

A DEFENCE OF REDUCTIVE NIHILISM

Simon Thunder

PhD Thesis, University of Nottingham

Acknowledgements

Lots of people have given me help and support of various kinds as I researched and wrote this thesis, and still more people helped me to become the sort of person who could do a PhD in the first place. Here I'll just acknowledge a few.

Thanks to my PhD supervisors, Jonathan Tallant and Nikk Effingham. I couldn't really have asked for a better supervisory team.

Thanks to Dave Ingram for his mentorship and friendship.

Thanks to my friends in the PhD community, both at Nottingham and elsewhere. Particular thanks to Lucía Ontiveros Llamas and Marcus Lee.

Thanks to the Midlands3Cities DTP, and by extension, the AHRC, for funding both my MA and the first three years of my doctoral studies; thanks to the Mind Association for funding my write-up year.

And thanks to my parents, for all of their love and support.

Contents

1. <i>Introduction</i> _____	p. 4
2. <i>Introducing Reductive Nihilism</i> _____	p. 8
3. <i>Reductive Nihilism and the Ramsey-Lewis Treatment of Theoretical Terms</i> _____	p. 35
4. <i>Reductive Nihilism and the Causal Theory of Reference</i> _____	p. 53
5. <i>Defending the Semantics</i> _____	p. 77
6. <i>An Intuitive Answer to the Special Composition Question</i> _____	p. 89
7. <i>A Parsimonious Answer to the Special Composition Question I: In Defence of the Razor</i> ____	p. 116
8. <i>A Parsimonious Answer to the Special Composition Question II: Keeping it Simple</i> _____	p. 131
9. <i>Conclusion</i> _____	p. 144

Chapter 1

Introduction

In 1990, Peter van Inwagen asked a question. He called his question the *Special Composition Question* (SCQ) (van Inwagen 1990a: 30-1). Here's an intuitive gloss (we'll encounter a more precise formulation in §2.3):

SCQ: Under what conditions do some smaller things compose a single larger thing?

Composition is just the relation that some parts collectively bear to the whole of which they are parts. It is what happens when some smaller things *make up* a single bigger thing, or when there is a bigger thing *made out of* some smaller things. So the SCQ is asking us for the conditions under which there is some bigger thing made out of some smaller things.

The reason (or rather, one reason) the SCQ is interesting is because, intuitively, there are some cases in which some smaller things do compose a single larger thing, and some cases in which some smaller things do not compose a single larger thing. That is, intuitively, there are some cases in which there is some larger thing made out of some smaller things, and some cases in which there is no larger thing made out of some smaller things.

Examples of cases where composition intuitively does occur are easy to think of. Consider any ordinary material object. Maybe it's a table, or a mug, or a laptop, or a microwave, or the Eiffel Tower. Chances are, the material object you're thinking of is made up of parts. The table on which I'm leaning, for example, is made up of four table-legs and a table-top (plus maybe some nails or screws, etc.); the Eiffel Tower is made up of a certain number of iron struts, among other things; etc. Zooming in a little, almost all material objects are made up of atoms; zooming in a little more, almost all material objects are made up of sub-atomic particles. Each of these cases seems to be a case in which some smaller things compose a single larger thing. For example, the four table-legs and the table-top seem to compose my table; the iron struts, together with other building materials, seem to compose the Eiffel Tower; various pluralities of atoms and/or sub-atomic particles seem to compose various larger objects; etc.

Examples of cases where composition intuitively doesn't occur are also easy to think of. Pick at random some objects that seem to have nothing at all to do with each other. Maybe it's the table you're leaning on and the oldest curtain in Buckingham Place; maybe it's a sea cucumber currently residing deep in the Mariana Trench and a particular paperclip located somewhere in Mozambique; maybe it's your right index finger, a particular atom at the centre of Venus, and the fifth page of your copy of *On the Plurality of Worlds*. You get the idea. Intuitively, there is nothing composed out of any of these pluralities of things. There is no thing that has as its parts a sea cucumber in the Atlantic Ocean and a Mozambiquean paperclip; nothing is made out of the table you're leaning on and a royal curtain; etc.

We'll encounter some more examples of cases in which composition intuitively does and doesn't occur in §6.3. For now, the point is just that composition does seem to occur in some cases, and doesn't seem to occur in others, and that that makes the question of specifying the conditions under which composition does and doesn't occur at least mildly interesting.

Here's another reason for thinking that the SCQ is interesting: the obvious, reflex candidate answers to it seem on reflection to be inadequate. For example, we might think that some smaller things

compose a single larger thing iff those smaller things are in *contact* with one another. We might then clarify that the *being in contact* relation at issue here is to be understood as a two-place, reflexive, symmetrical, transitive relation, and that some things can then be said to be in contact with each another iff each of those things bears the *being in contact* relation to every other one of those things. Thus the reason why the four table-legs and the table-top compose something is because they are all in contact with one another (the transitivity of *being in contact* ensuring that each of the four table-legs are in contact with one another in virtue of their all being in contact with the table-top); the reason why the sea cucumber and the paperclip do not compose anything is because they are not in contact with one another. But, as van Inwagen (1990a: ch. 3) demonstrates, this answer to the SCQ does not on reflection seem adequate, because 20th century physics apparently reveals that the table-legs and the table-top are not really in contact with one another at all, and in fact that *no* two objects are ever really in contact. If a necessary condition for composition were for some things to be in contact, then composition would never occur. Even if we revise our understanding of the relation of *being in contact* such that it matches the folk understanding of when some objects are in contact (assuming that the folk notion can be made sense of adequately), we could still cook up counterexamples fairly easily. For instance, suppose you and I were to shake hands. Then we would be in contact in the folk sense of that phrase, but it doesn't seem as though we would compose anything (van Inwagen 1990a: 35).

To other similarly prima facie plausible suggestions about when composition occurs (e.g. the suggestion that composition occurs when some things are sufficiently *well-bonded*), van Inwagen points out similar problems (1990a: 56-71). We appear to have a genuine philosophical puzzle on our hands.

Let us turn now away from obvious, reflexive candidate answers the SCQ, and consider instead the efforts of philosophers. My aim here is not to evaluate these efforts (for that, see chapters 6 and 8), nor to present every minor technicality or caveat associated with them (I will highlight such technicalities and caveats if and when they become relevant later on). Rather, my aim here is merely to introduce the most popular and/or promising answers to the SCQ proposed by philosophers.

I also note at this stage that I won't consider at all here *deflationist* views on the SCQ (see especially Hirsch 2002, 2009; Thomasson 2007, 2015). Deflationist views say that the SCQ is not a substantive question about the nature of reality, and that although competing answers to it appear to be genuinely disagreeing about the contents and nature of reality, this disagreement is really a merely verbal dispute. On this view, different answers to the SCQ only differ in the language in which they choose to *describe* reality. One of the assumptions of my project here is the falsity of deflationism. I'll sometimes refer to this assumption as that of *ontological realism*.

The most popular answer to the SCQ, then, is *universalism*. Universalism is the view that in *every* case in which there are some things, those things compose something. On universalism, not only is there something composed out of the four table-legs and the table-top, but there is also something composed out of your right index finger, a particular atom at the centre of Venus, and the fifth page of your copy of *On the Plurality of Worlds*. Universalism has been endorsed and defended by, among others, Lewis (1986: 211-3), Sider (2001), Rea (1998), and van Cleve (2008).

Various forms of *restrictivism* deny that every plurality of things compose, claiming instead that composition occurs in some cases but not in others. *Organicism* (aka *VIPA*, i.e. van Inwagen's Proposed Answer) is the view that some things compose iff their activities constitute a life (van Inwagen 1990a: ch. 9). *Merricks' answer* says that composition occurs iff the resultant composite

object is a conscious being (Merricks 2001). *Carmichael's answer* says that some things compose iff *either* they are lump-like and well-bonded *or* their activities constitute an event that imposes sufficient unity upon them (Carmichael 2015). *NIA* (or the Normal Intuiter Account) says that some smaller things compose a larger thing iff the smaller things are disposed to elicit in normal intuiters the intuition that the larger thing exists (Kriegel 2008).

Also relevant here is the view called *Brutal Composition* (Markosian 1998). Roughly, Brutal Composition is the view that there is no general principle governing when composition occurs and when it doesn't, and that composition either does or doesn't occur in specific cases as a matter of brute fact. This doesn't commit Brutal Composition to any specific account of which pluralities of things compose and which don't, but one might find attractive the version of the view that says things like: as a matter of brute fact, the table-legs and the table-top compose something (a table), but the sea cucumber and the paperclip do not.

Similarly worthy of note here is the view called *Composition as Identity*, or CAI for short (e.g. Wallace 2011a). Roughly, CAI is the view that the relationship between some parts the whole they compose is that of identity. It is perhaps tempting to think that CAI implies universalism, but there seems to be good reason for thinking that CAI is instead combinable with any answer to the SCQ (Cameron 2012, see also McDaniel 2010). The relevance of CAI will become clear later on.

Finally, there are various forms of *mereological nihilism*, or just nihilism for short. Nihilism is the view that composition never occurs. As we'll see (§2.3), it comes in various flavours, but the most familiar version, microphysical nihilism, says that the only material objects that exist are *simples*, i.e. partless particles such as quarks and leptons (Rosen and Dorr 2002; Sider 2013a). These simples never compose any larger objects. Where you might have thought that there was a table, or a microwave, or the Eiffel Tower, there are really just some simples that are (so to speak) arranged table-wise, or microwave-wise, or Eiffel-Tower-wise. In each case, the simples arranged *F*-wise jointly look, feel, behave, etc. exactly as the *F* that might appear to be composed out of them is supposed to look, feel, behave etc. (so it is at least not obvious from observation that microphysical nihilism is false).

The view that I will defend throughout this thesis is a form of microphysical nihilism. I call it *reductive nihilism*. Roughly speaking, reductive nihilism starts with the microphysical nihilist metaphysical picture just outlined, but claims that that picture suffices for the existence of composite objects. For reductive nihilism combines a microphysical nihilist metaphysic with a non-standard semantics according to which terms that purport to refer to composite *F*s instead refer plurally to pluralities of simples arranged *F*-wise. Given this semantics, reductive nihilists can hold on to their nihilism whilst accepting the existence of e.g. the Eiffel Tower: for them, 'the Eiffel Tower' refers to the simples arranged Eiffel-Tower-wise that they (*qua* nihilists) already posit, and because they take the simples arranged Eiffel-Tower-wise to exist, for them the Eiffel Tower exists too. In the same way, reductive nihilists accept the existence of tables, laptops, mugs, etc. because, given their semantics, accepting the existence of these things requires nothing more than accepting the existence of pluralities of simples arranged table-wise, laptop-wise, mug-wise, etc.

There's more to say about the semantics of reductive nihilism, but for now I just want to acknowledge that reductive nihilism does not represent the first attempt to try to combine the metaphysic of microphysical nihilism with a non-standard semantics of this sort. In particular, reductive nihilism is partly inspired by two previous such attempts: that of Liggins (2008), and that of Contessa (2014). When I come to explicating the reductive nihilist's semantics in greater detail (§2.4.2), I'll point out

the ways in which those semantics borrow from or build on Liggins' and/or Contessa's. That said, reductive nihilism's semantics differ in certain attractive ways from Liggins' and Contessa's, and cover bits of vocabulary that those accounts do not. For example, neither Liggins nor Contessa (nor any other nihilist account, to my knowledge) provides a semantic account of 'composition' that allows the nihilist to accept that composition occurs – the reductive nihilist does offer such an account (§2.4.3.3).¹

I don't expect anyone, be friend or foe of nihilism, to simply accept the semantics that I develop for reductive nihilism. Semantic claims, like any other claims that a theory makes, should be *argued* for. Consequently, three chapters of this thesis will be devoted to arguing that if the microphysical nihilist is right about the metaphysics, then the reductive nihilist's semantics follow. The final three chapters (conclusion aside) will then be devoted to arguing for reductive nihilism as a whole.

With that in mind, then, here, then, is the structure of the thesis to come.

In chapter 2 I'll introduce reductive nihilism, precisely differentiating it from rival answers to the SCQ and carefully formulating its semantic claims. In chapters 3 and 4 I'll argue that, if microphysical nihilism is true, then the reductive nihilist's semantics follow. I'll do so by appeal to established and influential accounts of meaning and reference. So in chapter 3 I'll argue that the reductive nihilist's semantics can be argued for in this way by appeal to the Ramsey-Lewis treatment of theoretical terms; in chapter 4 I'll argue instead by appeal to the causal theory of reference. In chapter 5 I'll consider four objections to the semantics of reductive nihilism, and argue that none of them succeed.

At this point I'll drop the assumption of nihilism, and shift to arguing for reductive nihilism as the right answer to the SCQ. In chapter 6 I'll argue that reductive nihilism is the most *intuitive* answer to the SCQ. Then in chapters 7 and 8 I'll present an extended argument for the claim that reductive nihilism is the (joint) most *ontologically parsimonious* answer to the SCQ. Chapter 7's contribution to this argument will be to argue for a particular account of what it is for a theory to be ontologically parsimonious in the (or at least, a) way that matters for theory choice; chapter 8 will then apply that account of ontological parsimony to reductive nihilism and its rivals.

Chapter 9 then summarises the two main conclusions of this thesis. The first is that, if microphysical nihilism is true, then the reductive nihilist's non-standard semantics follow. This has consequences for nihilists and non-nihilists alike: (microphysical) nihilists should be reductive nihilists; critics of (microphysical) nihilism should understand the view to be reductive nihilism, and tailor their criticisms accordingly. The second conclusion, offered only slightly more tentatively, is simply that since reductive nihilism is the (joint) most ontologically parsimonious and the most intuitive answer to the SCQ, we should all be reductive nihilists.

¹ How can the reductive nihilist both deny that composition occurs (in line with her nihilism) and accept that composition occurs (in line with her semantic account)? By differentiating two possible definitions of 'composition', and claiming that the sense of 'composition' according to which she, *qua* nihilist, denies that composition occurs is different from the sense of 'composition' according to which she, *qua* reductive nihilist, accepts that composition occurs. See chapter 2.

Chapter 2

Introducing Reductive Nihilism

2.1 – Introduction

As we saw in chapter 1, reductive nihilism combines the metaphysical picture of a certain kind of mereological nihilism with a non-standard semantics for English. This chapter focuses on elucidating the nature of this combination.

I'll proceed as follows. First (§2.2), I'll introduce some background concepts that are essential to understanding reductive nihilism: in particular, I'll present the crucial notions of a *single entity* and of *plural quantification*. Second (§2.3), I'll introduce some mereological concepts, and use them to characterise the metaphysical claims of various answers to the SCQ, including those of different kinds of mereological nihilism. Then (§2.4) I'll be in a position to present the semantic claims of reductive nihilism. §2.5 summarises and concludes.

2.2 – Background Concepts

I begin by presenting some concepts that will be crucial in elucidating reductive nihilism.

2.2.1 – Entity

Perhaps the most important such concept is that of a (*single*) *entity*. I'll present two ways to try to understand this notion.

First, to be a (single) entity is to be *one in number simpliciter*. What it is to be one in number *simpliciter* contrasts with what it is to be one *F*: i.e. what it is to be one in number relative to a sortal term *F*. When counting relative to a sortal term *F*, we are trying to answer the question of *how many Fs are there?*² Our answer to that question may vary depending on what *F* is. For example, if we were to be asked how many shoes there are in the cupboard, we might truly answer 'six'; if we were instead asked how many *pairs* of shoes there are in the cupboard, we might truly answer 'three'. These answers differ from one another apparently because in the first case we were counting relative to the sortal 'shoe', whereas in the second case we were counting relative to the sortal 'pair of shoes'. In general, to be one *F* is just to be counted as one relative to a sortal term *F*. It does seem, though, that it's legitimate and coherent to ask whether a single pair of shoes, for example, is *really* one in number or not. We might suspect that what we sometimes conceptualise as being a single pair of shoes is, *really, strictly speaking*, not one but *two things*. One way to make sense of this thought is to treat 'thing' as just another sortal that we can count relative to. But there's some plausibility to the thought that the function of 'really' and 'strictly speaking' here is to signal a kind of counting that does *not* depend on the sortals we use to conceptualise reality. Call this kind of counting *counting simpliciter*. The idea here is that the world has an objective, mind-independent, sortal-independent cardinality, which we attempt to report when we make claims about how many things there are. So we might think that there's a fact about whether the pair of shoes, for example, is really one in number or two in number (or 745 in number, or...) that obtains independently of our conceptual scheme. When we

² The notion of counting relative to a sortal traces back to Frege (1953: §46). Subsequent discussions of sortals have emphasised that one of the defining characteristics of a sortal is that they give a criterion for counting what they apply to (see e.g. Strawson [1959: 168] and Grandy [2016], though see e.g. Wiggins [2001: 75] for a difficulty with this idea).

attempt to report that fact we are counting *simpliciter*. To be one in number *simpliciter* is to be one in number according to such a count. Thus to be an *entity* is to be one in number according to such a count.

A second (though related) way to characterise what it is to be an entity is to say that to be an entity is to be apt for being the value of a singular variable of first-order logic (written ‘*x*’ or ‘*y*’, etc.), i.e. the sort of variable that can be bound by the singular existential quantifier (written ‘ $\exists x$ ’, or ‘ $\exists y$ ’, etc.). Equivalently: list all of the true claims that can be made using the singular existential quantifier; find out what the values of the variables bound by those singular existential quantifiers are; those are all and only the entities. Those who accept the Quinean thought that the singular existential quantifier adequately expresses what we mean and understand by *existence*³ might therefore want to understand ‘entity’ as synonymous with ‘existent’, i.e. *that which exists*. With an eye on the availability of the *plural* existential quantifier – using which we can say of *multiple* things at once that they exist (see next section) – we might tweak this slightly and say that ‘entity’ is synonymous with ‘*singular* existent’.⁴

Thus to be an entity is to be one in number *simpliciter*, and to be apt for being the value of a singular variable. So understood, I think the term is fairly ubiquitous and well-understood in metaphysics.⁵

2.2.2 – Plural Quantification

A second important notion to get to grips with here is that of plural first-order quantification. Doing so will give us the logical framework and vocabulary to properly explain the semantics of reductive nihilism. My introduction to this notion will be relatively brief – much more detail and discussion can be found in McKay (2006) and Oliver and Smiley (2013).

I note from the outset that some philosophers regard plural quantification with suspicion: in particular, *singularists* argue that apparent plural quantification is really singular quantification over sets, or mereological fusions, etc. But this view is becoming less and less common, and at any rate I’ll simply assume here that plural quantification is in good standing and that singularism is false.⁶

Plural first-order quantification is a way of extending ordinary first-order quantification in various ways. We start by adding the plural existential quantifier:

$$\exists xx$$

³ See Quine (1953: 13-4); see also van Inwagen (1998: 237).

⁴ A complication is that, in recent times, the notion that there are *multiple* singular existential quantifiers has become somewhat fashionable (e.g. Hirsch 2002, Sider 2011). In light of this, a fair question is: *which* singular existential quantifier is such that to be an entity is to be apt for being the value of a variable bound by that singular existential quantifier? Answer: if there are multiple quantifiers then we are to understand ‘entity’ in terms of the *metaphysically privileged* existential quantifier, where this amounts to the most *natural* quantifier that *carves the world at the joints* better than any other (see Sider 2011). Of course, this assumes that there *is* such a metaphysically privileged existential quantifier. But a commitment to there either being just one existential quantifier, or at least just one metaphysically privileged quantifier, is the hallmark of ontological realism, which is a background assumption of the present project (see chapter 1).

⁵ The task of clarifying terms in such a way as to satisfy *all* readers is probably a hopeless one. As Lewis once remarked: ‘any competent philosopher who does not understand something will take care not to understand anything else whereby it might be explained’ (1986: 203). I also think that anyone who purports not to understand ‘entity’ as I’ve characterised it here is likely to be a deflationist (see e.g. Thomasson [2007: ch. 6]), and as I said in chapter 1, in this project I simply assume the falsity of deflationism.

⁶ See Oliver and Smiley (2013: ch. 3), McKay (2006: ch. 2), and Boolos (1984) for (critical) discussion of various singularist strategies.

This can be read as ‘there are some *xx*’ or ‘there is a plurality, *xx*’. The latter reading is only acceptable, however, if ‘plurality’ is understood to be a genuinely plurally referring term (more on plural reference below), as opposed to referring to some single, set-like entity with multiple members. Indeed, throughout the present project, ‘plurality’ should be understood as being *plural* in this way, even though in most cases it requires a grammatically singular agreement.

There is also the plural universal quantifier:

$$\forall xx$$

This is best read as ‘for each plurality, *xx*’. It is defined in terms of the plural existential quantifier in the usual way (i.e. $\forall xx A \equiv \neg \exists xx \neg A$).

These plural quantifiers bind *plural variables*: *xx*, *yy*, *zz*, etc. Plural variables differ from singular variables in that whilst each singular variable can take as a value only a single entity, each plural variable can take as a value *either* a single entity *or many entities*.⁷

Of crucial importance to plural logic is the primitive ‘is one of’ predicate. This purports to report the relationship between a single entity and a plurality that contains it – for example (and assuming that my laptop is a single entity) the relationship between my laptop and the things in this room. Here I take the limit case of *being one of* to be identity, such that everything is trivially one of itself. This is consistent with the characterisation of ‘is one of’ as reporting the relationship between a single entity and the plurality that contains it, because, again, there can be *improper* pluralities.

More generally, plural logic introduces *plural predicates*, i.e. predicates that can be attributed to pluralities. Plural predicates come in two varieties: distributive and non-distributive. A distributive plural predicate, *F*, is such that it’s analytic that if some *xx* satisfy *F* then each *x* that is one of *xx* also satisfies *F*; a plural predicate is *non-distributive* otherwise. For example, the plural predicate ‘are running’ is distributive, because to say of a plurality of people that they are running analytically implies that each person within that plurality of people is running; on the other hand, the plural predicate ‘are surrounding’ is non-distributive, because to say of some people that they are surrounding something does not analytically imply that any one of those people is surrounding anything. When some entities satisfy a plural predicate that is non-distributive, we can make the non-distributivity clear by saying that those entities satisfy that predicate *jointly* or *collectively*.

At times, I’ll talk as if predicates *refer to* or *express* properties and relations (I suspect that a nominalist paraphrase is available, though I don’t assume this). When talking in this way, I’ll refer to properties (or relations) expressed by plural predicates as *plural properties*.

We’re now in a position to formally state what it is to be a plurality:

$$\text{PLURALITY: } xx \text{ is/are a plurality} \text{ =}_{df} \exists x (x \text{ is one of } xx)$$

An *improper* plurality is the special case where there is just one thing that is one of the plurality:

$$\text{IMPROPER PLURALITY: } xx \text{ is an improper plurality} \text{ =}_{df} \exists !x (x \text{ is one of } xx)$$

⁷ One might find it odd that a *plural* variable can take as a value a single entity. But it is standard in plural logic to think that there are two kinds of plurality: *proper* pluralities, containing more than one entity, and *improper* pluralities, containing just one entity (see below for a more formal definition of an improper plurality). Either kind of plurality can be assigned to a plural variable.

