the fact that there is a multiplicity of objects in the first place. Nothing about this is particularly revelatory – it is just the way our language works. But given that the way we communicate is so inherently pluralistic, it is really quite unsurprising that the way we *think* is so inherently pluralistic too. Language plays such a central role in our lives that it is quite ingrained in the way we think and the way we act; it is entirely second nature to us. Thus, again, it is unsurprising that the pluralistic picture of the world that our language paints is *also* ingrained in the way we think

[326] Quine (1960), 170

[327] When I talk of 'our language' I mean English, but I think that what I say will apply to most natural languages.

and act. The very structure of our language, a system which is so fundamental to our way of life, guides us inexorably towards pluralism.

The obvious line of response to these thoughts would be to say that our language is inherently pluralistic precisely *because* the world is also pluralistic. In other words, our language paints a pluralistic picture of the world *and it paints accurately.* But whilst this response might be reasonable, it certainly doesn't constitute a proof. Quine would surely have been unimpressed, for instance, if someone had objected: "but language *is* tensed, therefore reality *must* be too". The whole point of the thought being expressed is that the nature of language may mislead us into taking a mistaken view of the world. Now one may think that this kind of thought is mere idle speculation, and *sceptical* speculation at that. That is, one may concede that it is *possible* that our language is misleadingly pluralistic, and that the development of such a language is perfectly consistent with a monistic ontology, but one may also insist that that the best explanation of why our language developed as it did is because the world *really is pluralistic.* But I think that such a response would be unfair; because at this stage in the dialectic, the suggestion that our language might be misrepresentative of the world's true mereological structure is *not* mere speculation. I have spent the last six chapters of this thesis building a case for why we should believe that the world is, in fact, monistic; so to suppose that it *is* monistic can no longer be considered mere speculation. Rather, it is a serious metaphysical hypothesis backed by argument. And if this hypothesis is true, it will turn out that our language *is* misleadingly pluralistic after all.

The question still remains, however, of why language developed as it did. That is, why did we develop a language so inherently pluralistic if the world which we use it to describe is really monistic? I can't, of course, give a definitive answer to such a difficult question, but I can suggest some plausible candidates. Perhaps pluralistic languages such as English are just very effective means of communication? After all, it is beyond doubt that we do communicate very successfully. Perhaps we are simply hard-wired to think pluralistically, and this innate feature of human cognition is reflected in our language? Perhaps there is some evolutionary reason for why language has developed as it has? All of these explanations are perfectly plausible, I would suggest, and that is regardless of whether the world actually is pluralistic or monistic. Which of these explanations – *if any* – is right, however, is another question entirely. But hopefully the main point should be clear: it is not

unreasonable to suggest that our language *could* have developed in such a way that it grossly misrepresents the true nature of reality, and this could be the case regarding the reality (or not) of tense, or the world's mereological structure, or any number of (alleged) features of the world. Therefore, if our philosophical reasoning leads us to a conclusion that is at odds with the picture of the world painted by ordinary language then we should accept that the error could just as well lie on the linguistic side of things as it could in our reasoning.

To help press the point home, I invite you to consider a thought experiment that I have taken from Horgan & Potrč's book *Austere Realism*.[328] Imagine a world at which there was only a single object, a large extended simple. Suppose that this simple was made entirely of a jelly-like substance that was physically continuous regardless of what degree of magnification one viewed it at. Finally, suppose that this object displayed qualitative variation across its extension, in terms of properties such as colour, temperature and density. Of course, since this object is, by stipulation, mereologically simple, its qualitative variation is not to be explained in terms of its having parts which instantiate various different properties. It has no parts. Rather, it is to be explained by its instantiating various irreducible distributional properties. Thus it will have a certain *colour* distribution, a certain *mass* distribution, and so on. Given what has gone before, I trust that the reader will allow that such a world is at least possible. Indeed, for the experiment to work, one need only admit that such a world is *epistemically* possible.

Now the question is: how would one go about *describing* such a world? A natural way of doing this would be to employ a descriptive framework that posits certain entities (such as regions or points) and then ascribes certain properties to those entities. One might say, for instance, that the world has 'lumps', in order to describe the various fluctuations in density throughout the object. Or one might say it has 'stripes', for instance, in order to describe the colour variation across the object. Such descriptions would get something right about the world; they would track the genuine qualitative variation that the object exhibits. There would, as Horgan & Potrč rightly point out, "be systematic correspondence between statements couched in this discourse, on the one hand, and how things are with the jello-

[328] Horgan & Potrč (2008), 168-171. I have modified the details of the thought experiment slightly, but the gist if it is very much the same.

world, on the other hand".[329] But these descriptions would not be entirely perspicuous. Indeed, they would be really quite misleading, in that they suggest that the world has parts (lumps, and stripes, and so on) when in fact it has none. The point is, however, that it would be perfectly natural for us to describe the world in this way and, moreover, it would be perfectly natural for us to suppose that, on the back of our descriptions, the world *really did* have parts (i.e. lumps and stripes and so on). But, of course, the world, *ex hypothesi,* does *not* have parts. But if this scenario is reasonable, and I will assume that it is, then it takes only a short leap of the imagination to see that a similar scenario could obtain at the actual world. We mistakenly describe the world pluralistically, and thus we mistakenly take it to *be* pluralistic.

It is very hard for us to even conceive of a way of describing such a world that is *not* pluralistic. At least, it is certainly very hard to see how we could describe such a world in a perspicuous manner, *in English,* that did not employ nouns that purport to refer to sub-world objects or parts/regions of the world. Perhaps other types of languages are possible that would be perfectly suited to the job? If so, then I can only think that they would be very far removed from the ordinary languages we actually employ.[330] I won't be attempting to come up with such a language here; that would be a substantial project in itself. But whether such a language could be produced is really beside the point. If it turns out that a monistic world is difficult to describe (in English or *any* language), then that should have no bearing on whether such a world is *possible.* It simply means that it is hard to describe. Quantum Field Theory is, I am told, quite difficult to explain, but that doesn't make it any less likely to be true.

To conclude, then, there are all sorts of plausible stories that one could tell in order to explain how it is that we have gone so drastically wrong in our understanding of reality. But, when providing such stories, one does have to tread with caution. For all manner of wild and wonderful stories can be told that are perfectly *consistent* with our experiences and our understanding. Russell's teapot comes to mind here, for instance.[331] But merely *consistent*

[329] Horgan & Potrč (2008), 169

[330] Hawthorne & Cortens (1995) attempt to flesh out such a language by replacing talk of objects with suitable adverbs. E.g. "there is a white pebble" would be translated as "it is pebbling whitely there". I think that their attempt is most admirable, but it only serves to reinforce the point: a language that doesn't refer to multiple objects is going to be far-removed from the languages with which we are ordinarily acquainted.

[331] See Russell (1997), 543

stories are nothing more than idle speculation if there are no independent reasons to believe that they may be true. But the story told by the monist is not like this at all. There are many independent reasons to believe that monism is true, as I have elucidated over the past few chapters. So given that there is independent motivation for believing monism, and given that monism is at odds with ordinary understanding, it is the *duty* of the monist to provide a story of how it is that we have come to be so mistaken. I hope that the preceding few paragraphs have gone some way towards fulfilling that duty.

## §7.6. Concluding Remarks

I hope that the reader now has a much better idea of what I take the monistic world to be like. The rich and diverse qualitative variation that we experience in the world is to be explained by its instantiating irreducible structured distributional properties. I hope to have shown that distributional properties are in fact most intuitive and easy to grasp. Moreover, by expanding on the arguments of Josh Parsons, I hope to have shown that there are some compelling reasons to believe that there are such things, and that these reasons hold independently of one's views on the monism versus pluralism debate. I then provided a more detailed picture of how distributional properties are to be conceived. Specifically, I claimed that distributional properties should be taken to be structured universals, comprising at least one distributable property, and at least one distribution pattern. This structured conception gives the monist an easy and satisfying way of explaining the appearances; a method for explaining the appearance of plurality, given that there is, in fact, only singularity. As a final, but no less important, point, I hope to have impressed the importance of distinguishing mere linguistic difficulties from genuinely ontological difficulties. Given that our language and thought is so inherently pluralistic, it is no surprise that we struggle to provide accurate descriptions of the monistic world that don't involve pluralistic terms. Indeed, since our language is unavoidably pluralistic, it should be *expected* that any descriptions couched in that language will also be unavoidably pluralistic. But it would be a mistake to infer from this that the world must be pluralistic too.

# Part III

# Objections & Replies

# §8. Objections and Replies

As the *denouement* to this thesis it remains for me to address, and respond to, some objections that have been presented against monism. By far and away the most commonly given objection is the common sense objection i.e. that which states that monism is obviously false because it denies the existence of ordinary objects. But I have tackled that objection in a number of places in the preceding chapters already (in a manner that I hope the reader will have found satisfactory), so I won't be considering it again here. Instead, I will be considering objections of a somewhat more technical nature. Specifically, I will be responding to four objections that have been presented by Ted Sider.[332] Unlike many of his contemporaries, Sider at least takes monism seriously; whilst he does not endorse it, he accepts that it is a coherent metaphysical thesis and that it cannot simply be dismissed out of hand. Because of this, his objections are much better thought out than most, and represent a serious challenge for the monist. In the present chapter I will show how that challenge is to be met. As well as addressing Sider's objections, I will also spend a short time considering an alternative theory of monism: *priority monism*.[333] It is sometimes claimed that priority monism is a variant of monism far superior to the variant I am defending here, and thus in the battle between the two, it is priority monism that should prevail. In §8.1 I will show why this claim is false.