A plurality is a proper plurality iff it is not an improper plurality, i.e.:

PROPER PLURALITY: xx is/are a proper plurality $=_{df} \exists x (x \text{ is one of } xx) \wedge \exists y (y \text{ is one of } xx) \wedge x \neq y$

Four further predicates/relations that are of particular interest and use for the purposes of plural logic can now be defined. First, *amonghood*:

AMONGHOOD: xx are among $yy =_{df}$ for all z , if z is one of xx then z is one of yy

Three things to note about AMONGHOOD:

First, where xx is an *improper* plurality, *amonghood* and *being one of* are materially equivalent. That is, where xx is an improper plurality, xx is one of yy iff xx is among yy . Here's the proof, via an arbitrary case:

In the left-to-right direction: suppose that my laptop is one of the things in this room (so again, assume that my laptop is a single entity). It follows that any z that is one of my laptop is also one of the things in this room, because the only z that is one of my laptop is my laptop itself, which by stipulation is one of the things in this room. Thus my laptop is among the things in this room.

In the right-to-left direction: suppose now that my laptop is among the things in this room. Then for all z , if z is one of my laptop then z is one of the things in this room. My laptop itself is trivially one of my laptop, so my laptop itself is one of the things in this room.

Second, *amonghood* is transitive. Suppose that my laptop and my mug are among the things in this room, and that the things in this room are among the things in my flat. Then any z that is one of my laptop and my mug (i.e. any z such that $vz = \text{my laptop}$ or $vz = \text{my mug}$)⁸ is one of the things in this room, and any z that is one of the things in this room is one of the things in my flat. It follows from this that any z that is one my laptop and my mug is one of the things in my flat; it follows from *this* that my laptop and my mug are among the things in my flat.

Third, AMONGHOOD is reflexive: any plurality is among itself. Consider the plurality of things in this room. It is trivially the case that, for all z , if z is one of the the things in this room then z is one of the things in this room. Thus the things in this room are among the things in this room.

The second predicate to be defined is equivalent to *amonghood*. It is the notion of *being a subplurality of*:

SUBPLURALITY: xx is/are a subplurality of $yy =_{df} xx$ are among yy

Of particular note here is that, given the reflexivity of *amonghood*, every plurality is trivially a subplurality of itself.

We can also define an *irreflexive* notion of *amonghood*, which we can call *proper amonghood*:

PROPER AMONGHOOD: xx are properly among $yy =_{df}$ for all zz , if zz are among xx then zz are among yy , AND it's not the case that for all zz , if zz are among yy then zz are among xx .

Corresponding to *proper amonghood* is the notion of a *strict subplurality*:

⁸ Again, I'm assuming here that the laptop and the mug are each single entities.

STRICT SUBPLURALITY: xx is/are a subplurality of $yy =_{df} xx$ are properly among yy ⁹

The third predicate I want to define here is that of *many-many identity*, aka *plural identity* (unlike some other authors I'll use the same symbol for this, '=', as I use to refer to ordinary one-one identity – I'll clarify whether I mean many-many or one-one identity wherever necessary):

PLURAL IDENTITY: $xx = yy =_{df}$ for all zz , zz are among xx iff zz are among yy .¹⁰

Thus cases of *improper* (i.e. reflexive) amonghood are cases of many-many identity. Further, given AMONGHOOD, an equivalent definition of plural identity is: $xx = yy =_{df} z$ is one of xx iff z is one of yy .¹¹ Given that the limit case of *being one of* is (one-one) identity, the limit case of many-many identity is one-one identity.

Further, I understand PLURAL IDENTITY to obey the indiscernibility of identicals (aka Leibniz's Law):

INDISCERNIBILITY OF IDENTICALS: if $xx = yy$ then $\forall F (Fxx \leftrightarrow Fyy)$

Due to the fact that one-one identity is the limit case of many-many identity, this is a perfectly general definition of the indiscernibility of identicals.¹²

The fourth predicate to be defined is that of *plural overlap*.

PLURAL OVERLAP: xx overlap $yy =_{df}$ there are some zz that are among both xx and yy

The utility of PLURAL OVERLAP will become clear later on.

The last thing to add to this picture is plural *constants*, aka plural *names*, aka plural *referring expressions*: for example, *aa*, *bb*, *cc*, etc. Plural (or plurally) referring expressions refer to a plurality of entities. That is, the relation of plural reference holds between a referring expression (or perhaps an utterance thereof) on the one hand and a plurality of entities on the other. For simplicity's sake we'll allow that a plural referring expression can refer to an improper plurality, but for the most part when I talk of plural referring expressions I'll mean expressions that refer to *proper* pluralities. Intuitive natural language examples of plural referring expressions are 'the Channel Islands' and 'the Beatles', or concatenations of ordinary singular names such as 'Obama and Bush'.¹³ Despite these examples, I note here that plurally referring expressions need not be morphologically plural. Assuming that 'the

⁹ I use 'strict subplurality' here rather than 'proper subplurality' to avoid confusion, as one might think that the natural definition of 'proper subplurality' is 'a subplurality that is a proper plurality'.

¹⁰ For this formulation, see e.g. McKay (2006: 125, 127) and Oliver and Smiley (2013: 109).

¹¹ For this formulation, see e.g. Noonan and Curtis (2018), van Inwagen (1994: 211), and Yi (1999: 184, 2006: 243).

¹² Nevertheless, one might take INDISCERNIBILITY OF IDENTICALS to be *explanatorily posterior* to its ordinary singular counterpart, as would be natural if it turned out that plural properties supervene on singular properties.

¹³ Of course, one might contest both whether these are genuinely (proper) *names* and whether they are genuinely *plural* referring expressions. But they do, I think, serve to illustrate what is meant by a plural constant/name/referring expression. Further, I do think they are additionally plausible examples of plural referring expressions. For example, Rumfitt (2005: 120) argues that 'the Channel Islands' is not a definite description, as the Isle of Wight is an island in the English Channel but is not one of the Channel Islands, and whilst it might be possible to regard 'the Channel Islands' as instead naming a *mereological sum* of islands or perhaps a *set* of islands, its most natural interpretation appears to be as of simply referring plurally to some islands. Moreover, the apparent acceptability of 'Paul McCartney was a Beatle' or 'Paul McCartney was one of the Beatles' (defeasibly) supports the suggestion that 'the Beatles' is plurally referring.

Channel Islands’ refers plurally to a proper plurality of islands, it seems possible for me to stipulate that the morphologically singular ‘the Channel archipelego’ is to be synonymous and co-referential with ‘the Channel Islands’.

Finally, in some cases it will be useful to write plural constants so that they ‘wear their members on their sleeve’, so to speak. So for example, if I intend a plural constant to refer to a plurality such that *a* is one of it, and *b* is one of it, and no other entity is one of it, then I’ll write that plural constant ‘*ab*’. I will make this clear again, though, whenever it comes up.

2.3 - Mereology and the Special Composition Question

I turn now to introducing some mereological concepts, and showing how various answers to the SCQ can be characterised in terms of them.¹⁴

I’m going to put an asterisk next to the names of the mereological concepts I discuss in this section. This is because I want to make clear that these are the definitions of mereological concepts as *mereologists* define them. To pre-figure somewhat, I (*qua* reductive nihilist) disagree with mereologists about the right way to define mereological concepts, at least insofar as these are supposed to be more or less ordinary, folk concepts. So by using asterisks when we want to make clear that we’re using mereological concepts *as defined by the mereologist*, we’ll give ourselves a tool for removing ambiguity.

We take as primitive the notion of *parthood**.¹⁵ This is a two-place relation relating one material entity to another.¹⁶ I repeat this for emphasis: *parthood** relates a *single entity* to a *single entity*; neither relatum of *parthood** can be a proper plurality. The first relatum of *parthood** is the *part**; the second is the *whole**. *Parthood** is typically taken to be a partial ordering relation. Thus it is reflexive (everything is a *part** of itself), transitive (if *x* is *part** of *y* and *y* is a *part** of *z*, then *x* is a *part** of *z*), and anti-symmetric (no distinct *x* and *y* can both be *parts** of the other).¹⁷

*Overlap** is defined as having a *part** in common:

OVERLAP*: *x* overlaps* *y* =_{df} there is a *z* such that *z* is a *part** of *x* and *z* is a *part** of *y*

*Disjointness** is defined as failing to have a *part** in common:

DISJOINT*: *x* and *y* are disjoint* =_{df} *x* and *y* don’t overlap*

I’ll sometimes also say things like ‘*xx* are disjoint*’: this should be understood to mean that the *members* of *xx* are disjoint* from one another. When *xx* is an improper plurality, I’ll understand ‘*xx* are disjoint*’ to mean ‘the one *x* that is one of *xx* is disjoint* from itself’, which is trivially false (trivially, every *x* has a part in common with itself, namely *x* itself).

¹⁴ Much more detail and discussion of mereological concepts and principles can be found in Varzi (2019). Here I limit myself to explaining only as much about mereology as is necessary for the present project.

¹⁵ As Parsons (2014) notes, we don’t have to take *parthood** as primitive: either *proper parthood** or *overlap** would do instead.

¹⁶ Though non-material entities – such as events, actions, etc. – might enter into *parthood** relations with one another, we here focus only on material entities.

¹⁷ Though the reflexivity of *parthood** is generally accepted as simply a stipulated feature, the alleged transitivity and anti-symmetry of *parthood** have been disputed (see Varzi [2019: §2.1] for citations). But we can here simply note the dissent and move on, as regarding *parthood** as transitive and anti-symmetric is standard, and at any rate nothing here turns on these assumptions.

We can also define an irreflexive notion of parthood*, called *proper parthood**:

PROPER PARTHOOD*: x is a proper part* of y =_{df} x is a part* of y and $x \neq y$

Any adequate system of mereology must apparently adopt some principle that (roughly speaking) disallows entities from having just one proper part*, in order to capture the apparent truism that, if x is a proper part* of y then there is *more to y* than x . Various supplementation principles have been proposed as axioms to secure this result. We don't need to take a stand on which such principle should be taken to characterise standard thinking about mereology, but I provide *weak supplementation** as a plausible and representational example:

WEAK SUPPLEMENTATION*: if x is a proper part* of y , then there is some z that doesn't overlap* x and that is a proper part* of y .¹⁸

Thus entities that have at least one proper part* have at least two disjoint* proper parts*. Entities that have proper parts* are *composite** entities:

COMPOSITE*: x is composite* =_{df} x has proper parts*

The relation that some proper parts* (collectively) bear to the x they compose* is called *composition**:

COMPOSITION*: xx compose* y =_{df} (i) each x that is one of xx is a part* of y , AND
(ii) no two of xx overlap*, AND
(iii) every part* of y overlaps* at least one of xx .

Intuitively, composite* entities will typically have multiple *decompositions**: that is, for most y , there are multiple pluralities xx that satisfy COMPOSITION*. For example, suppose that a given table is a composite* entity. The table decomposes* into four table-legs and a tabletop; it also decomposes* into a plurality of atoms; it also decomposes* into a plurality of molecules; etc.

To fail to be composite* is to be *simple**:

SIMPLEHOOD*: x is simple* =_{df} x has no proper parts*

If some x is simple* then we can say that it is *a simple**. Note that, awkwardly, although simples* are not composites*, they are *composed**: given COMPOSITION*, everything trivially composes* itself.

With all of this in mind, we can formulate the SCQ as follows:

SPECIAL COMPOSITION QUESTION: under what conditions do some xx compose* some y ?

We can then define various of the candidate answers to the SCQ that we encountered in the Introduction. First, universalism:

UNIVERSALISM: for all disjoint* xx , there is a y such that xx compose* y .¹⁹

¹⁸ For discussion, see Varzi (2019: §3.1) and Simons (1987: 26-28). Note that Varzi calls what Simons and I call *weak supplementation*[*] just 'supplementation', and uses 'part*' in the consequent rather than 'proper part*'. But Varzi notes (2019: fn 13) that his formulation is equivalent to Simons'/mine given the reflexivity and transitivity of parthood*.

¹⁹ Plausibly, all answers to the SCQ should be understood as being putative *necessary* truths (though see e.g. Cameron 2007), and I'll tacitly understand them as such.

Second, restrictivist answers can be characterised schematically:

RESTRICTIVISM: for all disjoint* xx , there is a y such that xx compose* y iff ϕxx

RESTRICTIVISM says that xx compose* some y whenever xx satisfy a certain condition, ϕ . By filling in the schema with a specification of the condition (which could be a long conjunction or disjunction of conditions), we arrive at various versions of restrictivism. Thus we have, for example:

ORGANICISM: for all disjoint* xx , there is a y such that xx compose* y iff the activity of xx constitutes a life.

CARMICHAEL'S ANSWER: for all disjoint* xx , there is a y such that xx compose* y iff (the xx are either both lump-like and bonded, or the activities of the xx constitute an event that imposes sufficient unity on xx).

Other restrictivist answers will replace ϕ with other conditions.²⁰ For our purposes, the only restriction that we'll place on what ϕ can be is that it must be at least *possible* for some xx to satisfy it, and also at least possible for some xx to *fail* to satisfy it. This is to ensure that restrictivism isn't compatible with universalism or mereological nihilism.

Speaking of mereological nihilism:

MEREOLOGICAL NIHILISM: for all disjoint* xx , there is no y such that xx compose* y .²¹

Thus mereological nihilists believe that there are no composite* entities. It follows that they think that nothing has any proper parts*. For suppose some entity y has a proper part*. Then by WEAK SUPPLEMENTATION* y has n disjoint* proper parts* for some n greater than 1. By COMPOSITION*, y is composed* of those n proper parts, and is composite*. But mereological nihilists deny that there are composite* entities. So nothing has any proper parts*. It also follows from this that mereological nihilists think that every entity is a simple*. That is, mereological nihilists are also committed to:

UNIVERSAL SIMPLEHOOD*: $\forall x (x \text{ is simple}^*)$

Indeed, UNIVERSAL SIMPLEHOOD* is equivalent to MEREOLOGICAL NIHILISM.

I'll often use 'nihilism' as an abbreviation of 'mereological nihilism'. Let me stress, then, that *all* nihilists must *by definition* accept MEREOLOGICAL NIHILISM, and therefore UNIVERSAL SIMPLEHOOD* (also, anyone who accepts these two principles is by definition a nihilist).

Even ignoring what semantics the nihilist adopts, there are at least three versions of nihilism *qua* metaphysical view. They all accept the claim that

$\forall x (x \text{ is simple}^*)$

But they differ on what the simples* are.

²⁰ One might wonder how a restrictivist version of Markosian's (1998) Brutal Composition fits into RESTRICTIVISM's schema. Whilst Brutal Composition doesn't explicitly offer any answer to the SCQ, and rather states that there is no finite, non-trivial answer to the SCQ, it should be combinable with restrictivism. One way to characterise such a combination of Brutal Composition and restrictivism is as claiming that the ϕ that the xx must satisfy to compose a y is simply that xx compose a y .

²¹ Recall that when xx is an improper plurality then xx is trivially not disjoint*, so MEREOLOGICAL NIHILISM is consistent with the fact that, by COMPOSITION*, every x trivially composes some y , namely itself.

The most familiar version of nihilism is:

MICROPHYSICAL NIHILISM: $\forall x (x \text{ is simple}^*) \wedge \forall x (x \text{ is a simple}^* \text{ iff } x \text{ is either a quark or a lepton})$

Strictly speaking, MICROPHYSICAL NIHILISM is probably at best only a materially adequate characterisation of the most familiar version of nihilism. For the familiar version of nihilism is probably best understood as committing itself only to identifying simples* with whatever the physicists tell us are the smallest, most fundamental physical particles: if the physicists tell us that they've discovered that quarks and leptons are made out of smaller particles, then the familiar version of nihilism will identify simples* with those smaller particles; similarly, if the physicists tell us that there are fundamental particles other than just quarks and leptons, then the familiar version of nihilism will agree that those particles are simples* too. But having noted this, for simplicity we'll stick with MICROPHYSICAL NIHILISM.

Further, though MICROPHYSICAL NIHILISM by itself doesn't require this, we should understand microphysical nihilism to accept something like the following principle: wherever observation or empirical study together with the assumption of universalism tell us that there is a composite* F , there are instead some simples* arranged F -wise. Thus where universalists posit tables, microphysical nihilists posit pluralities of simples* arranged table-wise; where universalists posit an entity composed* by Bill Clinton's nose and the Eiffel Tower, microphysical nihilists posit a plurality of simples* arranged nose-Eiffel-Tower-wise; etc.

Perhaps the second most familiar version of nihilism is *existence monism*, which is the view that there is just one, extended simple*, which is identical to the universe:²²

EXISTENCE MONISM: $\exists x (x \text{ is simple}^* \wedge x \text{ is the universe} \wedge \forall y (y = x))$

The third type of nihilism I'll consider is *emergence nihilism*.²³ Emergence nihilism characteristically claims that there are all sorts of different simples*, ranging from the familiar quarks and leptons to extended simples* such as tables, chairs, human beings, planets, etc. Indeed, emergence nihilists can accept any ontology of material entities they like, as long as they claim that every element in that ontology is simple*. Thus, for example, one kind of emergence nihilist will agree with a certain sort of restrictivist that there exist tables, chairs, mountains, molecules, etc., but disagrees with the restrictivist that these things are *composite** entities, maintaining instead that they are simples*. It's not obvious how to characterise emergence nihilism in a pithy way. I will settle here for:

EMERGENCE NIHILISM: $\forall x (x \text{ is simple}^*) \wedge \exists x (x \text{ is not a quark} \wedge x \text{ is not a lepton} \wedge x \text{ is not the universe})$

The first conjunct of EMERGENCE NIHILISM entails MEREOLOGICAL NIHILISM; the other conjunct distinguishes it from microphysical nihilism and existence monism.

That concludes my short introduction to the existing definitions of mereological concepts and theories that are important to the present project.

²² See Schaffer (2007) and Horgan and Potrč (2008) for arguments for existence monism.

²³ See Williams (2006), though Williams doesn't endorse that emergence nihilism is actually true.

2.4 – Reductive Nihilism

Reductive nihilists accept MICROPHYSICAL NIHILISM. They combine this metaphysical claim with certain non-standard semantic claims that distinguish their view from other nihilist views.

The reductive nihilist’s non-standard semantics are helpfully understood by contrasting them with the *standard* semantics that microphysical nihilists typically endorse. For reasons that will become clear, I’ll call the conjunction of microphysical nihilism with these standard semantics *eliminative nihilism*. In §2.4.1 I’ll detail the semantic claims of eliminative nihilism. In §2.4.2 I’ll present and explain the semantic claims that the reductive nihilist makes instead. One of these claims will require more discussion, which I’ll provide in §2.4.3.

2.4.1 – Eliminative Nihilism

Microphysical nihilism is typically regarded by friend and foe alike, as a form of *eliminativism* about composite objects.²⁴

ELIMINATIVISM: Composite objects do not exist.

Here’s the sort of reasoning that usually underlies the idea that microphysical nihilists are committed to ELIMINATIVISM. Assume microphysical nihilism and suppose for *reductio* that eliminativism is false. Thus there exists a composite object. By COMPOSITE*, that object must have proper parts*. But MICROPHYSICAL NIHILISM says that everything is simple*, and SIMPLEHOOD* says that simples have no proper parts*. Contradiction. So, one might think, if microphysical nihilism is true then eliminativism must also be true.

This apparent commitment to eliminativism is also typically taken to require the microphysical nihilist to deny the existence of familiar objects such as tables, chairs, mountains, molecules, human beings, honey bees, the Eiffel Tower, etc. For such things appear to have parts, and as such appear to be composite objects, and as such cannot exist, given ELIMINATIVISM.

We might not even think that we need to appeal to ELIMINATIVISM to show that microphysical nihilism must reject the existence of such familiar objects. Instead we might simply note that microphysical nihilists by definition accept:

$$\forall x (x \text{ is a simple}^*) \wedge \forall x (x \text{ is a simple}^* \text{ iff } x \text{ is either a quark or a lepton})$$

Since tables, chairs, etc. are neither quarks nor leptons, their existence seems incompatible with microphysical nihilism.

What underlies these sorts of arguments for the claim that microphysical nihilism is a form of eliminativism seems to be a background commitment to a certain package of semantic claims, namely:

(EN1) ‘Existence’ and its cognates express the singular existential quantifier, $\exists x$.

(EN2) ‘Object’ is synonymous with ‘entity’, as I characterised it in §2.2.1.

²⁴ Examples of microphysical nihilists who characterise the view in this way are Rosen and Dorr (2002), Grupp (2006), Cornell (2017), Brenner (2018). See also Merricks (2001), who isn’t nihilist, but defends a view closely related to microphysical nihilism. We can loosely understand Merricks as being a microphysical nihilist about all non-conscious composite objects, and he characterises himself as an eliminativist about such objects.

- (EN3) Mereological vocabulary such as ‘parthood’ and ‘composition’ are adequately defined by the definitions I gave in §2.3. Thus ‘parthood’ expresses parthood*, ‘composition’ expresses composition*, etc.
- (EN4) Sortal predicates of the form ‘is an F ’²⁵ are satisfied by a single entity, if at all.
- (EN5) Predicates that putatively ascribe properties to single things are satisfiable only by single entities.
- (EN6) Putatively singular referring expressions each refer to a single entity, if they refer at all.

As I’ll shortly argue, by rejecting (EN1)-(EN6) and adopting alternative semantic claims in their place, the microphysical nihilist can avoid commitment to eliminativism (as well as – as we’ll see – to other counterintuitive claims that accepting (EN1)-(EN6) implies).

Call the conjunction of microphysical nihilism with semantic claims (EN1)-(EN6) *eliminative nihilism*.²⁶ Reductive nihilists replace (EN1)-(EN6) with their own semantic claims. I turn now to presenting these alternative, reductive nihilist semantic claims.

2.4.2 – The Semantic Claims of Reductive Nihilism

In this section I’ll detail what claims the reductive nihilist replaces (EN1)-(EN6) with. Along the way I’ll show how this allows the reductive nihilist to avoid commitment to ELIMINATIVISM, as I think doing so is a useful way of making her semantic claims clear.

Before I begin, however, I want to note that the reductive nihilist does not deny that it is *possible* to use terms and predicates in the way described by (EN1)-(EN6). Instead the reductive nihilist’s point is as follows. We have an intuitive, pre-theoretic grasp of the meanings of the bits of vocabulary described by (EN1)-(EN6). For example, we have an intuitive grasp of what it is for something to *be a part of* something. Philosophers, including eliminative nihilists, have typically assumed that (EN1)-(EN6) succeed in defining or characterising the meaning of these bits of intuitively-grasped vocabulary. But, I argue, these philosophers are mistaken, at least given the background assumption of microphysical nihilism. Given microphysical nihilism, (EN1)-(EN6) do not satisfactorily define or characterise the ordinary notions that they concern. Philosophers may feel free to use terms and predicates that *are* defined by (EN1)-(EN6), if they like, but they are thereby using metaphysical terms of art. Insofar as we’re interested in defining vocabulary that we had a pre-theoretic, intuitive grasp of, we should ignore (EN1)-(EN6) and replace them with alternative semantic claims.

The alternative semantic claims of reductive nihilism build on and extend two previous microphysical nihilist accounts in the literature: that of Liggins (2008), and that of Contessa (2014).²⁷ I’ll note which claims are taken from or similar to claims of these accounts.

²⁵ Strictly speaking these quotation marks should be replaced with Quinean corner quotes. I’ll continue to use ordinary quotation marks, and trust the reader to interpret them as instances of quasi-quotation where appropriate.

²⁶ It’s worth noting, though, that (EN1)-(EN6) are not semantic claims that are *unique* to eliminative nihilists: rather, they are semantic claims that the vast majority of the disputants in the debate over the SCQ accept. For example, (EN1) is accepted by anyone who endorses a Quinean metaontology (see Quine 1953: 13-4; van Inwagen 1998), and (EN3) is absolutely standard among mereologists (see Varzi 2019).

²⁷ I note that Liggins doesn’t explicitly *endorse* the account he outlines, but rather suggests it on behalf of van Inwagen. I note also that there is a wider tradition of trying to use semantic considerations to show that microphysical nihilism or similar views can avoid ELIMINATIVISM. It’s outside the scope of the present project

First, reductive nihilists replace (EN1) with:

(RN1) ‘Existence’ and its cognates express the plural existential quantifier, $\exists xx$.²⁸

Second, they replace (EN2) with:

(RN2) ‘Object’ is not synonymous with ‘entity’, but is rather a plural predicate that can be satisfied by proper pluralities of entities. Roughly, xx satisfy ‘object’ iff it is intuitive that the region of space exactly occupied by xx is exactly occupied by an object.²⁹

For example, reductive nihilists say that any plurality of simples* arranged table-wise are jointly an object. For consider the region of space exactly occupied by a given plurality of simples* arranged table-wise: intuitively, that region of space is exactly occupied by an object (namely, a table).

Now, it seems plausible the predicate ‘[is an] object’ can have a variety of meanings in ordinary language. Reductive nihilists maintain that, contra (EN2), none of these ordinary meanings correspond to ‘entity’ (at least on the assumption of microphysical nihilism). (RN2) then concerns arguably the most inclusive meaning of ‘object’ in our ordinary vocabulary, namely the sense that corresponds roughly to ‘unified, single whole’. It’s the sense of ‘object’ that we withhold when confronted with what seems to us like nothing more than a disparate collection of things (like a nose and the Eiffel Tower) but that we’d apply when we’re confronted with what seems to us like a genuine unified whole (like a nose, or the Eiffel Tower, or perhaps a constellation of stars, or a penguin, etc.). This is consistent with there also being an available sense of ‘object’ in ordinary language on which it is analytic that objects are inanimate, or internally well-bonded, etc. Those senses can be arrived at by tweaking the characterisation of ‘object’ in (RN2) (for example, by adding the condition that in order to count as an object, xx must be inanimate, or internally well-bonded, etc.). But the ordinary sense of ‘object’ that it is most interesting and instructive to focus on here is the one corresponding roughly to ‘unified, single whole’, as this is the sense of ‘object’ that tracks our intuitions about when composition occurs. So from now on we can ignore other senses of ‘object’ and focus on the one at issue in (RN2).

I also note that the biconditional in (RN2) is not supposed to supply a reductive nihilist *analysis* of ‘object’ (if it were an analysis it would be a bad one, as ‘object’ appears on the right-hand-side). Rather, the biconditional simply establishes the relationship between our intuitions about when ‘object’ applies and when ‘object’ really applies, given reductive nihilism. Indeed, reductive nihilists regard the project of specifying non-trivial necessary and sufficient conditions for satisfying ‘object’ as a likely hopeless one. Nevertheless, drawing on what others have said, they can offer the following *hallmarks* of objecthood – that is, they can provide a rough idea of the sort of features that pluralities that satisfy ‘object’ tend to have:

to discuss these other accounts here. For the record, though, see in particular Cameron’s (2008) and (2010), in which he claims that a microphysical nihilist ontology suffices to *truthmake* claims such as ‘the Taj Mahal exists’, and van Inwagen’s remarks about why his organicism doesn’t contradict ordinary belief (1990a: ch. 10) (as Liggins says, Liggins’ [2008] account *could* have been what van Inwagen had in mind, but equally might well not have been). See also Morton (1975) and Hossack (2000).

²⁸ See Liggins (2008: §2 esp. p. 181 fn11) and Contessa (2014: 206).

²⁹ See Contessa (2014: 216-7). In addition, there are several other precedents for the thought that ‘object’ (or ‘material object’, which is what is meant by my ‘object’) might differ in meaning to ‘entity’, though not in the way that I or Contessa suggest: see e.g. Nolan (2014), Lowe (1997: esp. 35), or the systems of ontological categories detailed in Westerhoff (2005: ch. I).

- xx are well-bonded together.³⁰
- xx display a unifying causal relation.³¹
- xx are adjacent to one another,³² or display spatiotemporal continuity.³³
- xx contrast with their surroundings more than they do with each other.³⁴
- xx act jointly³⁵ or have a joint function.³⁶
- The activities of xx constitute an event that imposes sufficient unity on the xx .³⁷

Additionally, reductive nihilists claim that all entities are objects. Thus (RN2) should not be taken to imply that ‘object’ is a plural predicate that can *only* be satisfied by proper pluralities of entities: it can also be satisfied by improper pluralities, i.e. individual entities.

I relegate two more minor comments about (RN2) to a footnote.³⁸ Here, I’ll move on.

Third, reductive nihilists reject (EN3). They claim that ordinary terms such as ‘parthood’, ‘composition’, etc. do not receive their definitions from PARTHOOD*, COMPOSITION*, etc. They don’t deny that *mereologists* often mean parthood* by ‘parthood’, and composition* by ‘composition’, but claim that insofar as PARTHOOD*, COMPOSITION*, etc. are supposed to be definitions of ordinary mereological notions rather than metaphysical terms of art, they fail. In place of PARTHOOD*, COMPOSITION*, etc., reductive nihilists offer the following definitions.³⁹

PARTHOOD: xx is/are a part of $yy =_{df} xx$ are among yy

Thus (according to reductive nihilists), to be a *part* is to be a *subplurality*. Parthood so understood inherits the formal features of amonghood: thus, among other things, it is transitive and reflexive.

The irreflexive version of parthood is:

PROPER PARTHOOD: xx is/are a proper part of $yy =_{df} xx$ are properly among yy .⁴⁰

³⁰ Thomasson (2009: 458; 2015: 110); Carey (2009: ch. 3); see also Lewis (1986: 211).

³¹ Hoffman and Rosenkrantz (1997: 4).

³² Lewis (1986: 211).

³³ Carey (2009: ch. 3).

³⁴ Lewis (1986: 211).

³⁵ Lewis (1986: 211).

³⁶ c.f. Rose and Schaffer (2017).

³⁷ Carmichael (2015: 479).

³⁸ Comment 1: the reference to regions of space in (RN2) is for ease of expression, and should be eliminable as long as reference to reified locations or regions of space is in general eliminable (a question I don’t intend to take a stand on here).

Comment 2: I intend to bracket and ignore borderline cases in which intuition is undecided as to whether a given region of space is exactly occupied by an object or not; I’ll also ignore the possibility of cases in which our intuitions about whether a given region of space is exactly occupied by an object or not vary depending on contextual factors. Such cases are peripheral to the reductive nihilist’s headline semantic claim.

³⁹ Liggins (2008: 190) agrees with the reductive nihilist’s claims about *parthood* (with the caveat I note in the next footnote), but doesn’t discuss other mereological vocabulary. Contessa (2014: 217-8) also makes some brief comments suggestive of the reductive nihilist’s PARTHOOD.

⁴⁰ Liggins (2008: 190) defines ‘parthood’ as what I’m calling *proper* parthood. I’m happy to grant that, typically, when we say ‘parthood’ we mean what I’m calling proper parthood, but I’ll stick with my definitions

Thus to be a *proper part* is to be a *strict subplurality*.

Just as *overlap** is defined as having a *part** in common, and *disjointness** is defined as failing to have a *part** in common, so *overlap* is defined as having a part in common and *disjointness* is defined as failing to have a part in common:

OVERLAP: xx overlap yy =_{df} there are some zz such that zz is/are a subplurality of both xx and yy

Thus *overlap* equates to *plural overlap* (see §2.2.2).

DISJOINT: xx and yy are disjoint =_{df} xx and yy don't overlap

What's more, just as *proper parthood** must obey *weak supplementation** (or some similar principle, at least), so *proper parthood* obeys *weak supplementation*:

WEAK SUPPLEMENTATION: if xx is/are a proper part of yy , then there are some zz such that zz don't overlap xx and zz is/are a proper part of yy

In fact, WEAK SUPPLEMENTATION follows from PROPER PARTHOOD (note that WEAK SUPPLEMENTATION* does not follow from PROPER PARTHOOD*). Suppose that xx is/are a proper part of yy . By PROPER PARTHOOD, xx are properly among yy . So, by PROPER AMONGHOOD, some uu are among yy without being among xx . Thus, some u that is one of uu is one of yy but not one of xx (otherwise uu would be among xx). u is among yy (because u is among uu , and uu are among yy , and amonghood is transitive), and therefore is a proper part of yy . u does not overlap xx (because the only way for u to overlap xx would be for u to be one of xx , which we've ruled out). Now let $vu = vzz$ (this is legitimate because, recall, plural variables like zz can take *improper* pluralities as their values): thus there is a zz such that zz don't overlap xx and zz is/are a proper part of yy . So WEAK SUPPLEMENTATION follows from PROPER PARTHOOD.

Next:

SIMPLEHOOD: xx is simple =_{df} xx has no proper parts

Equivalently, to be a simple is to be an entity. For suppose a given quark is simple. By SIMPLEHOOD it has no proper parts; by PROPER PARTHOOD and STRICT SUBPLURALITY, it has no strict subpluralities. It follows from our definitions that any xx that are n in number have $2^n - 2$ strict subpluralities, so if the quark were n in number for some n greater than 1 then it would have at least two strict subpluralities. So our quark must be one in number *simpliciter*, i.e. it must be an entity. Similarly suppose a given lepton is an entity, and thus one in number *simpliciter*. Then it has no strict subpluralities, and so, by SIMPLEHOOD, is simple.

Equivalently, since to be an entity is to be an improper plurality: xx is simple =_{df} xx is an improper plurality.

Note that, given MERELOGICAL NIHILISM, xx is simple iff xx is simple*. Suppose xx is simple. Then it is a single entity. MERELOGICAL NIHILISM is equivalent to UNIVERSAL SIMPLEHOOD*: $\forall x (x \text{ is simple}^*)$. Since *singular* variables range over all and only *single entities*, we can read this as: every single entity is simple*. Thus xx is a simple*. Now suppose that xx is simple*. Strictly speaking,

of 'parthood' and 'proper parthood' purely in order to make clear how my definitions are analogous to the mereologist's definitions.

SIMPLEHOOD* doesn't tell us about the conditions under which some xx are simple*, as it only defines what it is for an x to be simple*. But I take it that the fact that simplehood* is defined in this way using a singular variable suggests that mereologists take *being a single entity* as a necessary condition on being simple*. Thus xx is a single entity. Thus xx is simple.

Despite this, simplehood and simplehood* are not logically equivalent, as if MERELOGICAL NIHILISM were false then there is at least one entity (i.e. one simple) that is not simple*.

I save the reductive nihilist's definition of *composition*, COMPOSITION, for the next section, as it requires a longer discussion. That just leaves:

COMPOSITE: xx is/are composite =_{df} xx has proper parts

Equivalently: xx is/are composite =_{df} xx is a proper plurality. I note that 'is/are composite' is a non-distributive plural predicate.

In sum, and on the understanding that the reductive nihilist's definition of composition, COMPOSITION, will be provided shortly, reductive nihilists replace (EN3) with:

(RN3) Mereological terms such as 'parthood' and 'composition' are defined by PARTHOOD, COMPOSITION, etc.

Now, we already have enough to show that reductive nihilists are not committed to ELIMINATIVISM. For reductive nihilists can agree that there are composite objects. Given her semantics, 'there are composite objects' expresses

$\exists xx (xx \text{ are an object} \wedge \exists yy (yy \text{ are properly among } xx))$

Assume microphysical nihilism, and consider the plurality of simples* arranged table-wise that I am currently leaning on. There are certainly some yy that are properly among the simples* arranged table-wise: the simples* arranged table-leg-wise, for example. Moreover, as above the reductive nihilist takes the simples* arranged table-wise to satisfy 'are an object'. Thus reductive nihilists can point to the plurality of simples* arranged table-wise as an example of a composite object, and as such reductive nihilists have a counterexample to ELIMINATIVISM.

However, we *haven't* yet shown that reductive nihilists can accept the existence of tables, chairs, mountains, molecules, human beings, honey bees, etc., because we haven't yet said anything about what the reductive nihilist takes 'is a table', 'is a chair', 'is a mountain', etc., to mean.

In my discussion of eliminative nihilism I called predicates like these *sortal* predicates, because they purport to tell us what *sort* of thing whatever they apply to is. Whereas eliminative nihilists, via (EN4), claim that sortal predicates can only be satisfied by single entities, reductive nihilists instead accept the following principle:

(RN4) Sortal predicates that purport to apply to single composite F s express properties that are instantiated by xx iff xx are both arranged F -wise and are an object.

For example, reductive nihilists claim that 'is a table' is satisfied by the (proper) plurality of simples* that are arranged table-wise, since those simples* are both arranged table-wise and, as above, count as an object.

Together with her claim that ‘existence’ expresses the plural existential quantifier, this claim about sortal predicates means that the reductive nihilist accepts the existence of tables, chairs, mountains, molecules, human beings, honey bees, etc. For example, the reductive nihilist analyses the claim ‘tables exist’ as

$\exists xx$ (xx are a table)

Since there are some simples* arranged table-wise which are additionally an object, there are some xx that are a table, and thus ‘tables exist’ is true.

The reductive nihilist’s story about sortal predicates also allows her to deny the existence of gruesome composites. In general, this is because no plurality of simples* arranged gruesome-composite-wise counts as an object. For example, consider the sortal predicate ‘is a nose-Eiffel-Tower’, which is satisfied by whatever is composed of a human nose and the Eiffel Tower (cf. Hirsch 2002: 51). Given what the reductive nihilist says about sortal predicates, this sortal predicate is satisfied by xx iff xx are arranged nose-Eiffel-Tower-wise and are an object. But there is no such xx . It is not intuitive that regions of space that are exactly occupied by pluralities of simples* arranged nose-Eiffel-Tower-wise are exactly occupied by objects, and hence, by (RN2), ‘object’ doesn’t apply to any plurality of simples* arranged nose-Eiffel-Tower-wise. Thus there is no xx such that xx are a nose-Eiffel-Tower; thus the nose-Eiffel-Tower does not exist, given the reductive nihilist’s semantics.^{41,42}

The reductive nihilist also has a story to tell about *non*-sortal (and *non*-mereological) predicates, such as ‘is 324m tall’ or ‘is over a mile away from’, etc. In each such case, she claims that these predicates express corresponding plural properties, such as *being (jointly) 324m tall* or *are over a mile away from*. Thus, in place of (EN5), she accepts:

(RN5) Predicates that putatively ascribe properties to single things express plural properties that are satisfiable by proper pluralities (thus, for example, ‘weighs 10kg’ express *(jointly) weighing 10kg*).⁴³

Thus reductive nihilists claim that most predicates that purport to ascribe properties to single things are really plural predicates. In most cases, these will be *non-distributive* plural predicates. For example, in the sentence, ‘the Eiffel Tower is 324m tall’, ‘the Eiffel Tower’ refers plurally to the simples* arranged Eiffel-Tower-wise, and ‘is 324m tall’ is a non-distributive plural predicate that says of those simples* that they are jointly 324m tall, without thereby implying that any simple* that is one of the simples* arranged Eiffel-Tower-wise is 324m tall. However, note that such plural predicates will not be non-distributive in every case. Take for example the predicate ‘is over a mile away from’, as it appears in e.g. ‘the Eiffel Tower is over a mile away from Big Ben’. Here, ‘is over a mile away from’ is a plural predicate expressing a relation between the simples* arranged Eiffel-Tower-wise and those arranged Big-Ben-wise. It seems as though it’s analytic that if the simples* arranged Eiffel-Tower-wise are over a mile away from the simples* arranged Big-Ben-wise, then each simple* that is one of the simples* arranged Eiffel-Tower-wise is over a mile away from the simples* arranged Big-

⁴¹ I think this element of the reductive nihilist’s semantics is an improvement over that of Liggins, who claims instead that ‘is an F ’ is synonymous with ‘are arranged F -wise’ (2008: 182). For unlike the reductive nihilist, Liggins’ microphysical nihilist must apparently accept that the nose-Eiffel-Tower exists (or else deny that there is any plurality of simples arranged nose-Eiffel-Tower-wise).

⁴² See also §6.3, where I explicitly defend the claim that, given reductive nihilism, there exists a (composite) F iff it is intuitive that exists a (composite) F .

⁴³ See Liggins (2008: 182) and Contessa (2014: 20fn15), who include something like this claim in their semantics too.

Ben-wise. Thus ‘is over a mile away from’ is distributive in its first place (and for similar reasons, it is also apparently distributive in its second place).

Lastly, for this section, reductive nihilists replace (EN6) with:

- (RN6) Referring expressions of the form ‘N’ that purport to refer to a single composite F refer plurally to xx iff xx are arranged N -wise and xx are an object.⁴⁴

For example, the reductive nihilist claims that ‘the Eiffel Tower’ refers to the simples* arranged Eiffel-Tower-wise (for those simples*, given (RN2), are jointly an object). Since she thinks that these simples* exist, the reductive nihilist therefore accepts the existence of the Eiffel Tower.

However, consider the referring expression ‘Nose-Eiffel-Tower’, which I hereby stipulate to be the name of the thing, should there be one, composed of Bill Clinton’s nose and the Eiffel Tower. The simples* arranged Clinton’s-nose-Eiffel-Tower-wise do not jointly count as an object, given (RN2). Thus ‘Nose-Eiffel-Tower’ fails to refer, which the reductive nihilist takes to mean that Nose-Eiffel-Tower does not exist.

2.4.3 – Composition for Nihilists

We are not quite finished with the exposition of the semantic claims of reductive nihilism. In this section I present the reductive nihilist’s definition of *composition*. I do so by first noting that reductive nihilists reject that ‘composition’ expresses composition*, and also that they reject a definition of ‘composition’ that results from tweaking COMPOSITION* in a natural way (§2.4.3.1). I propose instead that resources from superplural logic are needed to adequately define composition for the reductive nihilist. In §2.4.3.2 I introduce superplural logic; in §2.4.3.3 I present the reductive nihilist’s definition of composition.

2.4.3.1 – What Composition Doesn’t Mean

First, recall again the *mereologist’s* definition of composition (which I here present in a slightly more formalised form than above, in order to facilitate my subsequent discussion):

COMPOSITION*: xx compose* $y =_{df}$

- (i) $\forall x$ (if x is one of xx then x is part* of y) AND
- (ii) $\forall x \forall z$ (if x is one of $xx \wedge z$ is one of $xx \wedge x \neq z$, then x doesn’t overlap* z) AND
- (iii) $\forall x$ (if x is a part* of y then $\exists z$ (z is one of xx and x overlaps* z))

As we’ve seen, eliminative nihilists take ‘composition’ to express composition*. As such they deny that composition ever takes place, because (among other reasons) they deny that there are any x and y such that x is a part* of y .

But reductive nihilists do not accept that ‘composition’ expresses composition*. This is unsurprising, for at least two reasons. First, composition is supposed to be a relation defined in terms of parthood and overlap, and since reductive nihilists deny that ‘parthood’ and ‘overlap’ expresses parthood* and

⁴⁴ Minus the claim that the referents of putatively singular referring expressions must be objects, this is also a central component of Liggins’ account (2008). Contessa (2014) doesn’t explicitly mention referring expressions (his discussion centres on sortal predicates instead), but would plausibly be sympathetic to (RN6).

overlap* respectively, it's no surprise that they don't take composition to be defined in terms of parthood* and overlap*. Second, composition is supposed to be a relation that a whole's non-overlapping parts bear to it. But given reductive nihilism's other semantic claims, neither the whole object nor any of its parts need be single entities (indeed, an object with parts *cannot* be a single entity, on reductive nihilism), whereas COMPOSITION*'s use of singular variables (and its use of 'is one of') requires that both the whole object and each of its parts are single entities.

The natural (but unsuccessful) way to try to adapt COMPOSITION* to end up with a definition of composition acceptable to reductive nihilists, then, would be to replace all singular variables in COMPOSITION* with plural ones, all instances of 'is one of' with 'are among', and parthood* and overlap* with parthood and overlap. That would give us the following candidate definition of composition:

xx compose yy =_{df}

- (i) $\forall zz$ (if zz are among xx then zz is/are part of yy) AND
- (ii) $\forall zz \forall ww$ (if zz are among $xx \wedge ww$ are among $xx \wedge zz \neq ww$, then zz and ww don't overlap) AND
- (iii) $\forall zz$ (if zz is/are a part of yy then $\exists ww$ (ww are among zz and zz overlap ww))

But this candidate definition of composition is objectionable, because (ii) is much too hard to satisfy.

For example, intuitively, it ought to be the case that my mug is composed of a handle and a cylindrical body. That is, intuitively, the handle and the cylindrical body compose my mug. Given reductive nihilism's other semantics, that means that we want the plurality of simples*, call it aa , made up of the simples* arranged handle-wise and the simples* arranged cylindrical-body-wise, to compose the simples* arranged mug-wise. That is, we want it to be the case that when we assign $vxx = aa$ and $vyy =$ the simples* arranged mug-wise, (i)-(iii) are satisfied. But (ii) is not satisfied in this case, for there are a distinct zz and ww such that both are among aa , and that overlap. For example, let $vzz =$ the simples* arranged handle-wise, and let $vww =$ the plurality of simples* made up of the simples* arranged handle-wise *plus* one simple* of the simples* arranged cylindrical-body-wise. Given this variable assignment, zz and ww are both among aa , and are not identical to each other, but they overlap (the simples* arranged handle-wise are a subplurality of both). So, on this definition of composition, it would come out as false that the handle and the cylindrical body compose the mug. This is intuitively the wrong result.

Happily, reductive nihilists do not endorse this candidate definition of composition. Her definition instead utilises resources from *superplural logic*. I'll now introduce the notion of superplural logic, as I understand it, before using it to present the reductive nihilist's definition of composition.

2.4.3.2 – Superplural Logic

Superplural logic is supposed to be to plural logic what plural logic is to ordinary singular logic. Whilst plural logic allows us to talk and reason about many individuals at once, superplural logic purports to allow us to talk and reason about many *pluralities* of individuals at once.⁴⁵

⁴⁵ Hazen (1997: 247) is an early discussion of superplurals; see also e.g. Rayo (2006), Linnebo and Nicolas (2008), Oliver and Smiley (2013: §8.4), and Linnebo (2017: §2.4).