It is not possible for me to address *every* objection that either has been, or could be, levelled against monism. So there are some objections in the literature that I am aware of, that I will not be considering in what follows, and there may well be other objections of which I am not even aware. However, the objections that I *will* respond to have been chosen because I take them to be the most significant, and because I take them to represent the most serious threat to my project. Those that I have decided to omit have been neglected because I take them either to have fairly obvious and easy solutions, or because I don't take them to

[332] Sider (2007)

[333] This is the thesis proposed and defended by Jonathan Schaffer. See Schaffer (2007; 2009; 2010a; 2010b; 2010c)

pose much of a threat to monism.[334] With that said, let me get on to addressing the objections.

## §8.1. Priority Monism

In recent years, Jonathan Schaffer has brought monism very much into the limelight of contemporary metaphysics.[335] He has, almost singlehandedly, revived interest in what had until recently become little more than a relic from deep within the annals of the history of philosophy.[336] Indeed, it is rare that monism is even mentioned in contemporary philosophical literature *without* Schaffer's name being mentioned too. So it would probably be amiss of me were I not to devote at least a brief amount of time to Schaffer's work, given the topic of this dissertation. Schaffer has identified two very different forms of ontological monism. They are:

**Existence Monism**: There is only one concrete object in existence – the world itself
**Priority Monism**: There is only one *fundamental* concrete object in existence – the world itself

The former of these doctrines is the one I am defending in this dissertation. It is the latter, however, that Schaffer himself endorses. The central difference between the two theses is that Priority Monism allows that the world has parts; it endorses a multiplicity of concrete objects. The Priority Monist can allow that there are cats and dogs, and trees and rocks, and so on. But she merely maintains that these objects are not fundamental. They are ontologically posterior to, and ontologically dependent upon, the only fundamental object there is: the world itself. Priority Monism, then, involves a "hierarchical view of reality",

[334] Goff (2012), for instance, has presented an argument against monism, but I think it amounts to little more than a formal reconstruction of the common sense argument, and therefore can be overcome in exactly the same way as that argument. Essentially, he argues: I exist; you exist; therefore there is more than one thing.
[335] See Schaffer (2009; 2010a; 2010b; 2010c). Also, for an interesting introduction and overview of varying forms of monism in ontology, see Schaffer (2007b).
[336] I say 'almost' singlehandedly because there have been a few other key contributions. In particular, I am referring to the work of Horgan & Potrč (2000; 2008)

where objects are ordered in terms of *ontological priority*, and where the world itself is ontologically prior to everything else.[337]

It is often suggested that Priority Monism is by far the more reasonable of these two monistic theses. Schaffer himself, for instance, claims Existence Monism to be a "crazy view".[338] Elsewhere it has been said that "Existence Monism is an obviously weaker thesis [than Priority Monism] that can only be considered little more than a caricature".[339] Indeed, it has often been suggested to me in conversation that, given the availability of Priority Monism, I should feel extra pressure to justify my persistent interest in Existence Monism. In essence, the question that is often put to me, and indeed that is implicit in this type of view, is: *why be an Existence Monist when you could be a Priority Monist instead?*

The answer to this question is really very simple. One of the central motivations for my endorsement of monism is a strong conviction in the truth of compositional nihilism. Indeed, I spent the first four chapters of this dissertation arguing that compositional nihilism is true. But Priority Monism is entirely incompatible with nihilism – it states that the world has parts! – and it is for that reason that I reject it.[340] Furthermore, I think the question of why one would endorse Existence Monism when Priority Monism is available is somewhat misguided. The question implies that the two theories are somehow direct rivals, when they are in fact not. Of course, in a *trivial* sense they are rivals, in that they are both ontological views which differ from one another. But on that understanding of rivalry, all sorts of disparate and unrelated ontological theses would be considered direct rivals, and thus very little would be gained by so asserting. But Priority Monism and Existence Monism are often seen to be involved in a much closer rivalry than that, *because they are both forms of monism*. This, I claim, is where the central mistake lies, for one of these theses is an imposter; Priority Monism is monistic in name only.

This claim is not intended to discredit, or undermine, Priority Monism in any way. On the contrary, I think it is an independently very interesting thesis. But just because it has the term 'monism' in its title, does not mean it should be considered a direct rival of, or an

[337] Schaffer (2009), 351
[338] Schaffer (2010a), 324
[339] Morganti (2009), 272
[340] Strictly speaking, this is not true; Priority Monism is not *incompatible* with nihilism. One could be a priority monist *and* a nihilist, by claiming that the world is the only object there is. Such a view would simply 'collapse' into Existence Monism, however. Moreover, the Priority Monism that Schaffer endorses *is* incompatible with nihilism, since he endorses the existence of multiple concrete objects that are taken to be parts of the world.

available alternative to, Existence Monism. Suppose, for instance, I postulated a thesis called 'Chromatic Monism'; the thesis which states that objects can only be a single colour. Even if I could provide a whole dossier of reasons to believe such a thesis plausible (which is unlikely), it would still not be right to consider it a direct rival of, or any sort of *alternative* to, Existence Monism, for its central focus is on something entirely different: the *colour* of objects, rather than the *number* of them and their *mereological structure*. It would be absurd, would it not, to ask: but why are you an Existence Monist when you could be a Chromatic Monist? Such a question would seem guilty of making a category mistake. And it is likewise between Priority Monism and Existence Monist; they focus on entirely different things. The former is concerned with the hierarchical dependence relations that obtain between objects, whereas the latter is concerned with the mereological structure of objects and their number. Indeed, as Ted Sider has noted, "priority monism isn't really an ontology at all".[341] Sider is right; Priority Monism is only concerned with priority relations; it is not at all concerned with *what* exists. There may well be any number of independently compelling reasons to be attracted to Priority Monism, but none of this should affect the Existence Monist. For the Existence Monist, Priority Monism should be rejected for exactly the same reason as many other pluralistic ontologies: *it admits that the world has parts.*

## §8.2. Ted Sider's Objections

In a 2007 paper, Ted Sider presents four separate arguments against monism. Sider's arguments are insightful and challenging, and they represent a pressing problem for the monist to deal with. Indeed the monist *needs* to be able to deal with these types of problem if her thesis is to remain plausible. In what follows I will show what Sider's objections are, and how they are best overcome.[342],[343]

[341] Sider (2008), 130n

[342] The first two responses given (i.e. §82.1 & §8.2.2) are taken from Cornell (2013).

[343] It should also be noted that Sider himself has published his own responses to his own objections. See Sider (2008). I think Sider's solutions work, to an extent, but they are markedly different and much more complex than the solutions I will offer here.

## *§8.2.1 Statespace Size*

Sider's first objection is that a monistic view of a given world cannot provide a satisfactory explanation of the size of the statespace at that world.[344] To illustrate his point, he asks us to consider a simple two-dimensional world (let's call it a 'Screenworld') made up of sixteen pixels, arranged in a 4x4 grid, each of which can be either 'on' or 'off'. The statespace at such a world would have $2^{16}$ members, as illustrated below:

*Fig. 9*

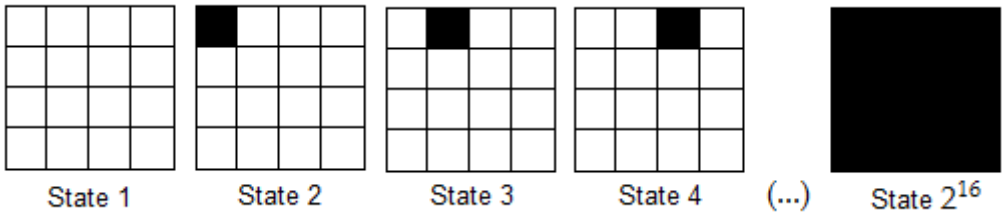


Sider argues that only a pluralistic view of such a world can satisfactorily explain *why* the statespace size is $2^{16}$. The Screenworld has $2^{16}$ possible states precisely *because* it is made up of sixteen pixels, each of which has two possible states. More generally, the size of the statespace at any Screenworld will be derived *combinatorially* from more fundamental facts about the number of objects at that world (the pixels) and the number of properties those objects can instantiate: "the possibilities for the entire system are generated combinatorially from the number of entities within the system and the fundamental states those entities can inhabit".[345]

Since the monist does not countenance the existence of the pixels, it is supposed that she will not be able to avail herself of this explanation. In fact, according to Sider, the monist has no option but to accept the statespace size as a brute fact. For the monist, every property at the Screenworld must be a property of the *entire* Screenworld. (This is the case since there are no other objects at the world available to instantiate properties). Therefore, each of the $2^{16}$ members of the statespace must involve a distinct and fundamental (i.e. irreducible) property of the entire world. So when asked *why* there are $2^{16}$ members of the statespace,

[344] 'Statespace' refers to the set of physical possibilities at a world.
[345] Sider (2007a), 3. One should of course remember here that we are using Screenworlds only as a model, due to their simplicity. It is assumed that any principles uncovered should also be applicable to more complex worlds, like the actual world.

according to Sider, the monist can do little more than shrug her shoulders in response. This is most unsatisfactory, especially when contrasted with the pluralist explanation, or so the argument goes.