Consider, for example, the referring term ‘the countries of Great Britain and the countries on the Irish landmass’. This is a *prima facie* good example of a superplural term, in that it’s a term that can be understood as referring to two pluralities at once: the countries of Great Britain (plurality number one), and the countries on the Irish landmass (plurality number two). Claims in which ‘the countries of Great Britain and the countries on the Irish landmass’ features in the subject position therefore seem to be best analysed through superplural logic.

A crucial feature of superplural logic, as I understand it, is that employing it does not require us to expand our ontology to include *superpluralities* in addition to individuals, any more than employing *plural* logic requires us to expand our ontology to include *pluralities* in addition to individuals. That is, superplural logic is *ontologically innocent*. Just as plural logic (i.e. plural quantification, plural predication, plural reference, etc.) is generally understood as simply giving us a new way to *talk about* the same old individuals that we posited all along, so superplural logic (i.e. superplural quantification, superplural predication, superplural reference, etc.) purports to give us a *different* new way to talk about those same old individuals.⁴⁶

Roughly speaking, superplural logic purports to let us talk about a certain number of individuals *conceived of as a plurality of pluralities*. For example, consider again ‘the countries of Great Britain and the countries on the Irish landmass’. This refers to the individual countries of England, Scotland, Wales, Northern Ireland, and Ireland. But it does so in a way that treats these five countries as dividing into two subpluralities: England, Scotland, and Wales (plurality number one), and Northern Ireland and the Republic of Ireland (plurality number two). Here’s another example. Suppose there are nine people in a room, three of whom are (adult) men, three of whom are (adult) women, and three of whom are children. Suppose also that four of those people are German and five of them are Austrian. We can refer to those nine people with (at least) two different superplural terms: ‘the men, the women, and the children’, and ‘the Germans and the Austrians’. Whilst both of these referring expressions invite us to conceive of the nine people as being a plurality of pluralities, they differ from one another with regard to *which pluralities* they invite us to conceive of the nine people as being a plurality of: ‘the men, the women, and the children’ invites us to conceive of the nine people as being a plurality of three pluralities, namely the plurality of the three men, that of the three women, and that of the three children; ‘the Germans and the Austrians’ invites us to conceive of the nine people as being a plurality of two pluralities, namely the plurality of the four Germans and the plurality of the five Austrians.

The notion of conceiving of some individuals as being a plurality of pluralities may be somewhat vague. But we can get a stronger grip on this notion by considering the effect it seems to have on *truth conditions*. For example, consider again the nine people we discussed above, and consider the following sentences:

P: The Germans and the Austrians are separated from each other.

Q: The men, the women, and the children are separated from each other.

⁴⁶ Of course, it’s open to philosophers to dispute that plural logic is genuinely ontologically innocent in this way in the first place (as I noted in §2.2.2), and also for philosophers to grant the plural logic is ontologically innocent but nonetheless insist that superplural logic is not ontologically innocent (additionally, see Linnebo (2004) who argues *from* the claim that superplural logic is not ontologically innocent to the claim that plural logic isn’t either). I cannot defend the ontological innocence of superplural logic here, and will rather simply assume it. I do note, though, that everything I have to say about superplural logic is at least *prima facie* consistent with an ontology containing only individuals (where no such individual is a superplurality).

Intuitively, P and Q have different truth conditions; at least, I think it's clear that there is an available reading of P and Q such that they have different truth conditions. For example, consider a situation in which the four Germans, including a man, a woman, and two children, are put in one room, and the five Austrians, including two men, two women, and a child, are put in another. Then P seems to be true, but Q seems to be false; at least, there is an available reading of P and Q that gives us this result. Now consider a situation in which all three men are put in a room together, all three women are put in a different room together, and all three children are put in a third room together. Now P seems to be false (because some Germans and Austrians remain unseparated) but Q seems to be true. Thus it seems that the way in which we superplurally refer to the nine people – and thus the way in which we conceive of those nine people as dividing up into pluralities – affects the truth values of sentences about those nine people.

Note further that, given what I've said, the predicate 'are separated from each other' as it appears in both P and Q appears to be a *superplural predicate*, in that it takes a superplural term in its only argument place.

Another way to get a grip on the idea behind superplural logic is to consider the difference between P and Q, on the one hand, and a corresponding sentence that uses only resources from *plural logic*. Consider again the nine people from the above example, and consider R:

R: The nine people are separated from each other.

Here we have a merely plurally referring term 'the nine people', and thus a merely plural predicate, 'are separated from each other'. Now, the 'from each other' appears to force us to interpret 'are separated from each other' as a distributive plural predicate, such that R's truth implies that no two of the nine people are together. This interpretation of R is clearly not equivalent to either P or Q, both of which are consistent with various subpluralities of the nine people being together. Further, even supposing we ignore the putative distributing effect of 'from each other' and interpret 'are separated' as it appears in R as a *non-distributive* plural predicate, R apparently remains non-equivalent to either P or Q. For then R would say that the nine people are *jointly* separated, presumably from some contextually specified thing or things; at any rate, R on this reading would be consistent with the nine people all being together, whereas both P and Q are inconsistent with this. Finally, suppose we think that the effect of 'from each other' is to say that some of the subpluralities of the nine people are separated from some of the other subpluralities of those nine people. Still, R doesn't seem equivalent to either P or Q, since P and Q seem to stipulate *which* of the subpluralities of the nine people are separated from which others, whereas R does not. In general, then, superplural logic allows us to say things about a plurality of individuals, the truth of which is sensitive to the way in which we conceive of that plurality of individuals as dividing into subpluralities.

Now, despite all of this, I don't take myself to have given an argument for the *indispensability* of superplural logic. For all I've said, the meanings expressed by P and Q might well be expressible in purely *plural logic*. For example, we might paraphrase P and Q in plural logic, replacing the ostensibly one-placed superplural predicate '___ are separated from each other' with the two-placed plural predicate '___ are separated from ___', as follows:

P*: The Germans are separated from the Austrians

Q*: The men are separated from the women, and the men are separated from the children, and the women are separated from the children.

P* and Q* seem to be extensionally equivalent to P and Q respectively. We might then wonder, and indeed argue about, whether there are acceptable paraphrases in plural logic of *every* such sentence that ostensibly requires superplural analysis.⁴⁷ But answering this question is irrelevant to the present project. I want to use superplural logic to give the reductive nihilist's definition of composition, and in order to do so I have tried to clarify what is meant by superplural reference, and superplural predication, etc. It's completely consistent with this that superplural logic is reducible to merely plural logic, in the sense that any sentence that apparently requires superplural analysis is equivalent to a sentence requiring only plural analysis. That would entail only that the definition of composition that I'm going to give in terms of superplural logic is equivalent to a definition of composition in terms of merely plural logic (or perhaps that I'm *really* giving a definition of composition in terms of plural logic, but that I'm doing so using some non-standard notation). This result need trouble the reductive nihilist not a bit.

Now, an important notion of superplural logic is the notion of being what I'll call *superplurally among*. This notion is expressed by the *superplural amonghood* predicate, that is distinct from the ordinary plural amonghood predicate characterised by AMONGHOOD. The superplural amonghood predicate takes a plural term or variable in its first argument place, and a superplural term (written e.g. 'aaa') or variable (written e.g. 'xxx') in its second argument place. Intuitively, *xx* are superplurally among a given superplurality iff *xx* are one of the pluralities that the superplurality is a plurality of. Here are some examples that should help to illustrate:

1. Just two pluralities are superplurally among the countries of Great Britain and the countries on the Irish landmass: the countries of Great Britain (plurality number one), and the countries on the Irish landmass (plurality number two). Importantly, the plurality consisting of e.g. Scotland and Northern Ireland are *not* superplurally among the countries of Great Britain and the countries of the Irish landmass.
2. Just three pluralities are superplurally among the men, the women, and the children: the plurality of men (plurality number one), the plurality of women (plurality number two), and the plurality of children (plurality number three).
3. Just two pluralities are superplurally among the Germans and the Austrians: the Germans (plurality number one) and the Austrians (plurality number two).

Note that 2 and 3 are the case *even if* 'the men, the women, and the children' refer to the same individuals as the ones that 'the Germans and the Austrians' refers to.⁴⁸

In general, suppose we have a plurality *xx*, that we can conceive of as a plurality of pluralities, i.e. a superplurality, *xxx*. Any *yy* that is superplurally among *xxx* is among *xx*, but the converse doesn't hold: just because some *yy* are among *xx*, does not mean that *yy* are superplurally among *xxx*. For example, although the Germans are among the nine people in the room in the example above, it's not

⁴⁷ For an argument that some sentences can *only* be adequately analysed by appeal to superplural logic, see Linnebo and Nicolas (2008).

⁴⁸ Thus the superplural amonghood predicate is hyperintensional. For another illustration of this, consider the natural numbers. We can refer to the natural numbers using, for example, 'the odd natural numbers and even natural numbers', or 'the natural numbers smaller than 10 and the natural numbers greater than or equal to 10'. That is, *the odd natural numbers and even natural numbers* are identical to *the natural numbers smaller than 10 and the natural numbers greater than or equal to 10*. Despite this, whilst it's true to say that e.g. the odd natural numbers are superplurally among *the odd natural numbers and the even natural numbers*, it's false to say that the odd natural numbers are superplurally among *the natural numbers smaller than 10 and the natural numbers greater than or equal to 10*.

the case that the Germans are superplurally among the men, the women, and the children (despite the fact that the men, the women, and the children *just are* the nine people in the room conceived of as a plurality of pluralities).

Two final notes to make about superplural logic, as least as I'm understanding it here. First, recall that a *plural* variable, xx can in the limit case have as a value a single entity, and that in that case the plurality assigned to xx is an *improper* plurality. I'm also going to allow for *improper superpluralities*. A superplurality can be understood as a plurality of pluralities; an *improper* superplurality is a superplurality of *improper* pluralities. For example, consider the superplural term, 'all of the people in my living room and all of the people in the garden'. Suppose that there is just one person in my living room, and just one person in my garden (and that each person is a single entity, i.e. an improper plurality). Then 'all of the people in my living room and all of the people in the garden' is an improper superplural term: it is a plurality of improper pluralities. There can also be superpluralities 'containing' a mixture of proper and improper pluralities (for example, imagine that there are *two* people in my garden, but still only one in my living room).

Finally, I take superplural logic to validate all instances of the *superplural comprehension principle*:

SUPERPLURAL COMPREHENSION: $\exists xx (\Phi xx) \rightarrow \exists xxx \forall xx (xx \text{ are superplurally among } xxx \leftrightarrow \Phi xx)$

Intuitively, SUPERPLURAL COMPREHENSION says that whenever there is at least one plurality, xx , that satisfies a certain condition, Φ , there is a superplurality, xxx , such that all and only pluralities that satisfy Φ are superplurally among xxx .

That short exposition of superplural logic gives us the tools necessary to give the reductive nihilist's definition of composition.

2.4.3.3 – The Superplural Definition of Composition

Here is the reductive nihilist's definition of composition:

COMPOSITION: $xxx \text{ compose } yy =_{df}$

- (i) $\forall xx$ (if xx are superplurally among xxx then xx is/are a part of yy) AND
- (ii) $\forall xx \forall zz$ (if xx are superplurally among $xxx \wedge zz$ are superplurally among $xxx \wedge xx \neq zz$, then xx doesn't overlap zz) AND
- (iii) $\forall xx$ (if xx is/are a part of yy then $\exists zz$ (zz are superplurally among xxx and xx overlaps zz)) AND
- (iv) yy are an object

Note that in the first argument place of COMPOSITION is a superplural variable. That's because reductive nihilists want to be able to say that e.g. the bricks compose the wall, where *each brick is itself a (proper) plurality* (in that 'is a brick' is in each case satisfied by a proper plurality). Thus reductive nihilists need COMPOSITION to be consistent with a plurality of pluralities (i.e. a

superplurality) composing something, and so need the first argument place of COMPOSITION to be a superplural variable.⁴⁹

Intuitively, COMPOSITION says that some pluralities compose *yy* iff

- (i) each of those pluralities are among *yy*, AND
- (ii) no two of those pluralities overlap, AND
- (iii) any *zz* that is among *yy* overlaps at least one of those pluralities, AND
- (iv) *yy* are an object.

Let's see an example to see how COMPOSITION works. Consider a case that we'd intuitively describe as one in which a wall has as its parts ten bricks, none of which overlaps any of the others. Reductive nihilists can redescribe this as a case in which there is a plurality of simples* arranged wall-wise, call them *aa*, that has ten pluralities of simples* arranged brick-wise, call them *bb*₁, ..., *bb*₁₀, as subpluralities. No two of *bb*₁, ..., *bb*₁₀ overlap. We can additionally suppose that there is no *xx* such that *xx* is a subplurality of *aa* but doesn't overlap any of *bb*₁, ..., *bb*₁₀. Given SUPERPLURAL COMPREHENSION, there is a superplurality, *ccc*, such that all and only each of *bb*₁, ..., *bb*₁₀ are superplurally among *ccc*.

Given this, *ccc* and *aa* meet all the conditions of COMPOSITION.

First, (i) is met: each plurality that is superplurally among *ccc*, i.e. each of *bb*₁, ..., *bb*₁₀, is/are a part of *aa*, because each of *bb*₁, ..., *bb*₁₀ is /are by stipulation a subplurality of *aa*.

Second, (ii) is met: of the pluralities that are superplurally among *ccc*, i.e. of *bb*₁, ..., *bb*₁₀, no two of those pluralities overlap, by stipulation.

Third, (iii) is met. We've stipulated that there is no *xx* such that *xx* is a subplurality of *aa* but doesn't overlap any of *bb*₁, ..., *bb*₁₀, which means that there is no *xx* such that *xx* is/are a part of *aa* but doesn't overlap one of the pluralities that is superplurally among *ccc*.

Fourth, (iv) is met: we've stipulated that *aa* are the simples arranged wall-wise, and intuitively any region of space exactly occupied by some simples arranged wall-wise is exactly occupied by an object (namely, a wall). So *aa* are an object.

So *ccc* compose *aa*. That is, the ten pluralities of simples* arranged brick-wise, *bb*₁, ..., *bb*₁₀ compose the plurality of simples arranged wall-wise. That is, given the reductive nihilist's other semantic claims, the bricks compose the wall.

Indeed, COMPOSITION is formulated in such a way that if we ignore condition (iv), the following holds: whenever COMPOSITION* would have implied that some *F*s compose a *G* on the assumption that each of the *F*s and the *G* are single entities, COMPOSITION implies that some *F*s compose a *G* even if we allow that the *F* and the *G* can each be (proper) pluralities. This can be seen by considering

⁴⁹ Since we're here allowing that there are *improper* superpluralities, i.e. a plurality of *improper* pluralities, COMPOSITION is consistent with some individual entities (i.e. some simples*) composing some *yy*. It is also consistent with a single entity composing itself. In that case, the superplural variable *xxx* takes as its assignment an improper superplurality, i.e. a plurality of improper pluralities, and the plurality of improper pluralities it takes as its assignment is an *improper* plurality of improper pluralities, i.e. a single entity. Indeed, given COMPOSITION, every entity trivially composes itself.

the notion of a ‘level’, where a singular variable is a variable of level one, a plural variable is a variable of level two, and a superplural variable is a variable of level three. For, again ignoring (iv), COMPOSITION is just the result of taking COMPOSITION* and shifting it up a level, as it were: that is, COMPOSITION takes every singular variable in COMPOSITION* and swaps it for a plural one, and likewise swaps every plural variable for a superplural one; additionally, where COMPOSITION* uses ‘is one of’ (a predicate that takes a variable of level 1 in its first place and a variable of level 2 in its second place), COMPOSITION uses ‘are superplurally among’ (an analogous predicate that takes a variable of level 2 in its first place and variable of level 3 in its second place); finally, where COMPOSITION* uses ‘parthood*’ and ‘overlap*’ (both of which take variables of level 1 in both of their argument places), COMPOSITION uses ‘parthood’ and ‘overlap’ (both of which taken variables of level 2 in both of their argument places). So it should be no surprise that, ignoring (iv), whenever COMPOSITION* implies that the *F*s compose the *G* on the assumption that what satisfies ‘is an *F*’ and ‘is a *G*’ in each case is a single entity, i.e. the sort of thing that can be the value of a *singular* variable, COMPOSITION implies that the *F*s compose the *G* on the assumption that what satisfies ‘is an *F*’ and ‘is a *G*’ in each case is a plurality, i.e. the sort of thing that can be the value of a *plural* variable.

Why does the reductive nihilist includes (iv) in her definition of composition? In order to rule out cases of *gruesome composition*. Consider the simples* arranged nose-wise, *nn*, and the simples* arranged Eiffel-Tower-wise, *ee*. Let *ddd* be the superplurality such that only *nn* and *ee* are superplurally among *ddd*. Let *ne* be the plurality of simples* made up of *nn* and *ee* (so *ne* are the simples arranged nose-Eiffel-Tower-wise). If COMPOSITION only included (i)-(iii), then *ddd* would compose *en*. For (i) would be met: each of *nn* and *ee* are, by stipulation, among *ne*, and therefore each of *nn* and *ee* is a part of *ne*. (ii) would also be met: *nn* and *ee* do not overlap. And (iii) would also be met: no plurality that doesn’t overlap either *nn* or *ee* is among *ne* (we can suppose), so no plurality that doesn’t overlap either *nn* or *ee* is a part of *ne*. So without (iv), COMPOSITION would imply that there is something, namely *ne*, composed out of a nose and the Eiffel Tower. Requiring that (iv) be met as well, however, allows us to deny this, since *ne* are not an object (because it’s not intuitive that the region of space exactly occupied by the simples* arranged Eiffel-Tower-wise is exactly occupied by an object).⁵⁰

A consequence of COMPOSITION is that *composition implies identity*. My proof of this trades on something I observed about superpluralities in §2.4.2.2: that a superplurality is just a plurality conceived of in a certain way. Metaphysically speaking, then, a superplurality is just (i.e. is *identical* to) a plurality. Intuitively, a given superplurality is identical to the plurality formed by (conceptually) pooling together all of the pluralities that are superplurally among that superplurality. For example,

⁵⁰ Including (iv) in the definition of COMPOSITION does have one perhaps slightly odd result, however. Consider again the nose and the Eiffel Tower. They do not compose anything, given COMPOSITION, but there is a plurality, i.e. *ne*, of which they are both parts. That’s because the nose and the Eiffel Tower – i.e. the simples* arranged nose-wise and the simples* arranged Eiffel-Tower-wise – are both among *ne*, and given PARTHOOD, that means that they are both parts of *ne*. Further, since *ne* has (proper) parts, by COMPOSITE it is composite. If it could be shown that this result is abhorrent, the reductive nihilist could avoid it by building into her definition of PARTHOOD that *xx* can only be a part of *yy* if *yy* are an object (she would then have to similarly alter her other definitions of mereological vocabulary). But I don’t find the result particularly abhorrent, because I think the following describes an intuitively possible state of affairs: the nose and the Eiffel Tower do not compose anything, as there is no object such that that object has the nose and the Eiffel Tower as parts; however, the nose and the Eiffel Tower, taken together, form a plurality that has the nose as one part and the Eiffel Tower as another part; the plurality of the nose and the Eiffel Tower is composite, in that it has parts. That is I don’t think it’s abhorrent to maintain that there can be some *xx* that are not composed by anything, but that nonetheless have (proper) parts and are composite. So I leave the reductive nihilist’s definitions of mereological vocabulary as they are.

consider again the superplurality of the Germans and the Austrians that we considered above. Two pluralities are superplurally among this superplurality: the plurality of the four Germans and the plurality of the five Austrians. If we pool the four Germans and the five Austrians together, we get a plurality of nine people. Thus the superplurality of the Germans and the Austrians is identical to a plurality of nine people.

With this in mind, consider again the arbitrary example of ten bricks composing a wall. That is, given reductive nihilism, consider again the ten pluralities of simples* arranged brick-wise, bb_1, \dots, bb_{10} , and the plurality of simples* arranged wall-wise, aa , and the superplurality, ccc , where the only pluralities that are superplurally among ccc are each of bb_1, \dots, bb_{10} . For reductive nihilists, then, to say that the ten bricks compose the wall is to say that ccc compose aa . That means that ccc and aa satisfy COMPOSITION. Now, ccc is identical to the plurality that results from pooling together the pluralities that are superplurally among ccc . Call this plurality cc . If $cc = aa$, then, since $ccc = cc$, we can conclude that ccc are identical to aa . So I'll show that $cc = aa$. I'll do so via *reductio*. For, given PLURAL IDENTITY, there are two ways for cc to fail to be identical to aa : first, there could be a plurality, xx , such that xx are among aa but not among cc ; second, there could be a plurality, xx , such that xx are among cc but not among aa . Both cases generate contradictions, as I'll now show.

Suppose first there there is a plurality xx such that xx are among aa but not among cc . It follows that some subplurality of xx , call it yy , doesn't overlap cc . Given the transitivity of amonghood, yy are among aa . Now, we said that cc is just the plurality you get by (metaphorically) pooling together bb_1, \dots, bb_{10} . So yy doesn't overlap any of bb_1, \dots, bb_{10} . But given that, by hypothesis, (iii) of COMPOSITION is satisfied, every subplurality of aa (including yy) overlaps at least one of bb_1, \dots, bb_{10} . Contradiction: yy both overlaps at least one of bb_1, \dots, bb_{10} and fails to overlap any of bb_1, \dots, bb_{10} .

Suppose instead that there is a plurality xx such that xx are among cc but not among aa . It follows that some subplurality of xx , call it yy , doesn't overlap aa . Given the transitivity of amonghood, yy are among cc . Now, we said that cc is just the plurality you get by pooling together bb_1, \dots, bb_{10} . So given that yy are among cc , yy overlap at least one of bb_1, \dots, bb_{10} . So some subplurality of yy , call it zz , is among one of bb_1, \dots, bb_{10} . Since zz is a subplurality of yy and yy doesn't overlap aa , zz is not among aa . But given that, by hypothesis, (i) of COMPOSITION is satisfied, each of bb_1, \dots, bb_{10} is a subplurality of aa . By the transitivity of amonghood, each subplurality of each of bb_1, \dots, bb_{10} must also be a subplurality of aa . So zz must be a subplurality of aa ; that is, zz must be among aa . Contradiction: zz is both among aa and not among aa .

Given that both ways for cc to fail to be identical to aa result in contradiction, we can conclude that $cc = aa$. Further, cc are identical to ccc . So $ccc = aa$. Translating, the bricks are identical to the wall that they compose.

This example was arbitrary. So we can conclude that, in general, whenever some parts compose a whole the parts are identical to the whole. In other words, composition implies identity.

Finally, I note briefly that the kind of identity that holds between parts and the whole they compose is not the *many-one* identity that those who endorse CAI say it is: that is, it's not identity between many entities on the one hand, and a single entity on the other (see §6.7.2 for more discussion of this claim of CAI's). Rather, the identity is ostensibly between a superplurality, e.g. ccc , and a plurality, aa . But since we're taking superpluralities to just be pluralities conceived of in a certain way, the identity at

issue here is really the identity between two pluralities. That is, composition entails not many-one identity, but *many-many, plural* identity.

2.5 – Conclusion

The reductive nihilist accepts MERELOGICAL NIHILISM, and the equivalent UNIVERSAL SIMPLEHOOD*. Thus she claims that composition* never occurs (or rather, only occurs insofar as every entity trivially composes* itself), and that every entity is a simple*. More specifically, she accepts MICROPHYSICAL NIHILISM, and accordingly thinks that every entity is either a quark or a lepton. However, she denies that these metaphysical claims commit her to ELIMINATIVISM, the claim that there are no composite objects; nor does she think that these metaphysical claims commit her to denying the existence of e.g. tables, chairs, mountains, molecules, human beings, honey bees, etc. For she rejects the semantic claims that microphysical nihilists usually accept, which together result in *eliminative* nihilism. These eliminative nihilism semantic claims are:

- (EN1) ‘Existence’ and its cognates express the singular existential quantifier, $\exists x$.
- (EN2) ‘Object’ is synonymous with ‘entity’, as I characterised it in §2.2.1.
- (EN3) Mereological terms such as ‘parthood’ and ‘composition’ are adequately defined by the definitions I gave in §2.3. Thus ‘parthood’ expresses parthood*, ‘composition’ expresses composition*, etc.
- (EN4) Sortal predicates of the form ‘is an *F*’ are satisfied by a single entity, if at all.
- (EN5) Predicates that putatively ascribe properties to single things are satisfiable only by single entities.
- (EN6) Putatively singular referring expressions each refer to a single entity, if they refer at all.

The reductive nihilist instead accepts the following semantic claims:

- (RN1) ‘Existence’ and its cognates express the plural existential quantifier, $\exists xx$.
- (RN2) ‘Object’ is not synonymous with ‘entity’, but is rather a plural predicate that can be satisfied by proper pluralities of entities. Roughly, xx satisfy ‘object’ iff it is intuitive that the region of space exactly occupied by xx is exactly occupied by an object.
- (RN3) Mereological terms such as ‘parthood’ and ‘composition’ are defined by PARTHOOD, COMPOSITION, etc.
- (RN4) Sortal predicates that purport to apply to single composite *F*s express properties that are instantiated by xx iff xx are both arranged *F*-wise and are an object.
- (RN5) Predicates that putatively ascribe properties to single things express plural properties that are satisfiable by proper pluralities (thus, for example, ‘weighs 10kg’ express *(jointly) weighing 10kg*).
- (RN6) Referring expressions of the form ‘*N*’ that purport to refer to a single composite *F* refer plurally to xx iff xx are arranged *N*-wise and xx are an object.

As I have detailed, accepting (RN1)-(RN6) allows the reductive nihilist to avoid ELIMINATIVISM and to accept the existence of the objects of common sense, including tables, chairs, and the rest (whilst also denying the existence of gruesome composites such as the nose-Eiffel-Tower).

That concludes my introduction to reductive nihilism.

Chapter 3

Reductive Nihilism and the Ramsey-Lewis Treatment of Theoretical Terms

3.1 – Introduction

In this chapter I'll argue that the semantic claims of reductive nihilism can be motivated by appeal to a modified version of the *Ramsey-Lewis treatment of theoretical terms*, which I'll often refer to for short as 'the Ramsey-Lewis treatment'. More carefully, I'll argue that on the assumption that the metaphysical claim of microphysical nihilism (given by MICROPHYSICAL NIHILISM – see §2.3) is true, the semantic claims of reductive nihilism that I introduced in chapter 2 follow from a suitably modified version of the Ramsey-Lewis treatment.

Before getting started, I want to emphasise that I don't attempt here to *defend* the Ramsey-Lewis treatment. Rather, since the Ramsey-Lewis treatment is an established and influential semantic theory, my strategy here is to simply assume that it is a good semantic theory, and show that the semantic claims of reductive nihilism are well-justified, given that semantic theory. That doesn't mean that I won't countenance any criticisms of the Ramsey-Lewis treatment: indeed, a key move in this chapter is to argue that the canonical version of the Ramsey-Lewis treatment is inadequate, and must be modified. But what it does mean is that I won't consider here the possibility that the Ramsey-Lewis treatment is fundamentally mistaken (this point becomes relevant to this chapter's argument in §3.4.3).

The plan for this chapter, then, is as follows. First (§3.2) I'll introduce and explain the Ramsey-Lewis treatment of theoretical terms, before arguing (§3.3) on grounds completely independent of consideration of nihilism that the canonical formulation of the Ramsey-Lewis treatment ought to be modified in certain ways. Then (§3.4) I'll argue that the semantic claims of reductive nihilism follow from the appropriately modified version of the Ramsey-Lewis treatment and the metaphysical claim of microphysical nihilism. In §3.5 I consider an important objection; §3.6 concludes.

3.2 – The Ramsey-Lewis Treatment of Theoretical Terms

The Ramsey-Lewis treatment of theoretical terms is an account of how theoretical terms, i.e. terms introduced without explicit definition by a new theory, come to have meaning. It was developed and popularised by Lewis (esp. 1970, 1972: 250-5), building on the work of Ramsey ([1929]: 1978) and Carnap (1958). According to this account, theoretical terms, or 'T-terms', are implicitly given definitions by the theory that introduces them, in that the theory in question implicitly specifies some roles and states that the T-terms are to name the occupants of those roles. By finding out (often by empirical means) what the occupants of those roles are, we can find out what those T-terms refer to.

The most obvious illustrative example of the Ramsey-Lewis treatment in action is that of a detective theory (inspired by Lewis 1972: 250-3). Suppose a detective investigating a murder knows *how* a murder was committed, including knowing that the murderer had an accomplice, but doesn't yet know who committed the murder nor who their accomplice was. Her theory might say, 'X hid in the bushes in the garden until midnight. At that point she crept out, picked the lock to the flimsy back door with her screwdriver, picked up the revolver that Y had earlier hidden in the kitchen cupboard, ...' etc. Though this theory doesn't explicitly define 'X' or 'Y', we understand precisely what the detective means by these T-terms. That's because our detective's theorising *implicitly* defined these T-terms. For having heard our detective's theory, we know that 'X' means something like *whoever hid in the*

bushes in the garden until midnight, and who picked the lock to flimsy back door with her screwdriver, etc. We also know that ‘Y’ means *whoever hid the revolver that X later retrieved from the kitchen, etc.* That is (and paraphrasing Lewis 1972: 251), in explaining her theory, our detective set forth two roles and said that they were occupied by X and Y; in doing so, she must have specified the meanings of the T-terms ‘X’ and ‘Y’, for they had meanings afterwards, they had none before, and nothing else was done to give them meanings. With this in mind, the detective is now in a position to go about trying to find the referents of ‘X’ and ‘Y’ – that is to say, the detective is now in a position to try to discover (by empirical means) who committed the murder and who their accomplice was. To discover who the murderer was, she just needs to find out who it was who hid in the bushes in the garden until midnight, and who picked the lock to the flimsy back door with her screwdriver, etc.; to discover who the accomplice was, she just needs to find out who it was that hid a revolver in the kitchen for the murderer to later retrieve.

An instructive feature of this example is the way in which ‘X’ and ‘Y’ are implicitly defined by our detective’s theorising *partially in terms of each other*. For example, ‘Y’ is defined as referring to whoever hid the revolver *that X later retrieved* from the kitchen. As a result, when we come to use the definitions of ‘X’ and ‘Y’ that our detective’s theory has implicitly supplied us with to try to find the referents of ‘X’ and ‘Y’, it’ll be no good to try to find the referents of these term *one by one*. That would leave us stuck: in order to know whether a certain person is the referent of ‘X’, we’d need to know whether that person retrieved the revolver that Y had hidden in the kitchen, which we can’t do without knowing what ‘Y’ refers to; but similarly in order to know whether a certain person is the referent of ‘Y’, we’d need to know whether that person hid the revolver that X later retrieved from the kitchen, which we can’t do without knowing what ‘X’ refers to.

So what we need to try to do is to find the referents of ‘X’ and ‘Y’ *at the same time*. That is, we need to try to find the ordered pair such that the *first* member of that ordered pair hid in the bushes, picked the lock to the flimsy back door, and retrieved the revolver that the *second* member of the ordered pair had hidden in the kitchen cupboard. Equivalently: we need to try to find the ordered pair such that, on the assumptions that ‘X’ refers to its first member and ‘Y’ refers to its second member, the detective’s theory is true. If we can find this ordered pair, then we can conclude that ‘X’ really does refer to the first member of that ordered pair, and that ‘Y’ really does refer to the second member of the ordered pair.

Now, it’s possible that we won’t be able to find the ordered pair that we’re looking for. Setting aside the possibility that we’re simply not very good detectives, this could be for one of two reasons. First, and most obviously, it might be that no-one did the things that the detective’s theory says were done by X and Y. That is, it might be that there is simply no ordered pair such that, on the assumptions that ‘X’ refers to its first member and ‘Y’ refers to its second member, the detective’s theory is true. If this is the case, then our detective was mistaken as to how the murder was committed. Then we should conclude that ‘X’ and ‘Y’ as they appear in the detective’s theory don’t refer to anything: they are like (or perhaps simply are) definite descriptions that are not satisfied by anything. Second, and less obviously, we might not be able to find *the* ordered pair we’re looking for because it might be that multiple ordered pairs are such that, on the assumptions that ‘X’ refers to their first member and ‘Y’ refers to their second member, the detective’s theory is true. If two pairs of people did everything that the detective’s theory said that X and Y did, then what the detective has said underdetermines which of those pairs of people ‘X’ and ‘Y’ refers to. Then we should again conclude that ‘X’ and ‘Y’ don’t

refer to anything: they are like (or perhaps simply are) definite descriptions that fail to be uniquely satisfied.⁵¹

So much for the intuitive gloss. I turn now to presenting the canonical formulation of the Ramsey-Lewis treatment of theoretical terms (as presented in Lewis 1972: 253-5; see also Lewis 1970), which purports to formalise the way in which theories implicitly define their T-terms.

Let T be a given theory, and let $[t_1, \dots, t_n]$ be the n theoretical terms that that theory introduces. Among $[t_1, \dots, t_n]$ will be nominalised versions of any new predicates, functors, etc. that T would introduce, thus ensuring that the Ramsey-Lewis treatment will provide definitions of every new bit of language that T introduces, regardless of its syntactic category (Lewis 1970: 429). We can then convert T into a single long conjunctive sentence, called the postulate of T, which we write like this:

$$(1) T[t_1, \dots, t_n]$$

to make clear that T uses $[t_1, \dots, t_n]$. The postulate says of the entities named by the T-terms it uses that they have certain properties, do certain things, and stand in certain relations to each other and to the entities named by the *O*-terms (which are *other* or *old* terms, i.e. terms that we already understand the meaning of and which aren't being introduced by the theory) used in the postulate. That is, it says of the referents of $[t_1, \dots, t_n]$ that they occupy certain corresponding roles $[R_1, \dots, R_n]$.⁵²

The next step is to replace each T-term of this postulate with a free variable, to get the *realisation* formula of T:

$$(2) T[x_1, \dots, x_n]$$

Any n -tuple of entities that satisfies this realisation formula, i.e. any n -tuple of entities whose members occupy the roles that T said the referents of the members of $[t_1, \dots, t_n]$ occupy, *realises* T. If we bind the variables with existential quantifiers, we get the Ramsey sentence of T:

$$(3) \exists x_1, \dots, \exists x_n T[x_1, \dots, x_n]$$

The Ramsey sentence of T tells us that T is realised by at least one n -tuple of entities. That is, at least one n -tuple is such that the entities it contains occupy the roles that T said that the referents of $[t_1, \dots, t_n]$ occupy. We're apparently licensed to infer the Ramsey sentence from T because T wouldn't be true if nothing realised it. For example, if a detective's theory alleges that X picked the lock to the kitchen door, but no-one picked the lock to the kitchen door, then our detective's theory would be false.

T would also fail to be true if more than one n -tuple were to realise it (see Lewis 1970: 432-5; 1972: 252-3). After all, T purports to describe the way the world is, and if it's neutral between more than one different ways the world might be, as it would be if it were multiply realised, then it has failed to

⁵¹ Note that even if it turns out that 'X' and 'Y' don't refer to anything, our detective's theorising has still given us definitions of these T-terms (otherwise we wouldn't understand the terms well enough to work out that they don't have referents).

⁵² Sometimes the roles that theories associate their *T*-terms with in this way are said to be purely causal or functional in nature (e.g. Lewis 1966, 1970, 1972), but there doesn't seem to be any principled reason for excluding roles that are characterised at least partially non-causally (e.g. Menzies 1996: 101 fn 25, Menzies and Price 2009: 183, Jackson 2000, Tooley 1987). A detective who theorises that 'X picked the lock with her screwdriver', for example, doesn't just thereby imply that X did certain causal work, i.e. picking a lock with a screwdriver: she also implies through her use of '*her* screwdriver' that X is a woman.

describe *the* way the world is. So we can apparently infer from T the *modified* Ramsey sentence, which says that T is *uniquely realised*:

$$(4) \exists!x_1, \dots \exists!x_n T[x_1, \dots, x_n]$$

The modified plural Ramsey sentence tells us that each role R_i associated with each t_i is uniquely occupied.

We've said in (1) that the entity that each t_i of $[t_1, \dots, t_n]$ refers to occupies a role R_i , and in (4) that each R_i is uniquely occupied. It follows that each t_i refers to the unique occupant(s) of each corresponding R_i . We can express this with the modified Carnap sentence, which says that if T is uniquely realised then it is realised by the entities referred to by $[t_1, \dots, t_n]$:

$$(5) \exists!x_1, \dots \exists!x_n T[x_1, \dots, x_n] \rightarrow T[t_1, \dots, t_n]$$

And we can complete our account of the meaning of $[t_1, \dots, t_n]$ by claiming that if it's *not* the case that T is uniquely realised (because it's either multiply or not at all realised) then $[t_1, \dots, t_n]$ don't refer.

3.3 – The Pluralised Ramsey-Lewis Treatment of Theoretical Terms

The above formulation of the Ramsey-Lewis treatment of theoretical terms is based on Lewis's presentation. It is canonical but, I think, mistaken in an important respect. In this section I explain Lewis' mistake, and fix it. For reasons that will become apparent, I'll call the version of the Ramsey-Lewis treatment that results from fixing this mistake the *pluralised Ramsey-Lewis treatment of theoretical terms*.

Lewis's mistake stems from an illicit assumption that motivates the step from (1) to (2). Here again is (1):

$$(1) T[t_1, \dots, t_n]$$

$T[t_1, \dots, t_n]$ is, recall, the postulate of T that makes clear that T uses theoretical terms t_1, \dots, t_n . The next step in the Ramsey-Lewis treatment is to replace each T-term with a free variable to get the realisation formula of T, which Lewis writes like this:

$$(2) T[x_1, \dots, x_n]$$

This is where Lewis' illicit assumption makes itself known. By using *singular* free variables, $[x_1, \dots, x_n]$, the value of each of which can by definition only be a *single* entity, Lewis assumes here that the reference of each t_i is a single entity; he assumes that $[t_1, \dots, t_n]$ are singularly referring. Nothing warranted this assumption. At this stage of the treatment, all we know is that our theory T has specified some roles, $[R_1, \dots, R_n]$, and has implied that each R_i is occupied by what the corresponding t_i refers to. But *plural reference*, whereby a single referring expression refers to many entities at once, is widely agreed to be possible.⁵³ Furthermore, there's no reason for thinking that roles can't in principle be *jointly* occupied by many entities. After all, what it is to occupy a role is simply to satisfy a certain description (for example, what it is to occupy the role associated with 'Y' above is to satisfy the description 'hid the revolver that Y later retrieved from the kitchen cupboard', and descriptions can be satisfied by many entities collectively (consider e.g. the description 'the people who surrounded the building') (see McKay 2006). It follows that there's no reason, semantic or metaphysical, to exclude

⁵³ See chapter §2.2.2. See also McKay (2006: esp. ch. 2), Oliver and Smiley (2013: esp. chs 3-4), and Yi (2005, 2006).

at this stage the possibility that some or all of $[t_1, \dots t_n]$ plurally refer to many entities that jointly occupy the relevant role.

So, I submit, we should use *plural* variables to represent the referents of $[t_1, \dots t_n]$. The value of a plural variable can be either a single entity *or* many entities, so by using plural variables we remain neutral as to whether $[t_1, \dots t_n]$ are plurally referring or not, and whether the roles associated with $[t_1, \dots t_n]$ are occupied by single entities or by many entities collectively. Thus we should replace Lewis's (2) with the *plural* realisation formula:

$$(2^*) \quad T[xx_1, \dots xx_n]$$

Swapping (2) for (2*) has knock-on effects for the rest of the treatment. The next step of Lewis' account is (3), i.e. the Ramsey sentence of T, which binds the free variables from (2) with the singular existential quantifier in order to express the claim that T is realised. We'll instead need to bind the plural variables of (2*) with the plural existential quantifier to arrive at the *plural* Ramsey sentence:

$$(3^*) \quad \exists xx_1, \dots \exists xx_n T[xx_1, \dots xx_n]$$

For the same reason that Lewis reasoned we were allowed to infer (3) from T, we're licensed to infer (3*) from T: T wouldn't be true if nothing realised it, i.e. if nothing occupied the roles that T said the referents of $[t_1, \dots t_n]$ occupied.

We should also now use the modified *plural* Ramsey sentence to express the claim that T is uniquely realised:

$$(4^*) \quad \exists!xx_1, \dots \exists!xx_n T[xx_1, \dots xx_n]$$

The modified plural Ramsey sentence tells us that each role R_i associated with each t_i is uniquely occupied either by some single entity, or by many entities collectively. As above, we can infer this sentence from T because T purports to describe *the* way the world is (again, see Lewis 1970: 432-5; 1972: 252-3).

Now, again as above, we've said in (1) that that the entity/entities that each t_i of $[t_1, \dots t_n]$ refers to occupies/occupy a role R_i , and in (4*) that each R_i is uniquely occupied. It follows that each t_i refers to the unique occupant(s) of each corresponding R_i . We can express this with the modified plural Carnap sentence, which says that if T is uniquely realised then it is realised by the entities referred to by $[t_1, \dots t_n]$:

$$(5^*) \quad \exists!xx_1, \dots \exists!xx_n T[xx_1, \dots xx_n] \rightarrow T[t_1, \dots t_n]$$

And, again mirroring the Lewisian formulation, we can complete our account of the meaning of $[t_1, \dots t_n]$ by claiming that if it's *not* the case that T is uniquely realised (because it's either multiply or not at all realised) then $[t_1, \dots t_n]$ don't refer.

Before moving on, let me note that it'd be a mistake to object to the pluralised version of the Ramsey-Lewis treatment of theoretical terms on the grounds that if it did turn out that the role associated with some t_i were uniquely occupied by many entities collectively then we should say that t_i failed to refer, not that t_i refers plurally to the many entities that occupy that role. This objection is mistaken for two reasons. First, there can be grammatically plural T-terms, and it'd be odd to say that they would fail to refer if the roles with which they're associated were occupied by multiple entities collectively. For example, if three people were to jointly uniquely occupy the role that a detective's theory associates

with the T-term ‘the Manchester muggers’, then we should say that ‘the Manchester muggers’ refers to those three people, not that it fails to refer. Second, nothing in the formulation of the pluralised Ramsey-Lewis treatment rules out that a given theory T might imply that some or all of its T-terms refer singularly if at all. For example, you might think that if a theory uses a grammatically singular form of a T-term then it thereby implies that one condition on occupying the role associated with that T-term is to be a single entity, in which case that T-term cannot refer plurally. I happen *not* to think this, as will become clear, but the point is that the pluralised Ramsey-Lewis treatment itself is perfectly compatible with the idea that some or all T-terms must refer singularly if at all, so the objection that it can’t accommodate this idea is misplaced.

3.4 – Applying the Treatment

Suppose microphysical nihilism is true. That is, assume the truth of MICROPHYSICAL NIHILISM:

MICROPHYSICAL NIHILISM: $\forall x (x \text{ is simple}^*) \wedge \forall x (x \text{ is a simple}^* \text{ iff } x \text{ is either a quark or a lepton})$

Given this assumption, I think that the pluralised Ramsey-Lewis treatment of theoretical terms supports the semantic claims of reductive nihilism.

3.4.1 – Preliminaries

The first step in my argument for this is to reiterate and emphasise that we can use the pluralised Ramsey-Lewis treatment of theoretical terms to get definitions of not only new referring expressions, but also new *predicates* that a theory introduces (as well as lexical items of other grammatical categories) (Lewis: 1970: 429). For example, suppose we introduce a new scientific theory that says ‘two atoms are covalently bonded to one another whenever they share an electron’, and suppose that ‘covalently bonded’ is a new bit of vocabulary introduced by our theory. In order to use the (pluralised) Ramsey-Lewis treatment to discover the definitions of new bits of vocabulary that are introduced by a theory but that are not nouns, we simply rewrite our theory in such a way that every occurrence of the new bit of vocabulary is replaced by an equivalent phrase in which that bit of vocabulary is nominalised. In our example, then, we can simply rewrite our theory so that it says, ‘two atoms *bear the relation of being covalently bonded* to each other whenever they share an electron’. Here ‘the relation of being covalently bonded’ is a noun phrase that is apt for being treated as a T-term and subjected to the (pluralised) Ramsey-Lewis theory in the usual way.

The next thing to point out is that theories (including term-introducing theories) need not be scientific or detective theories: there can be *folk* theories as well. This is a crucial insight, due to Lewis (1972: 253-6). A folk theory of X is the disjunction of conjunctions of most of the claims concerning X that the folk invariably take to be true (Lewis 1972: 256).⁵⁴ Plausibly, when we’re concerned with what a certain term *of ordinary language* means, it makes sense to find out how that term features in the *ordinary, folk* theory of the relevant phenomenon and to say that the term is defined as naming whatever uniquely occupies the role associated with it by that folk theory (Jackson 2000: 30-8). In this

⁵⁴ Why not treat the folk theory as simply being the *conjunction* of claims that the folk would take to be true? In order to allow for the folk to be wrong, to a limited extent, whilst still using their terms meaningfully (and referentially).

way, we can use folk theorising to discover the definitions and referents of ordinary terms of philosophical interest.⁵⁵

Now consider again the semantic claims of reductive nihilism:

- (RN1) ‘Existence’ and its cognates express the plural existential quantifier, $\exists xx$.
- (RN2) ‘Object’ is not synonymous with ‘entity’, but is rather a plural predicate that can be satisfied by proper pluralities of entities. Roughly, xx satisfy ‘object’ iff it is intuitive that the region of space exactly occupied by xx is exactly occupied by an object.
- (RN3) Mereological terms such as ‘parthood’ and ‘composition’ are defined by PARTHOOD, COMPOSITION, etc.
- (RN4) Sortal predicates that purport to apply to single composite F s express properties that are instantiated by xx iff xx are both arranged F -wise and are an object.
- (RN5) Predicates that putatively ascribe properties to single things express plural properties that are satisfiable by proper pluralities (thus, for example, ‘weighs 10kg’ express *(jointly) weighing 10kg*).
- (RN6) Referring expressions of the form ‘N’ that purport to refer to a single composite F refer plurally to xx iff xx are arranged N -wise and xx are an object.