What makes it worse for the monist, or so Sider supposes, is the *regularity* in statespace size across all Screenworlds. He claims that his explanation generalises to cover *all* Screenworlds. For instance, if we were to consider Screenworlds with different numbers, $n$, of pixels, then the statespace size at these worlds will always be $2^n$. Sider claims that "*no* world of this sort will *ever* have a statespace with a cardinality that is not a power of 2".[346] His point is that if one does not countenance the individual pixels, then one cannot explain this regularity, but rather, one must accept it as another (extraordinarily coincidental!) brute fact. For Sider, this is so wildly implausible that it constitutes a powerful objection against monism.

By way of a response, I should first concede that much of what Sider claims is correct. He is right to suggest that the statespace size at a given world will be determined by more fundamental facts about how many objects there are at the world and what fundamental states they can inhabit. Furthermore, he is right to suggest that this explanation generalises. Where he goes wrong, however, is by suggesting that this fact undermines the monist's position; it does no such thing. Let me explain.

Calculating the statespace size at a given world is a simple matter of calculating permutations given certain constraints. When calculating the number of permutations of a given number series, for instance, one need know only how many variables are in the series, and how many possible values each variable can take. Given this information, the total number of permutations will be $x^n$, where $n$ = the number of variables in the series, and $x$ = the number of values each variable can take.[347] Exactly the same process applies when determining the statespace size for a given Screenworld. The series itself represents the entire statespace, with each possible permutation of the series corresponding to an individual member of the statespace (i.e. a particular state of the entire system). The variables in the series correspond with the objects posited, and the possible values of those

[346] Sider (2007a), 3

[347] This is only the case for regular series, where each variable in the series has the same number of possible variables. For series whose variables have differing numbers of possible values, the formula to calculate the permutations will be more complex.

variables correspond with the possible states that those objects can inhabit. Thus the same formula that applies to number series will also apply to Screenworlds, such that we can say for *any* Screenworld:

Statespace Size = $x^n$

Where $n$ = the number of objects at the world, and $x$ = the number of properties each object can instantiate.

This formula vindicates Sider's claim that the statespace size is determined by, or derived from, facts about how many objects are at a world and how many properties they can instantiate. Furthermore, it demonstrates that the explanation is generalisable such that it covers all Screenworlds; for it shows that, given *any* Screenworld, one only need know how many things there are at the world, and what properties those things can instantiate, and one will be able to calculate the size of the statespace. However, what is equally clear is that this explanation is also available to the monist. Take our previous example Screenworld, for instance. Given this statespace, the monist will posit one object, the world, (i.e. $n = 1$) and $2^{16}$ fundamental distributional properties (i.e. $x = 2^{16}$), thus the statespace size = $x^n = (2^{16})^1 = 2^{16}$. And this will be the case at *any* given Screenworld.

So contrary to Sider's claim, the monist does *not* have to accept the statespace size as a brute fact, nor does she have to accept the regularity of statespace size as a brute fact. Instead, she can maintain that facts about the statespace size are derived directly from facts about how many objects there are at a world, and how many properties they can instantiate. Yes, it is a brute fact that on the Monistic picture there are $2^{16}$ fundamental properties and one object, but on the pluralistic picture it is a brute fact that there are two fundamental properties and sixteen objects. But why should the acceptance of the latter brute facts be any more agreeable than the acceptance of the former? There seems no reason why it should be, and unless there *were* some such reason, Sider's objection should carry no force at all.

## §8.2.2 Statespace Structure

According to Sider, it is not only facts about the *size* of the statespace which the monist is at a loss to explain. There are also important facts about the way in which the statespace is *structured* which cannot be satisfactorily accounted for by a monistic view, or so it is claimed. In particular, Sider claims there to be natural groupings, or subsets, within the statespace, which can only be explained in terms of facts about how many pixels are 'on' in each case. For instance, all those members at which only one pixel is on "go together", as do those members at which only two pixels are on, and so on and so forth. Sider claims that these 'natural' groupings can only be explained if one countenances the existence of pixels: "they emerge from the natural groupings of the most fundamental facts, facts concerning which pixels are lit".[348]

Once again, some of what Sider says is perfectly correct. There *are* certain members of the statespace that "go together" more naturally than others. This is because certain members of the statespace exhibit more objective similarity to one another than do others. But once again he is wrong to suggest that the monist cannot explain this phenomenon. The monist can explain these natural groupings perfectly well, and she has no need to posit sub-world objects (pixels) in order to do so. The reason that there are certain natural subsets of the statespace can be explained by appealing to the world's fundamental properties, and, crucially, by noting *what those properties are like.*

Sider's objection can be met, and overcome, through a consideration of what distributional properties the world instantiates, and more precisely, what those properties are like. To explain, consider State 2 in *fig. 9*. The pluralist will describe this state in terms of which pixels are on and which are off (probably by saying that one pixel is on and the rest off). The monist, by contrast, will describe it in terms of what distributional property the world instantiates. In my view, she will say that the world instantiates a structured distributional property that comprises the distributable properties of *off-ness* and *on-ness*, and a particular distribution pattern. What we *call* the property is not of great importance, provided one grasps the idea: the world instantiates some irreducible distributional property that gives rise to the appearance depicted in state 2 in *fig 9.* But this is not the limit of what the monist can say; for she can also describe in detail *what that property is like.* For

[348] Sider (2007a), 4

each of the world's distributional properties will have their own particular features; they will have *higher-order properties.* For instance, let's call the distributional property instantiated by the world in State 2, $DP_1$. We can say of $DP_1$ that it distributes more *off-ness* than *on-ness.* More precisely, we could say that the proportion of the world across which *off-ness* is distributed is fifteen times greater than the proportion of the world across which *on-ness* is distributed. I will call this feature of the world's properties a *distribution ratio.* Thus $DP_1$ has a distribution ratio of 15:1 of *off-ness* to *on-ness.* Indeed any particular state of the Screenworld which the pluralist would describe by saying one pixel is on and the rest off would, on the monist's picture, be described as instantiating a distributional property with a distribution ratio of 15:1. Those states which the pluralist would describe as having only two pixels on would, in monistic terms, have distribution ratios of 7:1; those described pluralistically as having only three pixels on would have distribution ratios of 13:3; and so on and so forth.

It is because of this that the monist has a perfectly good way of explaining the natural groupings, or subsets, within the statespace. As Sider rightly says, the natural groupings emerge from the natural groupings of the world's most fundamental facts. But these are not facts about what pixels are lit (there are no pixels!), rather, they are *facts concerning what the world's fundamental properties are like*. And these will include those facts about the distribution ratios of each member of the statespace. There will be natural subsets of the statespace, then, consisting of all those states which share the same distribution ratio. Thus the monist can readily accept that there are facts about objective similarity between certain members of the statespace; she can readily accept that there are natural groupings, or subsets, within the statespace. But contrary to Sider's claim, she does not need to posit sub-world objects in order to account for them.

## *§8.2.3 Haecceities and Probabilities*

Sider's third objection is that the Monist cannot account for purely haecceitistic differences between members of statespace. Indeed he claims that the Monist cannot even accept that there are such differences. To illustrate his point, Sider once more asks us to consider a hypothetical Screenworld as an example. This time we are asked to imagine a Screenworld

made up of just two pixels (call them 'A' and 'B'), both of which can be either on or off. Let's call this world $w1$. Now according to Sider, if haecceitism is true, then there will be a difference between the state in which A is on and B is off, and the state in which A is off and B is on. Because of this, there will be four members of the statespace at this world:

*Fig. 10*

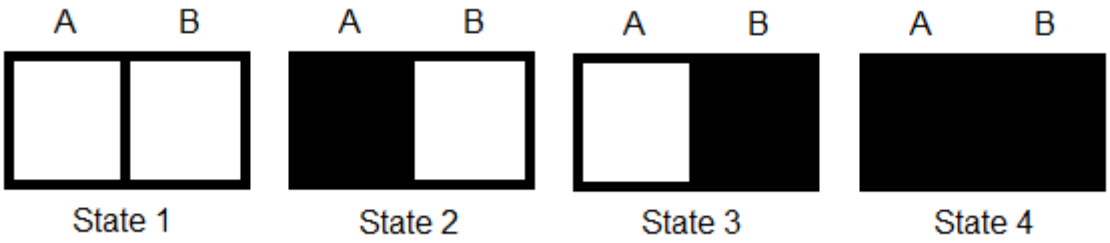


Furthermore, Sider suggests that it is quite important that there should be four members of the statespace at such a world. For it would play a part in explaining, for instance, why it is more likely that the world will be in a one-pixel-on-one-pixel-off state than in, say, a both-pixels-off state. In other words, one needs to recognise the difference between state 2 and state 3 in order to explain certain probabilistic features of the world. And in order to do that, he claims, one must accept *haecceities*. The reason for this is that states 2 and 3 are qualitatively indistinguishable. That is to say, one cannot distinguish between the two states in terms of what qualitative properties are instantiated, but only in terms of what objects instantiate them – i.e. in state 2 it is *that* pixel (A) which is on, and in state 3 it is *that* pixel (B) that is on. If one views the world monistically, then one cannot even distinguish between the states on those grounds, for in both states, not only are the qualitative properties indistinguishable, but they are instantiated by the very same object (the entire world) in both instances. For the monist, then, states 2 and 3 are identical. Thus for the monist, there are only three members of the statespace (as depicted in *fig. 11*). But if there are only three members of the statespace, then how can one explain the probabilistic features of the world to which Sider alludes? The pluralist has an easy answer, the monist does not.