All of these semantic claims concern ordinary, everyday bits of vocabulary.⁵⁶ Thus, as above, we should think that these bits of vocabulary receive their definitions from their place in the *folk* theory in which they feature. The relevant folk theory here is the folk theory of composite objects. Roughly speaking, this is the conjunction of the claims concerning composite objects that the folk invariably take to be true.⁵⁷ Presumably it includes such claims as ‘the Eiffel Tower is located in Paris’, ‘Obama’s left leg is a part of his body’, and so on. As above, we can rewrite the folk theory so that every occurrence of a predicate we want the definition of is replaced with an equivalent phrase in which that predicate is nominalised, and thus apt for being treated as a T-term (so our rewritten theory would contain, for example, ‘Obama’s left leg bears the *parthood* relation to his body’). Thus the T-terms of the folk theory of composite objects are terms that purport to refer to composite objects and (nominalised versions of) predicates that purport to ascribe properties and relations to those objects.

Now, let T be the folk theory of composite objects, which we write as

$$T[t_1, \dots, t_n]$$

⁵⁵ Indeed, philosophers have defended accounts of a variety of phenomena – including causation (Tooley 1987; Menzies 1996), colour (Johnston 1992; Lewis 1997), mental states (Lewis 1972), and moral properties (Jackson and Pettit 1995; Jackson 2000) – by arguing that the term for the phenomenon of interest is associated by the folk theory with a role that is uniquely occupied by the feature or entity specified by their account. Arguing in this way is to argue broadly in the tradition of the so-called *Canberra Plan* (on which see the essays collected in Braddon-Mitchell and Nola [2009]).

⁵⁶ To be sure, these terms are also used by philosophers. But I take it what when philosophers use these terms they understand themselves to be using ordinary, everyday terms (though perhaps with added stipulations, such as that ‘parthood’ is to be reflexive).

⁵⁷ As above, strictly speaking we should understand it as the disjunction of conjunctions of *most* of the claims concerning material objects that the folk take to be true (see Lewis 1972: 256).

where each t_i is a T-term introduced by the folk theory. The pluralised Ramsey-Lewis treatment tells us to replace each new T-term a given theory contains with a variable of a new type, which gives us the realisation formula of the folk theory of composite objects:

$$T[xx_1, \dots, xx_n]$$

Then the pluralised Ramsey-Lewis treatment then tells us to look for an n -tuple of entities that uniquely realises the folk theory; more accurately, the pluralised Ramsey-Lewis treatment tells us to look for an n -tuple of *pluralities* of entities that uniquely realises the folk theory. So we're looking for an n -tuple of the following form:

$$[ee_1, \dots, ee_n]$$

What it is for an n -tuple to realise the folk theory of composite objects is for that n -tuple to be such that its members, ee_1, \dots, ee_n occupy the roles associated by the folk theory with its T-terms t_1, \dots, t_n . What it is for an n -tuple to *uniquely* realise the folk theory is for it to be the *only* such n -tuple that realises the folk theory.

I'm going to argue that, on the assumption of microphysical nihilism, there is an n -tuple that uniquely realises the folk theory of composite objects. Call this n -tuple REALISER. REALISER's members are simply those (pluralities of) entities and properties that the reductive nihilist – via (RN1)-(RN6) – says are the referents of the T-terms of the folk theory of composite objects. So REALISER contains pluralities of simples* arranged F -wise, plural properties of the form [*collectively*] are F , the parthood relation (as defined by PARTHOOD), and so on.

In the next section, I'll argue for the claim that REALISER realises the folk theory of composite objects; in the section after that I'll argue that REALISER *uniquely* realises the folk theory.

3.4.2 – REALISER Realises the Folk Theory of Composite Objects

In the previous section I said that what it is for an n -tuple to realise the folk theory of composite objects is for that n -tuple to be such that its members, ee_1, \dots, ee_n occupy the roles associated by the folk theory with its T-terms t_1, \dots, t_n . This is equivalent to: what it is for an n -tuple to realise the folk theory of composite objects is for the folk theory's realisation formula to be true relative to the variable assignment that assigns to each xx_i of the folk theory's realisation formula the corresponding ee_i of the n -tuple. And, more simply, this is in turn equivalent to: what it is for an n -tuple to realise the folk theory of composite objects is the folk theory to be true relative to the assumption that each of its T-terms, t_i , refers to the corresponding ee_i of that n -tuple.

So in order to argue that REALISER realises the folk theory of composite objects, I'll argue that the folk theory is true relative to the assumption that each t_i of the folk theory refers to the corresponding ee_i of REALISER.

Now, I just characterised REALISER as simply containing the entities and properties that (RN1)-(RN6) nominate as the referents of the T-terms of the folk theory. So, we can check whether REALISER realises the folk theory of composite objects by assuming (RN1)-(RN6) and seeing whether the folk theory of composite objects is true or not, given that assumption.

I think it's very plausible that, given the assumption of semantic claims (RN1)-(RN6), the folk theory of composite objects is true. I'll argue for this by considering a representative sample of claims that the folk theory of composite objects seems to make, and demonstrating that they're true given the

assumption of (RN1)-(RN6) (and the assumption of MICROPHYSICAL NIHILISM). My argument will be necessarily inexhaustive, and I can't consider and rebut every potential counterexample and objection here. Nevertheless, the sample of claims I make here should make it generally speaking plausible that the folk theory of composite objects is true on the assumption of (RN1)-(RN6).⁵⁸

First, then, consider claims that the folk theory makes involving the predicate 'is an object'. Most of these claims, I think, will concern the *extension* of objecthood. For example, it's plausible that the folk theory says things like, 'pens are objects', and 'tables are objects', and 'trophies are objects', etc. These sentences are true, given (RN1)-(RN6). Recall that (RN4) tells us that sortal predicates that purport to refer to single composite *F*s such as 'is a pen', 'is a table', 'is a trophy', etc. are satisfied by *xx* iff *xx* are appropriately arranged (i.e. arranged pen-wise, table-wise, trophy-wise, etc., respectively) and *xx* are an object. Now, since pluralities of simples* arranged pen-wise, table-wise, trophy-wise, etc. do exactly occupy regions of space that are each intuitively exactly occupied by objects, those pluralities each count as an object (given (RN2)). Thus the simples* arranged pen-wise satisfy 'is an object', and therefore also satisfy 'is a pen'; the simples* arranged table-wise satisfy 'is an object', and therefore also satisfy 'is a table'; the simples* arranged trophy-wise satisfy 'is an object', and therefore also satisfy 'is a trophy', etc. Thus pens are objects, tables are objects, and trophies are objects, etc.

In general, since (RN2) tells us that the *xx* that exactly occupy a region of space are in the extension of 'object' iff it is intuitive that that region of space is exactly occupied by an object, the very fact that the folk theory (which is constituted by folk intuitions) says that an object is present in a given region of space guarantees that the simples* that in fact (given microphysical nihilism) jointly exactly occupy that region of space jointly count as an object. Thus whenever the folk theory says that an object is present in a given region of space it is saying something true, given (RN2).⁵⁹

It's also plausible that the folk theory includes claims about what *doesn't* count as an object, e.g. 'Clinton's nose and the Eiffel Tower do not jointly make up or count as an object'.⁶⁰ Given (RN1)-(RN6), sentences such as this come out true as well. First, note that, given (RN6), 'Clinton's nose' refers to the simples* arranged Clinton's-nose-wise and 'the Eiffel Tower' refers to the simples*

⁵⁸ My argument in this section dovetails nicely with my argument in chapter 6. In chapter 6 I argue that reductive nihilism vindicates all of our intuitions concerning composite objects, and to the extent that the folk theory of composite objects is just (roughly speaking) the conjunction of all of those intuitions, arguing that the folk theory of composite objects is true on the assumptions of MICROPHYSICAL NIHILISM and (RN1)-(RN6) is the same as arguing that reductive nihilism vindicates our intuitions concerning composite objects. Thus the more extended discussion in chapter 6 helps to further bolster the claims I make in this section.

⁵⁹ To be clear, the claim here isn't that the folk never make false claims about regions of space are exactly occupied by objects and which aren't. For example, it's possible that there be a region of space, *r*, that is occupied by nothing but fresh air, but (due to a widespread and persistent optical illusion) the folk invariably take to be occupied by a chair. Rather, the claim is, roughly speaking: assuming that the folk are right about the relevant non-object-y facts (which include enough information to preclude the folk being subject to widespread illusions, hallucinations, etc.), their intuitions about which regions of space are exactly occupied by objects match up exactly with which regions of space are exactly occupied by pluralities of simples* that satisfy 'object'.

⁶⁰ You might think that the folk theory of material objects wouldn't contain this sentence, as most people have never considered whether Clinton's nose and the Eiffel Tower together qualify as an object or not. But I think we should take folk theories to include claims that the folk have never actually considered but *would* invariably take to be true. For such hypothetical folk judgements plausibly reveal just as much about how the folk define the terms they use as their actual judgements (see e.g. Jackson [2000: 31-2], who cites folk judgements about e.g. Gettier cases as part of the folk theory of knowledge: plausibly, most people have never considered Gettier cases, but nonetheless it's plausible that the judgements they *would* make in such cases reveal something about the folk theory of knowledge).

arranged Eiffel-Tower-wise. Now consider the region of space occupied by the plurality of simples* made up of the plurality of simples* arranged Clinton's-nose-wise and the plurality of simples* arranged Eiffel-Tower-wise (that is, consider the plurality of simples* arranged nose-Eiffel-Tower-wise). Call this plurality *ne*. Since we're assuming microphysical nihilism, *ne* is the only occupant of that region of space (i.e. there are no composite* entities there too). For the sentence 'Clinton's nose and the Eiffel Tower do not jointly make up or count as an object' to be true, it has to be true that *ne* are not jointly an object. And this is true: it is not intuitive that the region of space exactly occupied by *ne* is exactly occupied by an object, so, given (RN2), *ne* is not in the extension of 'object'.

In general, since (RN2) tells us that the *xx* that exactly occupy a region of space are in the extension of 'object' iff it is intuitive that that region of space is exactly occupied by an object, whenever the folk theory says that an object is not present in a given region of space it is saying something true, given (RN2).

It is also somewhat plausible that the folk theory includes other sorts of claims relevant to determining the folk definition of 'object', such as 'objects are typically internally well-bonded', 'objects typically have a function' (cf. Rose and Schaffer 2007), 'there may exist objects that we haven't discovered yet, and that don't fall under any extant sortal' (cf. Korman 2019: 241), and so on. But since, as we've already seen, the folk theory's first-order judgements about which regions of space are filled by objects and which aren't come out true, given (RN1)-(RN6), it's safe to suppose that these more general observations about the typical features of objecthood also come out true, given (RN1)-(RN6).

So it seems safe to conclude that, given microphysical nihilism and (RN1)-(RN6), the folk theory's claims involving 'object' are true.

Second, consider claims that the folk theory makes involving *names* of composite objects, such as 'the Eiffel Tower'. For example, the folk theory presumably includes claims such as 'the Eiffel Tower is located in Paris', 'the Eiffel Tower is 324m tall', 'the Eiffel Tower was designed by Gustave Eiffel', etc. (or at least, the folk theory contains claims that imply these sentences, such as 'the Eiffel Tower is as tall as the experts say it is'). Given (RN6), 'the Eiffel Tower' refers to the simples* arranged Eiffel-Tower-wise, 'Gustave Eiffel' refers to the simples* arranged Gustave-Eiffel-wise, and 'Paris' refers to the simples* arranged Paris-wise; given (RN5), 'is 324m tall', 'was designed by', and 'is located in' all express corresponding plural properties and relations such as [*collectively*] *being 324m tall*, etc. Given this, the folk theory's claims about the Eiffel Tower are equivalent to 'the simples* arranged Eiffel-Tower-wise are collectively located in the simples* arranged Paris-wise', and 'the simples* arranged Eiffel-Tower-wise are collectively 324m tall', etc. It seems clear that these sentences are all true, given microphysical nihilism.

Third, consider claims the folk theory makes involving (nominalised versions of) sortal predicates, such as 'the property of being a tower'. Presumably the folk theory says things like, 'the Eiffel Tower instantiates the property of being a tower', 'the Leaning Tower of Pisa instantiates the property of being a tower', and 'Big Ben instantiates the property of being a tower'. It plausibly also says things like 'whatever instantiates the property of being a tower is taller than it is wide'. Given (RN1)-(RN6), these sentences are true. For example, 'the Eiffel Tower instantiates the property of being a tower' is true: given (RN6), 'the Eiffel Tower' refers to the simples* arranged Eiffel-Tower-wise; given (RN4), 'the property of being a tower' refers to property – call it the property of [*collectively*] *being a tower* – that is instantiated by any *xx* that instantiate *being arranged tower-wise* and *being an object*; the simples* arranged Eiffel-Tower-wise do instantiate *being arranged tower-wise* and *being an object*, and so they do instantiate the property of [*collectively*] *being a tower*; so 'the Eiffel Tower

instantiates the property of being a tower' is true. Further, for example, 'whatever instantiates the property of being a tower is taller than it is wide' is true: given (RN4) and (RN5), this sentence is materially equivalent to 'any *xx* that instantiate the properties of *being arranged tower-wise* and *being an object* are [collectively] taller than they are [collectively] wide', which is a sentence that microphysical nihilists can and should accept as true.

Another important class of claims the folk makes using sortal predicates are claims of *non-existence*. These divide into two categories. The first category contains claims such as 'unicorns don't exist', and 'dragons don't exist': these sentences are true, given (RN1)-(RN6), because there are no simples* arranged unicorn-wise, or dragon-wise, etc. The second category, roughly speaking, ascribes non-existence to putative gruesome objects composed of parts that do exist. Examples in this second category include 'the nose-Eiffel-Tower doesn't exist'. These sorts of sentences are also true, given (RN1)-(RN6), because although there are e.g. simples* arranged nose-Eiffel-Tower-wise, those simples* are not an object (given what (RN2) says about 'object'). Thus no *xx* satisfies 'is a nose-Eiffel-Tower', and thus 'the nose-Eiffel-Tower doesn't exist' is true.

Fourth, consider sentences of the folk theory that involve (nominalised versions of) non-sortal (non-mereological) predicates, such as folk claims involving 'the property of being green'. Presumably, the folk theory tells us that grass, certain apples, and certain people's eyes instantiate this property; plausibly the folk theory is also committed to the claim that 'the property of being viridian is a determinate of the property of being green, which is in turn a determinate of the property of being coloured'. Given (RN5), 'the property of being green' picks out a corresponding plural property that we can call [*collectively*] *being green*. Given (RN4), 'grass', 'apples', and 'eyes' pick out pluralities of simples* arranged grass-wise, apple-wise, and eye-wise respectively. Thus, for example, 'grass instantiates the property of being green' is equivalent to 'pluralities of simples* arranged grass wise instantiate the property of [*collectively*] *being green*', which is true, given microphysical nihilism. Likewise, 'certain apples instantiate the property of being green' and 'certain people's eyes instantiate the property of being green' come out true, given (RN4) and (RN5). Finally, since 'the property of being viridian' and 'the property of being coloured' pick out [*collectively*] *being viridian* and [*collectively*] *being coloured* respectively, and since [*collectively*] *being viridian* is a determinate of [*collectively*] *being green*, which is in turn a determinate of [*collectively*] *being coloured*, the folk claim that 'the property of being viridian is a determinate of the property of being green, which is in turn a determinate of the property of being coloured' also comes out true.

Fifth, consider claims of the folk theory involving (nominalised versions of) *mereological* predicates. The mereological predicate with which the folk are most familiar is plausibly 'parthood'. The folk theory presumably makes a lot of claims about what is a part of what. For example, the folk theory plausibly tells us that Obama's left leg bears the parthood relation to his body, that a certain wheel bears the parthood relation to a certain car, and that a certain table-leg bears the parthood relation to the table. Now, (RN3) tells us that 'parthood' expresses the relation defined by PARTHOOD; §2.4.1 tells us that PARTHOOD defines parthood as being equivalent to *amonghood*. So, together with the rest of (RN1)-(RN6), e.g. 'Obama's left leg bears the parthood relation to his body' is equivalent to 'the simples* arranged Obama's-left-leg-wise bear the amonghood relation to the simples* arranged Obama's-body-wise'. Nihilists regard this sentence as true, and so, given (RN1)-(RN6), 'Obama's left

leg is a part of Obama's body' is also true. In general, given (RN1)-(RN6), claims of the folk theory concerning the extension of parthood are true.⁶¹

I think that, on the assumption of (RN3), claims of the folk theory involving the extension of other mereological predicates also come out true. Take 'composition', for example. Plausibly, the folk theory is committed to the idea that this is the relation that is borne by some parts to the object that they are parts of as long as those parts don't overlap, and as long as no part of the whole fails to overlap at least one of those parts.⁶² On the assumption that 'composition' picks out the relation defined by COMPOSITION (see §2.4.3.3), the folk theory's claims about this are true. Similarly, the folk theory is committed to the idea that the property of being composite is instantiated iff the property of *having parts* is instantiated. Again, on the assumption that 'the property of being composite' picks out the relation defined by COMPOSITE (see §2.4.2), this is true.

Finally, with regard to mereological predicates, I think it's also plausible that the folk theory is committed to the idea that parthood is an *intimate* relation, in that, for example, wholes are located wherever their parts are and are nothing over and above those parts (see Sider 2007 and Cameron 2014). Assuming (RN3) (i.e. the reductive nihilist's claims about what mereological predicates express) also makes true these sorts of claims. For example, given (RN3) and PARTHOOD, 'parthood' picks out the amonghood relation, and the fact that the simples* arranged Obama's-left-leg-wise are among the simples* arranged Obama's-body-wise does seem to entail that the latter plurality of simples* are at least partially located wherever the former plurality is. Similarly, on the assumption that 'composition' picks out the relation defined by COMPOSITION, wholes really are nothing over and above the parts that compose them (this much is guaranteed by the fact that, given COMPOSITION, composition entails identity – see §2.4.3.3).

Sixth, and lastly, the folk theory makes claims concerning *existence*. One such sort of claim concerns the *inferential role* of existence. For example, the folk would presumably invariably accept claims such as 'given the claim that "a penguin is waddling happily", it is legitimate to infer that "a penguin exists"'. Given (RN1), 'existence' expresses the plural existential quantifier, $\exists xx$,⁶³ which has the sort of inferential role that the folk theory attributes to existence. Thus claims in the folk theory concerning the inferential role of existence will come out true, given (RN1)-(RN6). For example, it's legitimate to infer from 'a penguin is waddling happily' that 'a penguin exists', because, given (RN1)-(RN6), this inference is equivalent to the one from 'the simples* that are arranged penguin-wise that are also an object are collectively waddling happily' to ' $\exists xx$ (xx are simples* and xx are arranged penguin-wise and xx are an object)', which is a legitimate inference.

⁶¹ A slight complication here is that the folk theory plausibly denies that anything can be a part of itself, whereas by definition any plurality is among itself. Reductive nihilists can react to this by claiming that 'parthood' doesn't express the relation defined by PARTHOOD, but rather that defined by PROPER PARTHOOD.

Alternatively, she can maintain that she's interested in defining the notion of 'parthood' as it is used in philosophy: this notion is plausibly *parasitic* on the folk notion of parthood (so it's still instructive to see how the folk theory defines 'parthood'), but obeys the added stipulation that everything is trivially a part of itself.

⁶² You might think that the folk theory doesn't quite imply that 'composition' functions in this way: perhaps, for example, the folk theory assents to some parts composing a whole even when those parts overlap. But, as in the case of 'parthood' that I discussed in the previous footnote, the reductive nihilist can then claim that she is interested in the philosophical use of 'composition', which, though parasitic on the folk use, adds the stipulation that the composing parts cannot overlap.

⁶³ Perhaps more accurately: 'existence' expresses what the plural existential quantifier means. For ease of expression I'll continue to abbreviate this to the claim as it's written in the main text.

Other claims that the folk theory makes concerning existence are simply about *what exists*. For example, the folk theory presumably claims that the Eiffel Tower exists, that pens exist, that tables exist, etc. Given (RN1)-(RN6), these claims are equivalent to the claims that there are simples* arranged Eiffel-Tower-wise, pluralities of simples* arranged pen-wise, and pluralities of simples* arranged table-wise, and that these pluralities of simples* are all objects (as per the characterisation of objecthood in (RN2)). Nihilists accept these claims are true. For another example, the folk theory presumably claims that there is no nose-Eiffel-Tower. This claim is also true, given (RN1)-(RN6): given (RN1) and (RN4), for there to be a nose-Eiffel-Tower there would have to be some xx that both are arranged nose-Eiffel-Tower-wise and are an object; given (RN2), the xx that are arranged nose-Eiffel-Tower-wise are not an object.

In summary, then, I've considered an apparently representative sample of claims made by the folk theory of composite objects. I've argued that, on the assumption of (RN1)-(RN6), these claims are true. I think it's reasonable to conclude that the folk theory of composite objects as a whole is true, on the assumption of (RN1)-(RN6). Given the way I've defined REALISER, this is equivalent to saying that, on the assumption that each t_i of the folk theory of composite objects refers to the corresponding ee_i of REALISER, the folk theory of composite objects is true. This in turn is equivalent to saying that the realisation formula of the folk theory of composite objects is true relative to the variable assignment that assigns each ee_i of REALISER to the corresponding xx_i of that realisation formula. And finally, this in turn is equivalent to the claim that REALISER realises the folk theory of composite objects.

In §3.5 I'll attempt to further strengthen the claim that REALISER realises the folk theory of composite objects by considering and responding to an important objection to that claim.⁶⁴ For now, though, I want to move on to suggest why we should think that REALISER *uniquely* realises the folk theory of material objects.

3.4.3 From Realisation to Unique Realisation

To show that REALISER realised the folk theory of composite objects, I argued that the folk theory is true relative to the assumption each of its T-terms, t_i , refers to the corresponding ee_i of REALISER. In this section I argue that there is probably no other n -tuple that is like REALISER in this respect, and thus that it's plausible that REALISER uniquely realises the folk theory of composite objects.

Obviously it is impossible to consider every possible n -tuple that might realise the folk theory of composite objects. Instead, I'll present two general considerations that suggest that the folk theory is not multiply realised.