*Fig. 11*


Sider's argument rests on two important assumptions. They are:

Assumption 1: The probabilities at $w1$ are as follows:

> Both pixels off: 0.25
> Both pixels on: 0.25
> One pixel on one pixel off: 0.5

Assumption 2: All members of the statespace at $w1$ are equi-probable.

The conjunction of these two assumptions entails that there must be two distinct one-pixel-on-one-pixel-off states, and thus four possible states in total. However, I think that both assumptions can be questioned, and therefore that there are (at least) two possible lines of response available to the monist.

The simplest line of response would be to deny assumption 2. On this view, then, the monist would accept that there are only three members of the statespace (as depicted in *fig. 11*) but assert that they are not equi-probable. Instead, the probabilities would be stipulated thus:

> State 1: 0.25
>
> State 2: 0.5
>
> State 3: 0.25

Whilst this response would overcome the objection, it is evidently problematic. The problem being that there doesn't appear to be any good explanation of *why* state 2 has the probability it does. The monist will just have to accept it as a brute fact that state 2 has some kind of primitive bias associated with it such that it is twice as likely to crop up as either of the

others. Whilst this may not be a *fatal* concession for the monist to make, it looks suspiciously *ad hoc,* and thus would be something of an unsavoury pill to swallow.

A much more promising, and indeed interesting, avenue of response, however, would be to challenge assumption 1. That is, to question whether the probabilities at *w*1 are in fact as Sider suggests they are. According to assumption 1, the probability of *w*1 being in a one-on-one-off state is twice as likely as it being in a both-off state (or a both-on state). The question I suggest that we should ask is *why* should we assume that this is the case?

Here is a possible answer: consider flipping two coins, call them 'A' and 'B'. This situation seems to be analogous to the situation described at *w*1, for the two coins represent the two pixels and each possible state of the coins (heads or tails) represents each possible state of the pixels (on or off). Now if we flip both coins once each, then the probabilities of the possible results will exactly match those probabilities stated by Sider above:

Both heads: 0.25
Both tails: 0.25
One tail one head: 0.5

That the probabilities for the coins *are* like this is beyond doubt. (One only needs to flip enough coins to see this for oneself). So the monist certainly won't want to deny *that.* But she doesn't have to. When speaking strictly, the monist will obviously not accept that there are such things as coins (or indeed people to flip them) for she does not accept the existence of multiple concrete objects, but she can explain the coin flipping scenario perfectly well, in terms of the world instantiating certain distributional properties. (Very roughly, she can say that there are certain distributions of *coin-ness*, or *head-ness* and *tail-ness*, instantiated by the world, or something to that effect. To repeat what was stated earlier, what we *call* these properties is not really that important). Now the pluralist will say that the one-head-one-tail result has a probability of 0.5 because there are two distinct ways in which it can obtain (i.e. 'A-heads & B-tails' or 'A-tails & B-heads'). Whilst the monist won't actually posit the coins, her corresponding explanation will be very similar; there will be two different ways in which the result can obtain (i.e. there will be two distinct distributional properties, one corresponding to 'A-heads & B-tails' and the other corresponding to 'A-tails & B-heads').

The important point to recognise, however, is that neither pluralist nor monist will have to appeal to haecceities in order to distinguish these states. And that is because they are *qualitatively* distinct in both cases. In the pluralist case, for instance, because the coins are *sub-world* objects, there will be a multitude of relational properties in play which suffice to make the two states qualitatively distinct. For instance, the pluralist can distinguish between coin A and B by pointing to the relations in which they stand to other objects and/or observers, e.g. 'coin A is the one on the left' or 'coin B is the one closer to Roger' or something like that. The monist too can make such distinctions. For the world's distributional properties that correspond to the states 'A-heads & B-tails' or 'A-tails & B-heads' will likewise be qualitatively distinct. The distribution patterns of head-ness and tail-ness will be spatially, and thus qualitatively, different in each case.

In stark contrast, however, this is clearly not the case at $w1$. Because the pixels are the *only* things there, it would make no sense for the pluralist to start talking about things like 'the pixel on the left' or 'the pixel closest to Roger'. Likewise for the monist, there is no way to qualitatively distinguish between two half-on-half-off states at $w1$, because the distribution patterns will be perfect duplicates in both cases. Now what all this suggests, I propose, is that contrary to our initial presumption, the coin flipping scenario is *not* analogous to the situation described at $w1$ after all. And because the two scenarios are not analogous, or at least not as clearly analogous as they first seemed, then it seems questionable to infer that the probabilities in one scenario will be the same as in the other. Sider has made exactly that inference (at least, I am presuming he has made that inference, for I can see no other reason why he would assume the probabilities at $w1$ to be as he states), and then accused the monist of being unable to explain the results of it. But if the inference itself is unwarranted, then there is no pressure on the monist to provide the explanation that is being requested.

At this point, however, this response may still seem a little weak. For it just seems *obvious*, does it not, that the probabilities will be as Sider describes. If there are two pixels, both of which can be on or off, then *surely*, it is more likely that at any given time they will be in a one-on-one-off state that in, say, a both-off state! I sympathise with this intuition, but I think it may well be misguided. To show why, I would like to draw on an example from

contemporary physics.[349] Suppose that instead of considering the flipping of two coins, we instead focused on two elementary particles (e.g. electrons) each of which could be in one of two possible states (call them 'state X' and 'state Y'). One may think, that if the coin-flipping case is anything to go by, that the chance of both electrons being in state X at any given time would be 0.25 (as it would be for them both being in state Y), and the chance of them being in different states (i.e. one X and one Y) would be 0.5. Indeed if the intuition identified above is correct, then it is *obvious* that the probabilities should be like this. Interestingly, however, this is not the case. According to our current best theories, the probabilities of the electrons being in different states is actually 1/3, as is the probability of them both being in state X, and of them both being in state Y. That the probabilities *are* like this at the quantum level has been corroborated by empirical observation.[350]

There are two conclusions one could draw from these results. One would be to maintain that while there are still four different overall states our electron pair could be in at any given time, there is nevertheless some reason that they are not all equi-probable. That is, for some reason, there are primitive biases associated with the various states such that the probabilities come out as they do. The other conclusion would be to maintain that all possible states of the electron pair *are* equi-probable, but that there are actually only three of them (i.e. both X, both Y or one X one Y). The former of these conclusions is, I think, no better off than the first monistic response I considered to Sider's original objection. That is, it seems suspiciously *ad hoc.* The latter of the two conclusions seems much more reasonable, and furthermore, is the one favoured by experts in this field.[351] In other words, at the quantum level, if two states are qualitatively indistinguishable, they are taken to be a single state.

What all this shows is that we would have been wrong to infer from the probabilities that obtain in everyday cases (like the coin-flipping case), the probabilities that obtain in quantum cases (like the electron pair). Classical statistics just seem not to apply when considering the very very small. So the thought being proposed here, then, is perhaps they don't apply when considering the very very big either? Now I should make something clear here. I am not for a minute suggesting that the probabilities that obtain at the quantum level

[349] This example was brought to my attention when reading a recent paper by John Heil. See Heil (2012), 176-7.
[350] See French and Krausse (2006)
[351] See French & Krausse (2006), 83-84

give us any license whatsoever to infer what the probabilities may be at the maximal, worldly level. They do not. My argument is much more modest than that. I am merely highlighting a case where it would have been wrong for us to infer what probabilities obtain purely by observing the probabilities that obtain in ordinary circumstances (like the coin-flipping case) which are *supposedly* analogous. The fact that classical statistics and probabilities do not apply at the quantum level entails the more general conclusion that they do not always apply. But if we know that they do not always apply, then it doesn't seem too much of a problem for the monist to say that it is a consequence of monism that classical statistics do not always apply at a worldly level either.

To conclude, then, what Sider has successfully done is demonstrate some interesting probabilistic facts about monistic statespace. At the maximal level, if monism is correct, it seems as though classical statistics and probabilities do not always apply. But one can only really view this as an argument *against* monism if one has some prior reason to believe that classical statistics *should* apply at the maximal level. I cannot see what such a reason could be, other than some kind of inference from what the probabilities are in ordinary sub-world cases like the flipping of two coins. But since it has been demonstrated that such an inference would produce the wrong results when moving to considering the very very small, why should we have any faith that it would not produce the wrong results when we move to consider the maximally large? Indeed one might even think that there would be a certain satisfying symmetry if the probabilities at a maximal level exactly mirrored those at the minimal level. But, of course, we are now in the realms of speculation. The overriding point should be, however, that if the monist has to accept that some unusual probabilistic facts obtain in monistic statespace, then so be it; that is simply a consequence of monism. But in light of the preceding comments, it doesn't seem to be too costly a consequence. Indeed, I am inclined to believe that it isn't really a cost at all.

### *§8.2.4 Intrinsic and Extrinsic Properties*

Sider's final objection against monism is that only if one is a pluralist can one give a satisfactory definition of *intrinsicality.* But because we require a satisfactory definition of intrinsicality, pluralism should be preferred, or so the argument goes. Philosophers often

make claims about which of an object's properties are intrinsic and which are extrinsic. Consider, for example, the following two statements:

1. Being negatively charged is intrinsic
2. Being ten feet from an electron is extrinsic

Sider's claim is that although the monist won't take claims like this at face value (for according to monism there *aren't* any electrons), she will still want to make some kind of sense of them. Specifically, the monist needs to be able to accept that they are, *in some sense or other*, correct. This is supposed to be analogous to the sense in which the monist will accept that common sense claims like 'there are tables' are in some sense correct, even though they are, strictly speaking, false. But in order to accept this, Sider continues, the monist will need to have some kind of satisfactory definition or analysis of intrinsicality. Sider himself endorses a broadly Lewisian definition of intrinsicality:

> **INTRINSICALITY:** A property is intrinsic *iff* it can never differ between a pair of possible duplicates.

Where duplication is defined as follows:

> **DUPLICATION:** objects are duplicates *iff* their parts can be put in one-one correspondence preserving their perfectly natural properties and relations.[352]

Sider's argument is that these definitions are no good to the monist. And he is quite right. For the Lewisian account of intrinsicality, when coupled with monism, generates all sorts of counter-intuitive results when it comes to intrinsic and extrinsic properties. Take the property of *being ten feet from an electron*, for instance. According to the monist, nothing instantiates this property, because there are no electrons, *of necessity*.[353] Therefore, no possible duplicates will differ with respect to this property – because no possible objects will have it. So according to the Lewisian account, the monist must take *being ten feet from an electron* to be an intrinsic property. But if that is the case, then how on earth can the monist

[352] Sider (2007a), 5

[353] Presuming, of course, that monism is *necessarily* true. I will assume this to be the case.

account for the fact that statement 1 should be considered correct? It would seem that she can't. The monist, Sider therefore claims, owes us a story.[354]

The first thing to point out in response to this argument is that it doesn't show that the monist can't make sense of intrinsicality *per se*, but rather, it shows merely that the monist can't adopt Lewis's account in order to do so. That may represent a problem if the monist were, for some reason or other, particularly wedded to Lewis's account. But I can see no reason why she should be. For the Lewisian account is by no means the *only* way of understanding intrinsicality. (In fact, it would be fair to say that Lewis's account is a fairly controversial one). There are plenty of other ways one can try to define intrinsicality.[355] One fairly common way of doing so is to focus on the seemingly evident fact that most extrinsic properties tend to be *relational* in nature, and conversely, that most intrinsic properties tend to be *non-relational*.[356] Specific relational accounts of intrinsicality that can be found in the literature can be quite technical and complex, but the general thought that seems to pervade them all is that if *x*'s being *F* obtains in virtue of *x* standing in some relation to a distinct object then *F* is *extrinsic,* and if it does not, then *F* is *intrinsic.*

I think that the monist should endorse a relational account of intrinsicality. This is for two main reasons. Firstly, relational accounts seem to accurately capture a central and important intuition about the intrinsic-extrinsic distinction. That intuition is that an object's intrinsic properties are those it has in virtue of the way it is in itself, whereas its extrinsic properties are those it has in virtue of how it stands to its surroundings. This intuition seems to be accepted by *all* parties who engage in the debate over intrinsicality, regardless of how they choose to formally analyse the distinction.[357] The fact that relational accounts seem to capture this intuition so clearly must surely count in their favour. Secondly, a broadly relational account will enable the monist to generate, for the main part, the required facts about intrinsicality, and thus overcome Sider's objection.[358] To return to our previous

[354] Sider (2007a), 7

[355] See Cameron (2008c) for a good overview of the terrain.

[356] Relational accounts of intrinsicality have been given by Francescotti (1999) and Hoffmann-Kolss (2010), among others.

[357] Even Lewis agrees that much. See Lewis (1986), 61-2

[358] I say "for the main part", because even most relational accounts will be subject to some troublesome counter-examples. But I don't think this should count for too much, because *all* extant analyses of intrinsicality seem to be subject to at least some counter-intuitive counter-examples. On Lewis's account, for instance, the property of *being married to a round square* will come out as intrinsic, since no two possible duplicates will differ in respect to it.

examples, for instance, the property of *being ten feet from an electron* will, on a relational account, come out as extrinsic, just as it should do. And this is the case even for the monist. Sure, if you're a monist there are no electrons, so nothing could actually instantiate the property of *being ten feet from an electron*, but the monist can still recognise it as an extrinsic property, for the simple reason that if it *were* to be instantiated, whatever did instantiate it would do so in virtue of standing in some relation to some distinct object (specifically, by standing in the *ten feet away from* relation to an electron). More generally, the property of being ten feet away from *anything* will *always* come out as being extrinsic, since its instantiation will always depend on its bearer standing in some relation or other to some distinct object in its surroundings. Moreover, the property of *being negatively charged* will, rightly, come out as intrinsic, because something's having that property does not depend on that something standing in any relation to any distinct objects.

The specific way in which the monist should articulate her relational account of intrinsicality could vary. As I have already noted, there are a number of differing relational accounts already available in the extant literature. But I don't think the monist needs to get too bogged down in the specifics in order to show merely how a generally relational account of intrinsicality can overcome Sider's objection. Different relational accounts of intrinsicality will have differing pros and cons, for sure. Most likely, these pros and cons will be centred around the specific counter-examples that a particular analysis is (or is not) subject to. But this is to be expected in this particular debate. For I do not know of *any* account of intrinsicality (the Lewis-Sider account included) that is not subject to at least *some* counter-intuitive counter-examples. So the monist will simply need to make a reasoned and informed decision as to how she should formalise her relational account of intrinsicality, and take the resultant counter-examples on the chin. The important point is, however, that for the main part at least, any broadly relational account of intrinsicality should enable the monist to generate the required facts about intrinsic and extrinsic properties.

So here is the bottom line. If Sider's claim is that the monist can't make any sense *at all* of statements like 1 and 2 (and I think that it is), then I would suggest that he is just wrong. The monist can make perfect sense of them, at least, insofar as she can make sense of a generally relational account of intrinsicality, which she clearly can. But if, however, his complaint is the much stronger claim that the monist cannot give an adequate and formal

account of intrinsicality, then it is surely unjustified. Sure, it may be difficult for the monist to give a complete and satisfying account of intrinsicality that is entirely free from troublesome counter-examples. But that problem is not unique to the monist! Pluralists have been engaged in a similar struggle to analyse intrinsicality for years, and have yielded only limited success. It is an unfair request, therefore, to demand a complete and error-free account of intrinsicality from the monist, when one cannot provide one oneself. Overall, the monist has no more difficulty in defining intrinsicality than the pluralist. She is merely constrained by the fact that she cannot use Lewis's account in order to do so.

## §8.3. Final Concluding Remarks

I hope to have shown that none of the objections considered represent a serious threat to the monist, since all of them can be overcome. Priority Monism, despite its name, is not really a theory of monism, and should certainly not be seen as a straight alternative to the thesis of monism being considered here. Thus it presents no special threat to monism, at least, no more threat than any other ontological view with which monism is incompatible. Ted Sider's objections certainly constitute more of a serious challenge. But I hope to have shown that they too can be overcome, and at relatively little cost.

And so it is that I must conclude this ontological thesis. I hope that the reader will have been at the very least compelled by my reasoning and that they will concede that monism is a coherent and plausible thesis. Indeed I live in hope that the more convivial reader may have been convinced that monism is true. It certainly seems clear to me, at least, that we live in a monistic world; rich in properties, yet sparse in concrete objects. But I must concede that it is not without something of a heavy heart and a certain degree of resignation that I accept this conclusion. For no matter how compelling I find the reasoning, and no matter how convinced I am in the truth of my conclusions, I simply cannot seem to shake my inherently pluralistic view of the world. When away from the philosophy room, I cannot help but think and talk about the world as though I were a committed pluralist; I cannot help but act as though multiple objects abound. And I would be lying if I were to say that this somewhat hypocritical behaviour, this seemingly irrevocable tension between my

rational beliefs and my instincts, did not cause me any anxiety at all, for it most certainly does. But, of course, I am in good company for suffering from this turmoil. For I feel a certain sympathy for David Hume, who was so famously afflicted with melancholy by the realisation that his philosophical conclusions about causation were in such irremediable conflict with his natural and instinctual beliefs. And I think I have no option but to accept a similarly Humean conclusion; that our belief in a multiplicity of objects (just like our belief in the relation of cause and effect) is an innate and incontrovertible feature of the human mind. We can entertain the idea that it is false, but we can never truly believe it to be so. When going about one's daily life, one can never fail to act as though the ordinary objects of experience really and truly exist. And this is why I say that it is only with a certain degree of resignation that I conclude monism to be true. I am what one may call a *reluctant monist*. But I am a monist all the same.