First, note that the folk theory of composite objects is incredibly complex. It consists of every sentence accepted by the folk concerning composite objects and the properties, both mereological and otherwise, that those composite objects instantiate. In general, the more complex a theory is, the harder it is for a given n -tuple to realise it. For example, if my theory says only that 'X punched Y',

⁶⁴ Further support for the claim that REALISER realises the folk theory of composite objects can be found in §5.5, where I argue that reductive nihilists can respect the folk intuitions concerning the falsity of mereological essentialism and mereological sufficiency. For my argument there supports the claim that the folk theory's *modal* claims about composite objects will come out true relative to the assumption that each of its T-terms, t_i , refers to the corresponding ee_i of REALISER. I also note briefly in §5.5 that my argument there plausibly generalises to the *temporal* case, which helps to make plausible the idea that the folk theory's claims about the *persistence* of composite objects will also come out true relative to the assumption that each of its T-terms, t_i , refers to the corresponding ee_i of REALISER.

then any n -tuple such that its first member punched its second member will realise my theory; if I elaborate on my theory so that it instead says that ‘X punched Y at 10.05 on 28th May 2010, whilst Y was walking his dog’, then many of these n -tuples will no longer realise my theory. As I keep on adding detail and complexity to my theory, it will become less and less likely that my theory is multiply realised. Thus the incredible complexity of the folk theory of composite objects already makes it unlikely to be multiply realised.

Second, our starting assumption that microphysical nihilism is true leaves us with very few entities in our ontology out of which we could try to construct an n -tuple distinct from REALISER that also realises the folk theory of composite objects. For example, if the nihilist included composite* F s in her ontology in addition to simples* arranged F -wise, then she would have to worry that a rival n -tuple containing composite* F s in place of REALISER’s pluralities of simples* arranged F -wise (and the parthood* relation in place of REALISER’s parthood relation, etc.) might also realise the folk theory of composite objects. But of course, the nihilist *doesn’t* include composite* entities in her ontology, and so she doesn’t have to worry that this sort of rival n -tuple to REALISER might render the folk theory of composite objects multiply realised.

Now, despite these considerations, we might still worry that the folk theory of composite objects is multiply realised by n -tuples of *nihilistically acceptable* (pluralities of) entities that are very similar to the (pluralities of) entities that are the members of REALISER. Consider the member of REALISER that corresponds to the folk theory’s T-term ‘the Eiffel Tower’. This member of REALISER is a certain plurality of simples*, ee . Now consider $ee-1$, which is the result of subtracting one simple* from ee . Suppose we replace REALISER’s ee with $ee-1$. Now we have a new n -tuple – call it REALISER-1. Plausibly, REALISER-1 also realises the folk theory of composite objects. For $ee-1$ differs so minimally from ee that any claim of the folk theory of composite objects that is true on the assumption that ‘the Eiffel Tower’ refers to ee must also be true on the assumption that ‘the Eiffel Tower’ instead refers to $ee-1$. Since REALISER-1 is otherwise identical to REALISER, and since we’ve already seen that the folk theory of composite objects is true relative to the assumption that each t_i of the folk theory refers to the corresponding ee_i of REALISER, it seems clear that the folk theory is also true relative to the assumption that each t_i of the folk theory refers to the corresponding ee_i of REALISER-1.

In general, if REALISER realises the folk theory of composite objects, then so will a great many other very similar n -tuples that differ from REALISER only in that one or more of the pluralities of simples* they contain are one simple* smaller or one simple* larger than the corresponding pluralities of simples* that REALISER contains. So the claim that REALISER uniquely realises the folk theory of composite objects is under threat.

But I don’t think that reductive nihilists should lose sleep over this sort of difficulty with the claim that REALISER uniquely realises the folk theory of composite objects. The reason for this has to do with the fact that this sort of difficulty will apparently arise for *a great number* of potential applications of the (pluralised) Ramsey-Lewis treatment. That fact makes it dialectically appropriate in this context to assume that this sort of difficulty is resolvable. Let me explain.

First, the case of REALISER, REALISER-1, etc. is an instance of a more general difficulty for the project of finding a unique realiser for the relevant theory: the threat that, for any candidate unique realiser, there will be an alternative n -tuple whose members differ from the members of the candidate unique realiser’s members only very slightly, and that also realises the theory in question. Call this

general threat the *problem-of-the-many* difficulty, since it is inspired by Unger's *Problem of the Many* (1980).

The problem-of-the-many difficulty arises for a great number of potential applications of the pluralised Ramsey-Lewis treatment. Clearly, if microphysical nihilism is true, then almost every potential application of the pluralised Ramsey-Lewis treatment will be subject to the problem-of-the-many difficulty in just the same way as that exemplified by REALISER and REALISER-1: if some theory is realised by an n -tuple that has a plurality of simples*, xx , as its i -th member, then it will in almost all cases also be realised by another n -tuple that differs from the first only in that its i -th member is just the result of minusing one simple* from (or adding one simple* to) xx . But the same difficulty arises even without the assumption of microphysical nihilism. Suppose either universalism or some form of restrictivism is true. Then whenever a theory is realised by an n -tuple that has as its i -th member a composite* entity that is composed* of a plurality of simples* xx , that theory will in almost all cases also be realised by another n -tuple that differs from the first only in that its i -th member is a composite* entity that is composed* of a plurality of simples* $xx-1$, where $xx-1$ is the result of minusing one simple* from xx . For example, if a detective's theory is realised by [Mrs Peacock, Professor Plum], it will surely also be realised by [Mrs Peacock-minus-1, Professor Plum], where Mrs Peacock-minus-1 is the composite entity composed out of all *except one* of the simples* that compose* Mrs Peacock. Alternatively, if a form of restrictivism is assumed that denies the existence of Mrs Peacock-minus-1, the detective's theory will nevertheless still be multiply realised by [Mrs Peacock, Professor Plum] and [the simples* arranged Mrs-Peacock-minus-1-wise, Professor Plum].

What this shows, I think, is that the problem-of-the-many difficulty is a general difficulty facing the Ramsey-Lewis approach to defining terms. If the difficulty is insoluble then we will be prevented from using the pluralised Ramsey-Lewis treatment to find the referents of theoretical terms in a great number of cases, including in cases (such as that of a detective's theory) where, intuitively, the pluralised Ramsey-Lewis treatment *does* tell us what the referents of our terms are. For in order to use the pluralised Ramsey-Lewis treatment to find the referents of the T-terms of a theory, we need to find an n -tuple that uniquely realises that theory, but the problem-of-the-many difficulty will apparently almost always prevent us from doing so. As a consequence, if we want to maintain that the Ramsey-Lewis approach is in some form or other acceptable, we must suppose that the problem-of-the-many difficulty can be solved.

Therefore, since (as I pointed out at the start of this chapter [§3.1]) I'm simply assuming here that the Ramsey-Lewis treatment is not fundamentally flawed, it is dialectically appropriate to assume that the problem-of-the-many difficulty can be solved.⁶⁵ One suggestion on this front is to say that the claim

⁶⁵ Objection: the source of the problem-of-the-many difficulty is in the *pluralised* version of the Ramsey-Lewis treatment, not in the Ramsey-Lewis approach in general, so it's not dialectically appropriate after all to assume that the pluralised Ramsey-Lewis treatment can solve the difficulty. First response: even the ordinary, non-pluralised version of the Ramsey-Lewis treatment seems subject to the difficulty, as illustrated by the example, given above, of the detective's theory being realised both by [Mrs Peacock, Professor Plum] and [Mrs-Peacock-minus-1, Professor Plum], where all three of Mrs Peacock, Professor Plum, and Mrs-Peacock-minus-1 are single entities. One might respond by adopting a version of restrictivism according to which there are no pairs of composite* entities such that the first differs from the second only minutely, à la Mrs Peacock and Mrs Peacock-minus-one. But even if this rejoinder succeeds, consider this second response: the pluralised version of the Ramsey-Lewis treatment is the *only acceptable version* of the Ramsey-Lewis treatment (for the sorts of reasons I gave in §3.3 to do with the possibility of plural reference, the possibility of joint-role-occupancy, and the possibility of grammatically plural T-terms), so any difficulty for the pluralised version of the Ramsey-Lewis treatment is automatically a difficulty for the Ramsey-Lewis approach in general.

that a given n -tuple ‘uniquely’ realises a theory in this context allows for there to be other n -tuples that differ only very slightly to that n -tuple (allowing that it is therefore in general vague which of several very similar candidates the T-terms of a theory refer to); another potential solution would be to appeal to the notion of reference magnetism to break the tie between similar n -tuples that all seem to realise the theory in question – see Hawthorne (1994). But ultimately we don’t need to take a stand on what the solution: dialectically, we’re permitted to simply assume that there is one.

The upshot of all of this is as follows. First, the initial two points I made concerning the incredible complexity of the folk theory of composite objects and the sparse ontology that comes with accepting MICROPHYSICAL NIHILISM make it plausible that no n -tuple that isn’t very similar to REALISER realises the folk theory of composite objects, on the assumption of microphysical nihilism. Second, although there appear to be n -tuples that are very similar to but still distinct from REALISER and that also realise the folk theory of composite objects, any acceptable final version of the pluralised Ramsey-Lewis treatment will give us the resources to maintain that the claim that REALISER uniquely realises the folk theory is not threatened by the presence of very similar n -tuples such as REALISER-1.

With all of this in mind, I think there’s a strong case for the claim that REALISER uniquely realises the folk theory of composite objects.

3.4.4 – From Unique Realisation to Reference

Given that REALISER uniquely realises the folk theory of composite objects, the modified plural Carnap sentence – i.e. step (5*) of the pluralised Ramsey-Lewis treatment – allows us to infer that each t_i of the folk theory refers to the corresponding ee_i of REALISER. Since I characterised REALISER as containing those (pluralities of) entities and properties that the reductive nihilist claims to be the referents of the T-terms of the folk theory of composite objects, that means that we can infer the reductive nihilist’s semantic claims, i.e. (RN1)-(RN6).

That is, on the assumption of MICROPHYSICAL NIHILISM, (RN1)-(RN6) hold.

3.5 – Objection: Single Entities and Single Objects

I turn finally to considering and responding to an important objection to the claim that REALISER realises the folk theory of composite objects.

The objections can be put as follows. Isn’t the folk theory of composite objects committed to the idea that composite F s such as tables, pens, the Eiffel Tower, etc. are each *single entities*? This commitment might be *explicit* in the folk theory, as it would be if the folk would invariably agree to the claims that ‘each table is a single entity’, ‘the Eiffel Tower is a single entity’, etc. Alternatively, this sort of commitment might be *implicit* in the folk theory. For example, perhaps the *grammatical singularity* of ‘table’, ‘pen’, and ‘the Eiffel Tower’ as they appear in the folk theory implies that the referents of these terms must be single entities. Or perhaps this implication follows from the fact that (for example), when asked how many things there are in a room containing only a table, the folk would invariably answer ‘one’. At any rate, if the folk theory of composite objects *is* committed (either implicitly or explicitly) to the idea that composite F s are each single entities, then REALISER cannot realise the folk theory. For on the assumption that each T-term, t_i , of the folk theory refers to

we can avoid saying that the Ramsey-Lewis treatment is fundamentally flawed by rejecting the move from the canonical Ramsey-Lewis treatment to the pluralised version, and furthermore assuming a version of restrictivism

the corresponding *ee_i* of REALISER, the reference of e.g. ‘the Eiffel Tower’, ‘pen’, ‘table’ etc. is in each case a (proper) *plurality* of entities. The claims of the folk theory that imply that the Eiffel Tower, pens, tables, etc. are each single entities would therefore be false, in which case REALISER wouldn’t realise the folk theory.

The right response to this suggestion, I think, is to simply deny that the folk theory of composite objects is committed to the idea that composite *F*s are each single entities.

First, it isn’t plausible that the folk theory *explicitly* includes claims such as ‘the Eiffel Tower is a single entity’. That’s because ‘entity’, as we’re understanding it here, is a technical term of metaphysics, not one with which the folk are at all familiar (see §2.2.1, where I characterised this term). Plausibly, the term ‘entity’ doesn’t appear in the folk theory of composite objects at all. What’s more, if we were to try to introduce this term to the folk and ask them whether they think that e.g. the Eiffel Tower is a single entity or not, impressing on them that by ‘entity’ we here don’t simply mean ‘object’ but rather mean the notion defined in §2.2.1, I don’t think it’s plausible that they would answer with a resounding ‘yes’ – more plausible, I think, is that they would shrug their shoulders and profess to not really having an idea one way or the other.⁶⁶

Indeed, it’s especially implausible that the folk theory of composite objects includes the claim composite *F*s are single entities once we recall that the folk theory of composite objects does plausibly include the claim that composite *F*s are single *objects*. I think any temptation to ascribe to the folk a widespread belief that e.g. the Eiffel Tower is a single entity is rooted in the conflation of that belief with the belief that the Eiffel Tower is a single object. I’ve already granted (§3.4.2) that the folk theory of composite objects says of composite *F*s such as the Eiffel Tower that they are single objects (and I’ve already shown that this feature of the folk theory supports rather than threatens the claim that REALISER realises the folk theory). This seems to undermine any *prima facie* plausibility the distinct claim that the folk theory says of composite *F*s that they are single *entities* might have had.⁶⁷

Similar considerations tell against the idea that the folk theory of composite objects is implicitly committed to the idea that composite *F*s are each single entities. First, reductive nihilists hold that grammatical singularity in the folk dialect typically tracks single *objecthood*, and maintain that the grammatical singularity of e.g. ‘the Eiffel Tower’ as it appears in the folk theory is only indicative, if it is indicative of anything at all, of the folk theory’s commitment to the Eiffel Tower being a single object. To claim that grammatical singularity instead tracks single *entityhood* would be to implausibly claim that a technical concept of metaphysics with which the folk are not at all familiar plays an important role in the grammar of ordinary language. Similarly, reductive nihilists hold that folk claims concerning number typically track how many *objects* there are, not how many *entities* there are.⁶⁸ That is, the folk typically count relative to the sortal predicate ‘is an object’ (I briefly characterise the notion of counting relative to a sortal in §2.2.1; see also Frege [1953: §46]). Again, it seems

⁶⁶ For what it’s worth, in my experience, non-philosophers often tend to lean the other way and *deny* that e.g. a given table is a single entity, usually for reasons that have to do with the claim (reminiscent of Eddington [1929: x-xii]) that ‘physics tells us that the table is just a load of tiny particles whizzing about’.

⁶⁷ See also §6.4.3 and §6.5.3, where I say some more about ordinary belief about ‘object’ and ‘entity’.

⁶⁸ Perhaps it’s more accurate to say that folk counting typically tracks the number of *maximal* objects, where a maximal object is an object that is not a proper part of any other objects. This is because, in a room containing only a table, the folk would presumably say that there is *one* object in the room. In this case the folk apparently count the table (which is a maximal object) but do not count e.g. the table-leg (which is not a maximal object, as it is a proper part of a further object, the table).

implausible that folk counting instead tracks single entityhood, where single entityhood is a technical concept that has little relevance to non-philosophical life.

In sum, then, it isn't plausible to suggest that the folk theory is committed (either explicitly or implicitly) to the idea that composite *F*s such as the Eiffel Tower, tables, etc. are each single entities. Much more plausible is the claim that the folk theory is committed (perhaps both explicitly and implicitly) to the idea that composite *F*s such as these are single *objects*. But this latter claim is consistent with the claim that REALISER realises the folk theory of composite objects. For on the assumption that each t_i of the folk theory refers to the corresponding ee_i of REALISER, it is *true* that e.g. the Eiffel Tower, tables, pens, etc., are (single) objects.

The objection that REALISER fails to realise the folk theory of composite objects because the folk theory implies that composite *F*s are each single entities therefore fails.

3.6 – Conclusion

I've argued on independent grounds that the canonical version of the Ramsey-Lewis treatment ought to be replaced with what I've called the *pluralised* Ramsey-Lewis treatment, which differs from the canonical version in that it allows that the T-terms of a given theory might refer to a (proper) *plurality* of entities, rather than only to a single entity. I then argued that REALISER realises the folk theory of composite objects, which is the theory that implicitly defines the terms and predicates that are relevant to assessing the truth of (RN1)-(RN6). In fact, I argued that, on the assumption of microphysical nihilism REALISER *uniquely* realises the folk theory of composite objects. It follows from this claim, together with step (5*) of the pluralised Ramsey-Lewis treatment, that if MICROPHYSICAL NIHILISM is true, then the T-terms of the folk theory of composite objects refer to the corresponding members of REALISER. Given the way I characterised REALISER, that means that (RN1)-(RN6) follow.

In sum, then, I've argued that if MICROPHYSICAL NIHILISM is true, then the semantic claims of reductive nihilism follow from the pluralised Ramsey-Lewis treatment of theoretical terms.

Chapter 4

Reductive Nihilism and The Causal Theory of Reference

4.1 – Introduction

In this chapter I'll argue that the semantics of reductive nihilism can be motivated by the causal theory of reference.

The causal theory of reference purports to account for how our terms (and predicates) refer to certain things by appealing to facts about what the person who introduced a term was in causal contact with, and about how subsequent users of the term are causally related to the term-introducer. Very roughly (and we'll see some more of the details shortly), the idea is that terms are introduced by ostension, in that the introducer of the term in question gestures at some object and says that the term is to refer to *this*. In so doing, they make it the case that the term refers to object they gestured at. Then other language users acquire the term from its introducer (or from others who themselves acquired the term from its introducer, etc.), and their utterances of that term also refer to the object that the introducer gestured at when they introduced the term.

That's a rough characterisation of the causal theory of reference as it applies to *names*; the causal theory also apparently generalises to other bits of vocabulary. For example, consider the sortal predicate, 'is [a lump of] gold'. On the causal theory, the introducer of this predicate introduced it by gesturing at a certain lump of stuff, thereby fixing that the predicate is to apply to anything relevantly similar to *this* lump of stuff. In the case of gold, we might suppose for now that the relevant dimension of similarity is *being composed of atoms of the same atomic number*. Thus, since the lump of stuff that the introducer of 'is [a lump of] gold' in fact gestured at was composed of atoms with the atomic number 79, the predicate 'is [a lump of] gold' applies to any lump composed of atoms with the atomic number 79. This predicate can then be acquired by others in the same way as names are acquired.

Even with only this admittedly brief characterisation of the causal theory of reference in hand, it's clear that the theory grants the *way the world is* a significant role in determining the meanings/references of our terms and predicates. For example, suppose I introduce a name, 'Potty', thinking that I am gesturing at a potted plant. However, suppose I am really gesturing at a well-disguised Martian spy. Given the causal theory, 'Potty' in that case refers to the Martian spy, rather than a plant. Similarly, if what I gestured at when I introduced 'is a [lump of] gold' had the atomic number of 80 rather than 79, then my predicate 'is a [lump of] gold' would apply to any lump of stuff composed of atoms with the atomic number 80, rather than the atomic number 79.

With this in mind, then, it is perhaps also clear how the causal theory might support the semantic claims of reductive nihilism, given the assumption of the metaphysical claim of microphysical nihilism, i.e.:

MICROPHYSICAL NIHILISM: $\forall x (x \text{ is simple}^*) \wedge \forall x (x \text{ is a simple}^* \text{ iff } x \text{ is either a quark or a lepton})$

For given MICROPHYSICAL NIHILISM, what I gestured at when I introduced 'Potty', for example, was not a composite* plant, nor a composite* Martian spy, but rather a certain plurality of simples*. In particular (and ignoring the possibility of Martian spies) I gestured at a certain plurality of simples* arranged plant-wise. Given the causal theory, it seems that the reference of my term 'Potty' is

therefore a certain plurality of simples* arranged plant-wise, in line with the reductive nihilist's semantics. Similarly, if what I gestured at when I introduced the predicate 'is a [lump of] gold' was a certain plurality of simples*, then it seems that I have thereby made it the case that 'is a [lump of] gold' applies to certain pluralities of simples* that are relevantly similar to the plurality of simples* that I gestured at. Then 'is a [lump of] gold' would be a plural sortal predicate, in line with the reductive nihilist's semantics.

Indeed, I'm not the first to have noticed that this sort of nihilist semantics can be motivated by a causal theory of reference in this sort of way. Contessa (2014) has endorsed a nihilist semantics not dissimilar from that of the reductive nihilist, and has similarly appealed to semantic externalism to motivate his claims. My argument in this chapter is therefore of a piece with Contessa's work, to some extent. But my argument here differs from Contessa's in two key respects. First, the semantic claims that I think follow from the assumption of microphysical nihilism, given the causal theory, are different in some ways to ones that Contessa identifies. Here again are the semantic claims of reductive nihilism:

- (RN1) 'Existence' and its cognates express the plural existential quantifier, $\exists xx$.
- (RN2) 'Object' is not synonymous with 'entity', but is rather a plural predicate that can be satisfied by proper pluralities of entities. Roughly, xx satisfy 'object' iff it is intuitive that the region of space exactly occupied by xx is exactly occupied by an object.
- (RN3) Mereological terms such as 'parthood' and 'composition' are defined by PARTHOOD, COMPOSITION, etc.
- (RN4) Sortal predicates that purport to apply to single composite F s express properties that are instantiated by xx iff xx are both arranged F -wise and are an object.
- (RN5) Predicates that putatively ascribe properties to single things express plural properties that are satisfiable by proper pluralities (thus, for example, 'weighs 10kg' express *jointly weighing 10kg*).
- (RN6) Referring expressions of the form 'N' that purport to refer to a single composite F refer plurally to xx iff xx are arranged N -wise and xx are an object.

Whilst Contessa seems to endorse (RN1), (RN4), and (RN5), and appears open to accepting something like (RN2) (cf. the discussion in his [2014: §6]), he appears to reject (RN3): his explicit denial that (given microphysical nihilism) e.g. cats are composite objects (2014: 202, 217) suggests that he takes 'composite' to be defined by COMPOSITE*, and in general that pieces of mereological vocabulary are defined as the mereologist usually defines them (see §2.3), rather than in the way that the reductive nihilist defines them. Contessa also says nothing about *names*, and thus is neutral with regard to (RN6). The second difference between my argument here and Contessa's is that Contessa overlooks or ignores some complications that I'll discuss with the attempt to motivate the sort of nihilist semantics he endorses (this is particularly clear if we take Contessa to tacitly endorse something like (RN6), as motivating this claim via the causal theory raises some especially serious – if ultimately surmountable – difficulties, as we'll see).

Here, then, is the plan for the rest of the chapter. I'll start (§4.2) by considering all claims of the reductive nihilist *apart from* (RN6), arguing that they follow from the assumption of MICROPHYSICAL NIHILISM, given the causal theory. In the course of doing so, I'll elaborate on the

causal theory as it applies to the types of vocabulary at issue in (RN1)-(RN5); in particular I'll elaborate on how the causal theory must be modified so as to avoid the infamous *qua* problem that plagues simple versions of it. My argument will be that (RN1)-(RN5) are motivated by this modified version of the causal theory. In §4.3 I turn to arguing for (RN6). In doing so I'll elaborate on how the causal theory works with regard to *names*, and again, the version of the causal theory that I'll use to support (RN6) is a version designed to avoid the *qua* problem. §4.4 concludes.