# §9. Bibliography


Akiba, K. (2004) 'Vagueness in the World', *Noûs,* 38, 3, pp.407-429

Alexander, H. G. (1956) *The Leibniz-Clarke Correspondence,* (Manchester, MUP)

Aristotle (1998), *The Metaphysics,* trans. Lawson-Tancred, H. (London, Penguin)

Armstrong, D. M. (1978) *Universals and Scientific Realism vol. II,* (Cambridge, CUP)

Armstrong, D. M. (1980) 'Identity Through Time', in van Inwagen, P. (ed.) *Time and Cause,* (Dordrecht, D. Reidel), pp.67-78

Armstrong, D. M. (1986) 'In Defence of Structural Universals', *Australasian Journal of Philosophy,* 64, 1, pp.85-88

Armstrong, D. M. (1989) *Universals: an Opinionated Introduction,* (Boulder, Colorado; Westview Press)

Armstrong, D. M. (1997) *A World of States of Affairs,* (Cambridge, CUP)

Armstrong, D. M. (2004) *Truth and Truthmakers,* (Cambridge, CUP)

Arntzenius, F. (2008) 'Gunk, Topology and Measure' in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics vol. 4,* (Oxford, OUP), pp.225-247

Barnes, E. & Williams, J. R. G. (2009) 'Vague Parts and Vague Identity', *Pacific Philosophical Quarterly,* 90, 2, pp.176-187

Barnes, E. & Williams, J. R. G. (2011a) 'A Reply to Eklund' in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics, vol. 6,* (Oxford, OUP)

Barnes, E. & Williams, J. R. G. (2011b) 'A Theory of Metaphysical Indeterminacy' in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics, volume 6,* (Oxford, OUP)

Barnes, E. (2005) 'Vagueness in Sparseness: A Study in Property Ontology', *Analysis,* 65, 4, pp.315-321

Barnes, E. (2009) 'Indeterminacy, Identity and Counterparts', *Synthese,* 168, 1, pp.81-96

Barnes, E. (2010a) 'Ontic Vagueness: A Guide for the Perplexed', *Noûs,* 44, 4, pp.601-627

Barnes, E. (2010b) 'Arguments against Metaphysical Indeterminacy and Vagueness', *Philosophy Compass,* 5, 11, pp.953-964

Barnes, E. (2012) 'Emergence and Fundamentality', *Mind,* 121, 484, pp.873-901

Barnes, J. (2001) *Early Greek Philosophy (2nd Ed.),* (London, Penguin)

Baxter, D. (1988a) 'Identity in the Loose and Popular Sense', *Mind,* 97, 388, pp.575-582

Baxter, D. (1988b) 'Many-One Identity', *Philosophical Papers,* 17, 3, pp.193-216

Beckerman, A., Flohr, H. & Kim, J. (eds.) (1992). *Emergence or Reduction?* (Berlin, de Gruyter)

Bigelow, J. (1996) 'Presentism and Properties', *Philosophical Perspectives,* 10, pp.35-52

Black, M. (1951) 'Achilles and the Tortoise', *Analysis,* 11, 5, pp.91-101

Bohm, D. (1957) *Causality and Chance in Modern Physics,* (London, Routledge and Kegan Paul)

Bohn, E. (2009a) 'An Argument against the Necessity of Universalism', *Analysis,* 69, 1, pp.27-31

Bohn, E. (2009b) 'Must there be a Top Level?' *The Philosophical Quarterly,* 59, 235, pp.193-202

Braddon-Mitchell, D. & Miller, K. (2006) 'The Physics of Extended Simples', *Analysis,* 66, 3, pp.222-226

Bradley, F. H. (1930) *Appearance and Reality,* (Oxford, Clarendon Press)

Bunzl, M. (1979) 'Causal Overdetermination' *The Journal of Philosophy,* 76, 3, pp.134-150

Cameron, R. (2007) 'The Contingency of Composition', *Philosophical Studies,* 136, 1, pp.99-121

Cameron, R. (2008a) 'Truthmakers and Ontological Commitment' *Philosophical Studies,* 140, pp.1-18

Cameron, R. (2008b) 'Turtles all the way Down: Regress, Priority and Fundamentality', *Philosophical Quarterly,* 58, 230, pp.1-14

Cameron, R. (2008c) 'Intrinsic and Extrinsic Properties', in Le Poidevin, R., Simons, P., McGonigal, A., & Cameron, R. (eds.), *The Routledge Companion to Metaphysics,* (London, Routledge), pp.265-275

Cameron, R. (2010a) 'Quantification, Naturalness and Ontology', in Hazlett, A. (ed.) *New Waves in Metaphysics,* (Basingstoke, Palgrave MacMillan), pp.8-27

Cameron, R. (2010b) 'How to Have a Radically Minimal Ontology', *Philosophical Studies,* 151, 2, pp.249-264

Cameron, R. (2012) 'CAI Doesn't Settle the SCQ', *Philosophy and Phenomenological Research,* 84, 3, pp.531-554

Campbell, K. (1981) 'The Metaphysics of Abstract Particulars', *Midwest Studies in Philosophy,* 6, 1, pp.477-488

Campbell, K. (1990) *Abstract Particulars,* (Oxford, Blackwell)

Caplan, B. & Matheson, C. (2006), 'Defending Musical Perdurantism', *British Journal of Aesthetics,* 46, 1, pp.59-69

Cartwright, R. (1975), 'Scattered Objects', in Lehrer, K. (ed.) *Analysis and Metaphysics,* (Boston, Reidel Publishing Company), pp.153-171

Chalmers, D. (2009), 'Ontological Anti-Realism', in Chalmers, D., Manley, D. & Wasserman, R. (eds.) *Metametaphysics,* (Oxford, OUP), pp.77-129

Clayton, P. & Davies, P. (2006) *The Re-emergence of Emergence,* (Oxford, OUP)

Comesana, J. (2008) 'Could There Be Exactly Two Things?', *Synthese,* 162, 1, pp.31-35

Cornell, D. M. (2013) 'Monism and Statespace: A Reply to Sider', *Analysis,* 73, 2, pp.230-236

Cowling, S. (forthcoming) 'Ideological Parsimony', forthcoming in *Synthese*

Dainton, B. (2010) *Time and Space (2nd Ed.),* (Durham, Acumen)

Daly, C. (1994) 'Tropes', *Proceedings of the Aristotelian Society,* 94, pp.253-261

Descartes, R. (1934) *A Discourse on Method,* (London; J. M. Dent & Sons Ltd.)

Descartes, R. (1986) *Meditations on First Philosophy,* ed. Cottingham, J. (Cambridge, CUP)

Doepke, F. (1982) 'Spatially Coinciding Objects', *Ratio,* 24, pp.45-60

Doepke, F. (1996) *The Kinds of Things,* (La Salle, IL, Open Court Press)

Dorr, C. (2002) *The Simplicity of Everything,* PhD Thesis, available at <http://users.ox.ac.uk/~sfop0257/papers/SimplicityOfEverything.pdf>

Dorr, C. (2005) 'What We Disagree About When We Disagree About Ontology' in Kalderon, M. (ed.) *Fictionalist Approaches to Metaphysics,* (Oxford, OUP), pp.234-286

Dummett, M. (1975) *Truth and Other Enigmas,* (London, Duckworth)

Dunn, J. M. (1990) 'Relevant Predication 2: Intrinsic Properties and Internal Relations', *Philosophical Studies,* 60, pp.177-206

Eddon, M. (2011) 'Intrinsicality and Hyperintensionality', *Philosophy and Phenomenological Research,* 82, pp.314-336

Effingham, N. (2007) *The Restricted Composition of Material Objects,* PhD. Thesis. Available at http://www.nikkeffingham.com/publications.html

Field, H. (1980) *Science without Numbers,* (Princeton, Princeton UP)

Fine, K. (1995) 'Ontological Dependence', *Proceedings of the Aristotelian Society,* 95, pp.269-290

Fine, K. (2001) 'The Question of Realism', *Philosophers Imprint,* 1, 2, pp.1-30

Forrest, P. (1995) 'Is Space-Time Discrete or Continuous? An Empirical Question', *Synthese,* 103, pp.327-54

Forrest, P. (2004) 'Grit or Gunk?' *The Monist,* 87, 3, pp.351-370

Francescotti, R. (1999) 'How to Define Intrinsic Properties', *Nous,* 33, pp.590-609

Frege, G. (1892) 'Uber Sinn und Bedeutung', *Zeitschrift fur Philosophie und Philosophische Kritik*, C: 25–50. English Translation: "On Sense and Meaning", in McGuinness, B. (ed.), *Frege: Collected Works*, (Oxford, Basil Blackwell), pp.157–177

French, S. & Krausse, D. (2006) *Identity in Physics: A Historical, Philosophical and Formal Analysis,* (Oxford, Clarendon)

Garwood, C. (2007) *Flat Earth: The History of an Infamous Idea,* (London, St. Martin's Press)

Geach, P. (1962) *Reference and Generality,* (Ithaca, Cornell UP)

Geach, P. (1967) 'Identity', *Review of Metaphysics,* 21, pp.3-12

Geach, P. T. (1980) *Reference and Generality, 3rd Edition,* (Ithaca, Cornell UP)

Grice, P. (1989) *Studies in the Way of Words,* (Cambridge, Mass., Harvard UP)

Goff, P. (2012) 'There is More than One Thing', in Goff, P. (ed.) *Spinoza on Monism,* (Basingstoke, Palgrave MacMillan), pp. 113-122

Haslanger, S. (1989) 'Endurance and Temporary Intrinsics', *Analysis,* 49, 3, pp.119-125

Hawley, K. (2001) *How Things Persist,* (Oxford, OUP)

Hawthorne, J. & Cortens, A. (1995) 'Towards Ontological Nihilism', *Philosophical Studies,* 79, pp.143-165

Heil, J. (2009) 'Relations' in Cameron, R., Le Poidevin, R., McGonigal, A., & Simons, P. (eds.) *The Routledge Companion to Metaphysics,* (London, Routledge)

Heil, J. (2012) 'Substances Stressed' in Goff, P. (ed.) *Spinoza on Monism,* (Basingstoke, Palgrave MacMillan)

Heller, M. (1984) 'Temporal Parts of Four-Dimensional Objects', *Philosophical Studies,* 46, pp.323-334

Heller, M. (1990) *The Ontology of Physical Objects: Four-Dimensional Hunks of Matter,* (Cambridge; CUP)

Heller, M. (1996) 'Against Metaphysical Vagueness', *Philosophical Perspectives,* 10, pp.177-185

Hinchliff, M. (2000) 'A Defence of Presentism in a Relativistic Setting', *Philosophy of Science,* 67, 3, pp.575-586

Hirsch, E. (2002), 'Quantifier Variance and Realism', *Nous,* 36, 1, pp.51-73

Hirsch, E. (2005), 'Physical-Object Ontology, Verbal Disputes, and Common Sense', *Philosophy and Phenomenological Research,* 70, 1, pp.67-97

Hirsch, E. (2008), 'Language, Ontology, and Structure', *Noûs,* 42, 3, pp.509-528

Hirsch, E. (2009), 'Ontology and Different Languages', in Chalmers, D., Manley, D. & Wasserman, R. (eds.) *Metametaphysics,* (Oxford, OUP), pp.231-259

Hoffmann-Kolss, V. (2010) *The Metaphysics of Extrinsic Properties,* (Frankfurt, Ontos-Verlag)

Horgan, T. (1993) 'On What There Isn't', *Philosophy and Phenomenological Research,* 53, 3, pp.693-700

Horgan, T (1994) 'Robust Vagueness and the Forced-Marched Sorites Paradox', *Philosophical Perspectives,* 8, pp.159-188

Horgan, T. & Potrč, M. (2000) 'Blobjectivism and Indirect Correspondence', *Facta Philosophica,* 2, pp.249-270

Horgan, T. & Potrč, M. (2008) *Austere Realism,* (London, MIT Press)

Hudson, H. (2000) 'Universalism, Four-Dimensionalism, and Vagueness', *Philosophy and Phenomenological Research,* 60, 3, pp.547-560

Hudson, H. (2001) *A Materialist Metaphysics of the Human Person,* (London, Cornell UP)

Hudson, H. (2006) *The Metaphysics of Hyperspace,* (Oxford, OUP)

Hudson, H. (2007) 'Simples and Gunk', *Philosophy Compass,* 2, 2, pp.291-302

Johnston, M. (1987) 'Is There A Problem About Persistence?', *Proceedings of the Aristotelian Society,* vol. 61, pp.107-35

Jubien, M. (1993) *Ontology, Modality and the Fallacy of Reference,* (Cambridge, CUP)

Kim, J. (1993) *Supervenience and Mind,* (Cambridge, CUP)

Korman, D. (2007) 'Unrestricted Composition and Restricted Quantification', *Philosophical Studies,* 140, 3, pp.319-344

Kripke, S. (1980) *Naming and Necessity,* (Oxford, Basil Blackwell)

Le Poidevin, R. (1991) *Change, Cause and Contradiction,* (Basingstoke, MacMillan)

Leonard, H. S. & Goodman, N. (1940) 'The Calculus of Individuals and its uses', *Journal of Symbolic Logic,* 5, pp.45-55

Lesniewski, S. (1916), 'Foundations of the General Theory of Sets' in Lesniewski, S. *Collected Works,* eds. S. J. Surma, J. Srzednicki, D. I. Barnett, and F. V. Rickey, *trans.* D. I. Barnett, (Dordrecht, Kluwer, 1992) Vol. 1pp.129-173

Lewis, D. (1973) *Counterfactuals,* (Oxford, Blackwell)

Lewis, D. (1986a) *On the Plurality of Worlds,* (Oxford, Basil Blackwell)

Lewis, D. (1986b) 'Against Structured Universals', *Australasian Journal of Philosophy,* 64, 1, pp.25-46

Lewis, D. (1991) *Parts of Classes,* (Oxford, Basil Blackwell)

Loeb, L. E. (1974) 'Causal Theories and Causal Overdetermination', *The Journal of Philosophy,* 71, 15, pp.525-544

Markosian, N. (1998a) 'Simples', *Australasian Journal of Philosophy,* 76, 2, pp.213-228

Markosian, N. (1998b), 'Brutal Composition', *Philosophical Studies,* 92, 3, pp.211-249

Markosian, N. (2004) 'A Defence of Presentism', in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics, vol. 1,* (Oxford, OUP), pp.47-82

Markosian, N. (2008), 'Restricted Composition', in Hawthorne, J., Sider, T. & Zimmerman, D. (eds.) *Contemporary Debates in Metaphysics,* (Malden, Blackwell), pp.341-364

McDaniel, K. (2003) 'Against Max-Con Simples', *Australasian Journal of Philosophy,* 81, 2, pp.265-275

McDaniel, K. (2007) 'Extended Simples', *Philosophical Studies,* 133, 1, pp.131-141

McDaniel, K. (2009) 'Extended Simples and Qualitative Heterogeneity', *Philosophical Quarterly,* 59, 235, pp.325-331

McGrath, M. (1998) 'van Inwagen's Critique of Universalism', *Analysis,* 58, 2, pp.116-121

McKay, T. (2006) *Plural Predication,* (Oxford, OUP)

McKenzie, K. (2011) 'Arguing against Fundamentality', *Studies in Philosophy and History of Science part B,* 42, 4, pp.244-255

Merricks, T. (2001) *Objects and Persons,* (Oxford, OUP)

Merricks, T. (2005) 'Composition and Vagueness', *Mind,* 114, 455, pp.615-637

Moore, G. E., Stout, G. F. & Dawes Hicks, G. (1923) 'Are the Characteristics of Particular Things Universal or Particular?', *Proceedings of the Aristotelian Society, Supplementary Volumes,* 3, pp.95-128

Morganti, M. (2009) 'Ontological Priority, Fundamentality and Monism', *Dialectica,* 63, 3, pp.271-288

Mulligan, K. (1998) 'Relations, Through Thick and Thin', *Erkenntnis,* 48, 325-353

Mulligan, K., Simons, P. & Smith, B. (1984) 'Truth-Maker', *Philosophy and Phenomenological Research,* 44, pp.287-321

Nerlich, G. (1994) *The Shape of Space,* (Cambridge, CUP)

Nolan, D. (1997) 'Quantitative Parsimony' *The British Journal for the Philosophy of Science,* 48, 3, pp.329-343

Nolan, D. (2005) *David Lewis,* …

Parsons, J. (2000) 'Must a Four-Dimensionalist Believe in Temporal Parts?', *The Monist,* 83, 3, pp.399-418

Parsons, J. (2004) 'Distributional Properties' in Jackson, F. & Priest, G. (eds.) *Lewisian Themes: The Philosophy of David K. Lewis,* (Oxford, OUP), 173-180

Parsons, T. & Woodruff, P. (1995) 'Worldly Indeterminacy of Identity', *Proceedings of the Aristotelian Society,* 95, pp.171-191

Plato (1892) *Parmenides,* in Plato *The Dialogues of Plato,* trans. Jowett, B. (Oxford, Clarendon Press), Vol. IV, pp.1-107

Quine, W. V. (1960) *Word and Object,* (Cambridge, Mass., MIT Press)

Ramsey, F. P. (1925) 'Universals', *Mind,* 34, 136, pp.401-417

Rea, M. (1998) 'In Defence of Mereological Universalism', *Philosophy and Phenomenological Research,* 58, 2, pp.347-360

Rea, M. (2001) 'How To Be An Eleatic Monist', *Philosophical Perspectives,* 15, pp.129-152

Rescher, N. (1955), 'Axioms for the Part Relation' *Philosophical Studies,* 6, 1, pp.8-11

Rosen, G. & Smith, N. (2004) 'Worldly Indeterminacy: A Rough Guide', *Australasian Journal of Philosophy,* 82, 1, pp.185-198

Rosenberg, J. F. (1993) 'Comments on Peter van Inwagen's *Material Beings*', *Philosophy and Phenomenological Research,* 53, 3, pp.701-708

Russell, B. (1905) 'On Denoting', *Mind,* 14, 56, pp.479-493

Russell, B. (1910) 'Some Explanations in Reply to Mr. Bradley', *Mind,* 19, 75, pp.373-378

Russell, B. (1923) 'Vagueness' *Australasian Journal of Philosophy,* 1, 2, pp.84-92

Russell, B. (1997) 'Is There A God?' in Slater, J.G. & Kollner, P. (ed.s) *The Collected Papers of Bertrand Russell, vol.11: Last Philosophical Testament,* (London, Routledge)

Russell, J. B. (1991) *Inventing the Flat Earth,* (London, Greenwood)

Ryle, G. (1949) *The Concept of Mind,* (London, Hutchinson)

Sanford, D. (1993) 'The Problem of the Many, Many Composition Questions, and Naive Mereology', *Nous,* 27, 2, pp.219-228

Saucedo, R. (2012) 'Parthood and Location' in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics, vol.6,* (Oxford, OUP)

Schaffer, J. (2003) 'Is There a Fundamental Level?', *Nous,* 37, 3, pp.498-517

Schaffer, J. (2007) 'From Nihilism to Monism' *Australasian Journal of Philosophy,* 85, 2, pp.175-191

Schaffer, J. (2009), 'On What Grounds What', in Chalmers, D., Manley, D. & Wasserman, R. (eds.) *Metametaphysics,* (Oxford, OUP), pp.347-383

Schaffer, J. (2010a) 'The Least Discerning and Most Promiscuous Truthmaker', *Philosophical Quarterly,* 60, 239, pp.307-324

Schaffer, J. (2010b) 'The Internal Relatedness of All Things', *Mind,* 119, 474, pp.341-376

Schaffer, J. (2010c) 'Monism: The Priority of the Whole', *Philosophical Review,* 119, 1, pp.31-76

Shoemaker, S. (1999) 'Self, Body and Coincidence', *Proceedings of the Aristotelian Society,* Supp. Vol. 73, pp.287-306

Sider, T. (1993a) 'Van Inwagen and the Possibility of Gunk', *Analysis,* 53, pp.285-289

Sider, T. (1993b) 'Intrinsic Properties', *Philosophical Studies,* 83, 1, pp.1-27

Sider, T. (2001) *Four-Dimensionalism,* (Oxford; OUP)

Sider, T. (2003) 'What's So Bad about Causal Overdetermination?' *Philosophy and Phenomenological Research,* 67, pp.719-726

Sider, T. (2007a) 'Against Monism', *Analysis,* 67, 1, pp.1-7

Sider, T. (2007b) 'Parthood', *Philosophical Review,* 116, pp.51-91

Sider, T. (2008) 'Monism and Statespace Structure', in Le Poidevin, R. (ed.) *Being: Developments in Contemporary Metaphysics,* (Cambridge, CUP), pp.129-150

Sider, T. (2009) 'Ontological Realism', in Chalmers, D., Manley, D. & Wassermann, R. (eds.) *Metametaphysics,* (Oxford, OUP)

Sider, T. (forthcoming) 'Against Parthood' forthcoming in Bennett, K. & Zimmerman, D. (eds.) *Oxford Studies in Metaphysics,* vol. 8 (Oxford, OUP, 2013). Available at: http://tedsider.org/papers/nihilism.pdf

Silva, P. (2012) 'Ordinary Objects and Series-Style Answers to the Special Composition Question', *Pacific Philosophical Quarterly,* 93, 4, pp.69-88

Simons, P. (1987) *Parts: A Study in Ontology,* (Oxford, Clarendon)

Simons, P. (2004) 'Extended Simples: A Third Way between Atoms and Atomless Gunk', *The Monist,* 87, 3, pp.371-385

Simons, P. (2006) 'Real Wholes, Real Parts', *The Journal of Philosophy,* 103, 12, pp.597-613

Sorensen, R. (2003) *A Brief History of the Paradox,* (Oxford; OUP)

Strawson, P. (1950) 'On Referring', *Mind,* 59, 235, pp.320-344

Thomasson, A. (2006) 'Metaphysical Arguments against Ordinary Objects', *The Philosophical Quarterly,* 56, 224, pp.340-360

Thomasson, A. (2009), 'Answerable and Unanswerable Questions', in Chalmers, D., Manley, D. & Wasserman, R. (eds.) *Metametaphysics,* (Oxford, OUP), pp.444-471

Thompson, J.J. (1998) 'The Statue and the Clay', *Nous,* 32, pp.149-173

Turner, J. (2011) 'Ontological Nihilism' in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics, vol. 6,* (Oxford, OUP), pp.1-50

Unger, P. (1979a) 'There Are No Ordinary Things' *Synthese,* 41, pp.117-154

Unger, P. (1979b) 'I Do Not Exist', in MacDonald, G. F. (ed.) *Perception and Identity: Essays Presented to A. J. Ayer with his Replies to them,* (London, MacMillan)

Unger, P. (1979c) 'Why There Are No People', *Midwest Studies in Philosophy,* 4, 1, pp.177-222

Unger, P. (1980a) 'Skepticism and Nihilism', *Noûs,* 14, 4, pp.517-545

Unger, P. (1980b) 'The Problem of the Many', *Midwest Studies in Philosophy,* 5, 1, pp.411-468
Uzquiano, G. (2004) 'Plurals and Simples', *The Monist,* 87, 3, pp.429-451

Vallicella, W. (2002) 'Relations, Monism and the Vindication of Bradley's Regress', *Dialectica,* 56, 1, pp.3-35

Van Bendegem (1987) 'Zeno's Paradoxes and the Tile Argument', *Philosophy of Science,* 54, 2, pp.295-302

Van Cleve, J. (1986) 'Mereological Essentialism, Mereological Conjunctivism, and Identity Through Time', *Midwest Studies in Philosophy,* 11, 1, pp.141-156

Van Cleve, J. (2008) 'The Moon and Sixpence: A Defence of Mereological Universalism', in Hawthorne, J., Sider, T. & Zimmerman, D. (ed.s) *Contemporary Debates in Metaphysics,* (Oxford, Blackwell), pp.321-340

Van Inwagen, P. (1987) 'When are Objects Parts?' *Philosophical Perspectives,* 1, pp.21-47

Van Inwagen, P. (1990), *Material Beings,* (Ithaca, Cornell UP)

Van Inwagen, P. (1994) 'Composition as Identity', *Philosophical Perspectives,* 8, pp.207-220

Van Inwagen, P. (2002) *Metaphysics (2nd ed.),* (Oxford, Westview Press)

Varzi, A. (2000) 'Mereological Commitments', *Dialectica,* 54, pp.283-305

Wachter, D. (1998) 'On Doing Without Relations', *Erkenntnis*, 48, pp.355-361

Wasserman, R. (2004) 'The Constitution Question', *Nous*, 38, 4, pp.693-710

Wasserman, R. (2009) 'Material Constitution' in *The Stanford Encyclopedia of Philosophy*. Available at <http://plato.stanford.edu/entries/material-constitution>.

Weyl, H. (1949) *Philosophy of Mathematics and Natural Science*, (Princeton, Princeton UP)

Wheeler, S. (1979) 'On That Which Is Not', *Synthese*, 41, 2, pp.155-173

Whitehead, A. N. (1929) *Process and Reality: An Essay in Cosmology*, (New York, MacMillan)

Wiggins, D. (1968a) 'On Being in the Same Place at the Same Time', *Philosophical Review*, 77, pp.90-95

Wiggins, D. (1968b) *Identity and Spatio-Temporal Continuity*, (Oxford, Blackwell)

Williams, J. R. G. (2006) 'Illusions of Gunk', *Philosophical Perspectives*, 20, 1, pp.493-513

Winston, Chaffin & Herrman (1987) 'A Taxonomy of Part-Whole Relations', *Cognitive Science*, 11, pp.417-444

Wollheim, R. (1959) *F. H. Bradley*, (Harmondsworth, Pelican)

Yablo, S. (2006) 'Non-Catastrophic Presupposition Failure', in Bryne, A. & Jarvis Thomson, J. (eds.) *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, (Oxford, OUP), pp.164-191

Yablo, S. (2009) 'Must Existence-Questions Have Answers?' in Chalmers, D., Manley, D. & Wasserman, R. (eds.) *Metametaphysics*, (Oxford, OUP), pp.507-527

Zimmerman, D. (1996a) 'Could Extended Objects be Made out of Simple Parts? An Argument for Atomless Gunk', *Philosophy and Phenomenological Research*, 56, 1, pp.1-29

Zimmerman, D. (1996b) 'Indivisible Parts and Extended Objects', *The Monist*, 79, 1, pp.148-180

Zimmerman, D. (2008) 'The Privileged Present' in Hawthorne, J., Sider, T. & Zimmerman, D. (eds.) *Contemporary Debates in Metaphysics*, (Oxford, OUP)