Before jumping into the argument, though, I want to reemphasise a dialectical point that I have made already. It's no part of my project here to *defend* the causal theory. Rather, as I did with the Ramsey-Lewis treatment of theoretical terms (see chapter 3), my strategy here is to take an established and influential semantic theory and show that the semantic claims of reductive nihilism are well-justified, given that semantic theory. In effect, then, I am in this chapter endorsing a conditional claim: *if* the causal theory of reference is correct, then (if MICROPHYSICAL NIHILISM holds) (RN1)-(RN6) are true. As such, it's dialectically appropriate to assume the antecedent, i.e. assume the correctness of the causal theory, and try to derive the consequent. So, though I'll consider an objection to the causal theory (namely the *qua* problem) I'll only *sketch out* a solution to that problem on behalf of the causal theorist: I won't delve into the finer details of that solution, and nor will I attempt to show that that 'solution' is ultimately viable. Rather, since it is dialectically appropriate to assume the correctness of the causal theory here, it's appropriate to simply assume that the solution is workable, as it must be if the causal theory is correct.

With that in mind, I turn to arguing that (RN1)-(RN5) can be motivated by the causal theory of reference.

4.2 – *Predicates and Existence*

In this section I'll deploy the causal theory of reference to motivate the reductive nihilist's semantic claims concerning both predicates and the term 'existence'. I'll assume the truth of microphysical nihilism throughout. In §4.2.1 I'll elaborate on the causal theory, explain the *qua* problem for it, and sketch out how any acceptable form of the causal theory must be modified to solve the *qua* problem. In §4.2.2 I'll show how the causal theory, so modified, motivates (RN1)-(RN5).

4.2.1 – *The Causal Theory and the Qua Problem*

Proper names aside (as they are the subject of §4.3), most discussions of the causal theory concern its application to so-called *natural kind terms*. I'll follow suit, and introduce both the causal theory and the *qua* problem using the example of natural kind terms.

Roughly speaking, natural kind terms are general terms that purport to apply to members of particular kinds of naturally occurring things. Typical examples are 'giraffe', 'animal', 'gold', 'water', etc. In the terminology of the present project, natural kind terms are therefore best understood as a certain class of sortal predicates, including 'is a giraffe', 'is an animal', and so on. But nothing important hangs on whether we treat these bits of vocabulary as predicates or terms, so for now I'll follow the orthodoxy and treat them as terms.

I've already given a quick overview as to how the natural kind term 'gold' (or 'is [a lump of] gold') comes to have an extension, on the causal theory. Here I'll work through another example, that of 'giraffe' (my discussion throughout is based on the loci classici of the causal theory, Putnam [1975] and Kripke [1980]).

The introducer of the term ‘giraffe’, then, once came across an example of what we would now call a giraffe. Let’s call the thing that the introducer came across the *sample*.⁶⁹ Our introducer introduced the name ‘giraffe’ as a general term to refer to things relevantly similar to *that* (with a gesture at the sample). What is it to be relevantly similar to a sample in this sort of case? Typically, it is to have a similar underlying structure. In this case, the relevant underlying structure is plausibly *biological*, such that anything that has a similar biological structure to the sample that was gestured at or ostended when ‘giraffe’ was introduced counts as a giraffe (we’ve already seen that in the case of ‘gold’ it’s plausible that anything with the same *atomic* structure as the sample that was ostended when ‘gold’ was introduced counts as an example of gold).

Causal theorists then augment this account of the reference-*fixing* of ‘giraffe’ with an account of reference-*borrowing*. We need not worry about the details too much here, but roughly the idea is that other speakers acquire the term ‘giraffe’ via a chain of people leading back to the introducer of the term, and that any instance of ‘giraffe’ acquired in this way has the same extension as that that was fixed by the introducer of the term. Speakers who have successfully acquired or borrowed a term in this way are then said to be appropriately causally connected to the introduction of the term.

I also note that causal theorists can allow that some natural kinds have their referents fixed not by ostension but by description. A plausible instance of this is Putnam’s example of the general term ‘hunter’, which was plausibly introduced by the description ‘one who hunts’ (1975: 164), and that therefore has in its extension all and only those things that hunt.⁷⁰ Despite this, the causal theorists regard *most* natural kinds as being amenable to the sort of reference-fixing-by-ostension account outlined above.

Already we have enough detail on the causal theory here to appreciate the *qua* problem for it.⁷¹ The problem concerns the reference-fixing component of the causal theory. Simply put, the problem is that the act of ostending a sample cannot by itself establish a determinate extension of a natural kind term. This is for two reasons. Consider again the attempt to fix the extension of ‘giraffe’ by gesturing at a certain long-necked animal. The first reason is that that gesture almost certainly will not adequately determine what the sample is. Suppose our term-introducer just waves vaguely in the direction of the animal: then it’s indeterminate whether the sample is supposed to be the animal, or the air particles in between our term-introducer and the animal, or just the animal’s *legs*, or just the animal’s *fur*, etc.

The second reason is that, even if our term-introducer’s gesturing is specific enough to single out the giraffe itself as the relevant sample, gesturing does nothing to determine *what respect it is* in which other objects must be similar to that sampled giraffe in order to fall in the extension of the term ‘giraffe’. Our sample here is a member of a great many similarity classes: the class of long-necked things (which might also contain ostriches, diplodoci, geese, maybe certain wine bottles, etc.), the class of things that live in Africa (alongside certain elephants, people, lions, etc.), the class of

⁶⁹ Plausibly, in many case the introducer will ‘gesture at’ multiple samples when introduced the term in question. Perhaps they’ll also gesture at multiple *foils*, i.e. multiple things that they take the term they’re introducing *not* to apply to. For brevity’s sake I’ll often suppress this point in what follows.

⁷⁰ On a plausible reading of the relevant passage, Putnam doesn’t just think that the reference of ‘hunter’ is *fixed* by description, but also that subsequent uses of the term refer in virtue of the user associating it with the descriptive content ‘one who hunts’. But in the main text I use the example only to illustrate the fact that causal theorists can *at least* say that the reference of some terms is fixed by description (leaving open what story they want to tell about subsequent uses of that term).

⁷¹ On the *qua* problem, see Papineau (1979: 158-68), Devitt (1981), Sterelny (1983), Devitt and Sterelny (1999: 90-3), Stanford and Kitcher (2000), and Thomasson (2007: §2.3). Stanford and Kitcher (2000) also suggest that a proto-version of the *qua* problem is to be found in Locke ([1690] 1959).

mammals, the class of living beings, and so on. The act of ostension alone does nothing to determine that ‘giraffe’ applies to all and only the giraffes, as opposed to all and only the long-necked things, or all and only the living things, etc. Above I said that the relevant dimension of similarity in the case of ‘giraffe’ is plausibly something like *having a similar biological structure*. But nothing about the act of gesturing alone seems to make this the case.⁷²

There is some consensus that the best way to modify the causal theory so as to avoid the *qua* problem is to allow some role for *descriptive content* to creep into the picture of reference-fixing to a limited extent (see e.g. Devitt and Sterelny 1999: 90-3; Stanford and Kitcher 2000; Thomasson 2007: §2.3). In particular, the idea is that when the introducer of a natural kind term gestures at a sample whilst introducing that term, she has in mind some descriptive content that serves both to determine what exactly she is gesturing at (i.e. to determine what the sample is) and to determine a particular dimension and degree of similarity that other objects must bear to the sample in order to fall under the extension of the natural kind term in question.⁷³

To illustrate how this modified version of the causal theory might work, consider the following account, due to Stanford and Kitcher (2000: 101-2) (though Stanford and Kitcher go on to argue that this theory needs further refinement), and which they call ‘CTR1’.⁷⁴

CTR1: A natural kind term is associated with an open sentence, whose constituents are predicates applicable to observable properties and propositional connectives, together with a sample that satisfies that open sentence. The term then refers to all things having the inner constitution causing the sample to have the properties picked out by the conjuncts of the open sentence.

For example, the introducer of ‘giraffe’ associates the term with an open sentence that might say something like ‘has a long neck, has four legs, has yellow and brown fur, [...]’, and gestures at the same sample as before, which satisfies this open sentence. As a matter of fact that sample has an ‘inner constitution’, certain features of which cause it to possess the properties of having a long neck, four legs, etc.⁷⁵ ‘Giraffe’ then applies to all and only those things that possess an inner constitution with those features.

⁷² Nor can we simply *stipulate* a standard respect and degree to which other objects must be similar to a sample in order to fall under the extension of a given natural kind term that was introduced via ostension of that sample, as for different natural kind terms the respect and degree to which other objects must be similar to the relevant sample in order to fall under the extension of those natural kind terms seems to differ. Indeed, I might use the same sample to introduce natural kind terms with different extensions to one another. For example, I might use the same sample to introduce the term ‘giraffe’, as I use to introduce ‘animal’, or ‘mammal’, or ‘male’, etc., and for each of these terms a different respect and degree of similarity to the sample is apparently required for a given object to fall under the extension of that term.

⁷³ The idea here is not, or at least need not be, that all subsequent users of the term must have the same descriptive content (or indeed *any* descriptive content) in mind when they use it. We can say that the descriptive content of the term introducer helps to fix an extension of the natural kind term in question, and that all subsequent uses of that term have the same extension, regardless of whether the speaker associates any descriptive content with the term or not.

⁷⁴ Stanford and Kitcher suggest that CTR1 is suggested by Devitt and Sterelny’s discussion of the *qua* problem (see Devitt and Sterelny 1999: 91-2).

⁷⁵ Strictly speaking, and benefitting from the recent extensive discussions of metaphysical grounding, we might prefer to say that the features of the inner constitution *ground* the fact that the sample possess these properties, rather than *cause* the sample to possess them. Nothing important turns on this, so I’ll stick with Stanford and Kitcher’s use of ‘cause’.

There's reason for optimism that this sort of account ensures that 'giraffe' applies to all and only the giraffes, and not, for example, to all and only the mammals, as mammalian non-giraffes such as elephants do not have an inner constitution with the same features that caused the sample to have a long neck, etc. (otherwise such mammalian non-giraffes would typically *have* long necks, etc.), and so do not fall under the extension of 'giraffe'. If our term-introducer were to have introduced 'mammal' via ostension of our giraffe, then she would have associated that term with a different open sentence that would, if all goes to plan, determine that the features of the inner constitution that cause the sample to have the properties specified by that open sentence are shared by all and only the mammals.

That said, however, I don't intend to defend CTR1 (as I've noted, not even Stanford and Kitcher do that); nor do I intend to defend *any* particular modification of the causal theory that purports to allow that theory to avoid the *qua* problem. CTR1 is merely a helpful illustration of the sort of causal theoretic view that can hope to avoid *qua* problem. Nor will I defend here the idea that there even *is* an acceptable version of the causal theory that can avoid the *qua* problem. As I stressed at the end of §4.1, I'm assuming that some form of the causal theory is correct, which means simply assuming that there is some acceptable way of augmenting the simple causal theory with the idea that the introducer of a natural kind term associates that term with descriptive content, in such a way as to solve the *qua* problem. So I'm going to assume that there's an acceptable story to be told about the introducer of a natural kind term associating that term with descriptive content that determines both what the sample is and what the relevant dimension and degree of similarity is that other things must bear to the sample to fall under the extension of the natural kind term.

I want to make two further points with regard to the causal theoretic account of natural kind terms, the relevance of both of which will become clear later.

First, I note that a consequence of the sort of modified causal theory that I've been outlining is that if the sample that the introducer of a natural kind term doesn't satisfy the descriptive content that the introducer associates with that natural kind term, then the attempt at reference-fixing fails, and that natural kind term fails to apply to anything (see e.g. Devitt and Sterelny 1999: 91-2).

Second, consider the case of (putative) natural kind terms that apparently apply to things that do *not* share a similar inner constitution that causes (or grounds) them to have the superficial characteristics that allow them to satisfy the descriptive content associated with them by their introducer. The classic example is 'jade', which apparently has in its extension instances of two distinct minerals: jadeite and nephrite (see Putnam 1975: 160). That is, plausibly, 'jade' was introduced by ostending various samples that satisfied certain descriptive content, and as it turned out, some of those samples satisfied that descriptive content in virtue of having an inner constitution characteristic of nephrite, whilst others satisfied that descriptive content in virtue of having an inner constitution characteristic of jadeite.⁷⁶ What should causal theorists say about this sort of case? Here's Putnam:

'To sum up: if there is a hidden structure, then generally it determines what it is to be a member of the natural kind, not only in the actual world, but in all possible worlds [...]. But the local water, or whatever, may have two or more hidden structures – or

⁷⁶ Putnam also points out that some names for diseases share this characteristic: 'some diseases, for example, have turned out to have no hidden structure [i.e., in the idiolect of Stanford and Kitcher, no characteristic 'inner constitution'] (the only thing the paradigm cases have in common is a cluster of symptoms)' (1975: 160).

so many that “hidden structure” becomes irrelevant, and superficial characteristics become the decisive ones.’

(Putnam 1975: 160)

Following Putnam, then, causal theorists should say this about cases where the samples used to fix the reference of a (putative) natural kind term fail to display the same inner constitution: it is anything similar to the samples with respect to their relevant *superficial characteristics* (as specified by the descriptive content associated with the term) – rather than with respect to their inner constitution – that falls under the extension of the term.

Now, the focus up until now has been on natural kind terms, which in the terminology of the present project are a particular kind of sortal predicate (thus the natural kind term ‘giraffe’ is here understood to be the sortal predicate ‘is a giraffe’). But semantic externalists typically take their account to generalise to other kinds of predicates too (see e.g. Putnam 1975: 160, 164). Thus causal theorists can say that the predicate ‘is red’, for example, was introduced with a gesture at a particular red thing, such as a tomato, in order to fix that ‘is red’ is to apply to anything relevantly similar to the tomato. Again, although the *qua* problem clearly arises again here (for in what respect must some other object be similar to the tomato in order to fall under the extension of ‘is red’?), the same solution can apparently be offered: the descriptive content that the introducer of ‘is red’ associated with that predicate determines both what sample is being gestured at and in what respect other objects must be similar to the sample in order to fall under the extension of ‘is red’.

The final bit of vocabulary that I want to consider here is ‘existence’. It’s not clear to me what the standard causal theoretic account of how the existential quantifier comes to have a meaning is. Most discussions of ‘existence’ in the context of the causal theory concern not its *meaning* but rather the *truth conditions* of sentences in which it appears.⁷⁷ But consideration of one such account of the truth conditions of sentences in which ‘existence’ appears will tell us enough about the meaning of ‘existence’ to be getting on with. So consider the so-called *metasemantic* view of the truth conditions of sentences involving ‘existence’ (e.g. Thomasson 2008), according to which, for any general term *F*, a claim of the form ‘*F*s exist’ is true iff ‘*F*’ applies or refers. Thus, for example, ‘giraffes exist’ is true iff ‘giraffe’ applies (that is, iff the extension of ‘giraffe’ is non-empty). It seems plausible that *any* acceptable account of the meaning of ‘existence’ must preserve the truth of at least the right-to-left direction of this biconditional.⁷⁸ For if ‘giraffe’ applies, i.e. if the extension of ‘giraffe’ is non-empty, then there must *exist* something to which ‘giraffe’ applies, and so ‘giraffes exist’ must be true. That is all I’ll need to say about the meaning of ‘existence’ on the causal theory for my argument below to go through.

That concludes my exegesis of the causal theory as it pertains to predicates and ‘existence’. I turn now to arguing that (RN1)-(RN5) of reductive nihilism follow from the causal theory as I’ve characterised it.

⁷⁷ See, for example, the metalinguistic/metasemantic views discussed in Donellan (1974) and Thomasson (2008), or the ‘gappy proposition’ views defended in Braun (1993) and Salmon (1998).

⁷⁸ One reason for doubting the right-to-left direction of the biconditional is the thought that our terms can apply to entities that *used* to exist but don’t any more, or perhaps to *merely possible* entities (cf. Merricks 2015 ch. 5 §§IV-V). But this kind of counterexample is irrelevant to my argument, as I’m only concerned with existence claims concerning actual, present objects. So, if need be, we can suppose that what I’ve said about the meaning of ‘existence’ is only supposed to apply to a *limited* notion of existence according to which entities that used to exist, as well as merely possible entities, don’t exist.

4.2.2 – Applying the Causal Theory to Predicates and ‘Existence’ on the Assumption of Microphysical Nihilism

If microphysical nihilism is true, then the causal theory, as characterised in the previous section, supports the reductive nihilist’s semantic claims concerning sortal predicates, mereological predicates, non-mereological non-sortal predicates, and ‘existence’. I’ll consider each such claim in turn, though slightly out of order.

Let’s start, then, with (RN2):

- (RN2) ‘Object’ is not synonymous with ‘entity’, but is rather a plural predicate that can be satisfied by proper pluralities of entities. Roughly, *xx* satisfy ‘object’ iff it is intuitive that the region of space exactly occupied by *xx* is exactly occupied by an object.

Immediately we have hit upon a slightly tricky case. Is it plausible to treat the predicate ‘[is an] object’ like ‘[is a] giraffe’, and say that there is some inner constitution or characteristic common to all objects, such that when ‘object’ was introduced via ostension of some samples, it then has in its extension whatever has a relevantly similarly inner constitution or characteristic to those samples? Or is it instead better to treat ‘object’ more like ‘jade’, and say that there is likely no shared inner constitution or characteristic common to all instances of ‘object’ (and, in particular, common to all of the samples used to introduce ‘object’)? It isn’t clear to me which option we should favour. That said, it seems somewhat more plausible to me that ‘object’ is more like ‘jade’ than ‘giraffe’, in that it’s probable that the samples used to fix the extension of ‘object’ are probably heterogeneous with regard to their inner constitutions and characteristics, so I’ll treat it as such here (at any rate, it should be clear from my discussion how the reductive nihilist’s account would go if ‘object’ were instead like ‘giraffe’ in that the members in its extension possess a similar inner constitution).

Above I noted that, at least on Putnam’s causal theoretic account, terms like ‘jade’ (or like terms for certain diseases have no in common only a cluster of symptoms rather than any shared cause or biological structure) have in their extension anything that displays the same *superficial* characteristics as the samples, regardless of their inner constitution. With that in mind, given the causal theory, the extension of ‘object’ was fixed via the ostension of some samples (and some foils), such that ‘object’ applies to anything relevantly similar to the samples with respect to the superficial characteristics of those samples. As a matter of fact, given the truth of MICROPHYSICAL NIHILISM, the samples that were in fact ostended were pluralities of simples* - perhaps they were pluralities of simples* arranged mug-wise, or table-wise, or rock-wise, etc. That already suggests that proper pluralities can be in the extension of ‘object’, and thus that ‘object’ is a plural predicate. A plurality is then in the extension of ‘object’ iff it is similar to the sample pluralities with respect to certain superficial characteristics. What determines which superficial characteristics these are is the descriptive content associated with ‘object’ by the introducer of that predicate. Given the sense of ‘object’ at issue here (i.e. the sense of ‘object’ roughly corresponding to ‘unified, singular whole’ – see §2.4.2), I think it’s plausible that the descriptive content associated with ‘object’ during its introduction determined that any plurality that shares the samples’ superficial characteristics of *apparent unity*, *singularity*, and *wholeness* falls under the extension of ‘object’. What it is about a given plurality that causes it to display these superficial characteristics is irrelevant: perhaps *this* plurality superficially appears to be a unified, singular whole because its members are well-bonded to each other, whilst *that* plurality appears that way in virtue of its members jointly possessing a single function, and so on. What matters is that if, and only if, a given plurality displays the requisite superficial characteristics, then it is in the extension of ‘object’.

The upshot of this is, I think, the vindication of (RN2). Given what I've just said, a plurality of simples* *xx* is/are in the extension of 'object' iff *xx* appear to us jointly be a unified, singular whole. And given that the sense of 'object' at issue here is the one roughly corresponding to 'unified, singular whole', it's trivial that a given region of space is intuitively exactly occupied by an object iff what exactly occupies that region of space appears to us to be unified, singular, and whole. It follows that 'object' applies to a plurality of simples* *xx* iff it is intuitive that the region of space exactly occupied by *xx* is exactly occupied by an object, which is what (RN2) says.

Next, consider (RN4):

- (RN4) Sortal predicates that purport to apply to single composite *F*s express properties that are instantiated by *xx* iff *xx* are both arranged *F*-wise and are an object.

For example, consider the sortal predicate 'is a giraffe'. On the causal theory, the introducer gestured at a sample as they introduced this predicate. They also associated 'is a giraffe' with some descriptive content which served to specify what exactly they were gesturing at (i.e. which specified what exactly the sample was), and which also served to specify in what respect other things must be similar to the sample in order to fall under the extension of 'is a giraffe'.

The first point I want to make here is that, plausibly, the descriptive content that our term-introducer associated with the predicate 'is a giraffe' was the idea that the sample she ostended was a *single object*. The main evidence for this is that the sortal predicate introduced by our predicate-introducer was grammatically singular: it was '*is a giraffe*'. A second piece of evidence is this. Suppose our term-introducer is the victim of an optical illusion: standing near to our term introducer is a horse with its head hidden behind a tree; far away in the distance, there is a building site with a crane; by chance, the crane is positioned in just the right way relative to our term-introducer and the horse, and is coloured in just the right way, that to a person with bad eyesight, it looks as though the crane and the horse are together really a single long-necked animal standing near to our term-introducer. In this case, I think the intuitive thing to say is that the attempt to fix the reference/extension of 'is a giraffe' has failed: the alternative would be to say that 'is a giraffe' applies to all and only the combinations of nearby horses bodies and distant cranes, which seems absurd. Further, a plausible explanation as to *why* the attempt to fix the reference of 'is a giraffe' in this way has failed in this case is that the predicate-introducer was intending to introduce a predicate that applies to single *objects*, and associated 'is a giraffe' with descriptive content that is satisfiable only if the sample ostended is a single object (and the combination of a nearby horse and a distant crane is not a single object, given (RN2)).⁷⁹

Now, given microphysical nihilism (and supposing that no strange scenario involving optical illusions and distant cranes obtains), the sample that our introducer of 'is a giraffe' ostended was in fact a certain plurality of simples* arranged giraffe-wise. Given (RN2), that plurality of simples* is an object. Plausibly, the descriptive content the introducer associated with 'is a giraffe' specifies that this predicate will apply to all and only those objects that are sufficiently similar to the ostended sample with respect to *biological substructure*. Given microphysical nihilism, and given that the sample was a plurality of simples* arranged giraffe-wise, I think it's therefore plausible that 'is a giraffe' will

⁷⁹ Mightn't an alternative explanation of this be that the descriptive content associated with 'is a giraffe' requires that it may only apply to single *entities*, rather than single *objects*? I don't think that it's plausible that the descriptive content associated with ordinary sortal predicates such as 'is a giraffe' requires that such sortal predicates only apply to single entities, as I don't think it's plausible that ordinary speakers typically employ such metaphysically arcane terms as 'entity' – for more discussion, see §3.5.

apply to all and only the objects that are pluralities of simples* arranged giraffe-wise, since all and only those pluralities of simples* will jointly display the right biological substructure.⁸⁰

I see no reason for doubting that this result will generalise nicely to other ordinary sortal predicates. In each case, it's plausible that the sortal predicate in question was introduced by ostension of a sample, where that sample was a plurality of simples*. The descriptive content associated with the sortal predicate in question determines that the sortal predicate applies to all and only those objects that are pluralities of simples* arranged in an appropriately similar way to the sampled plurality.

However, I don't think the above discussion of 'is a giraffe' will generalise to cases of sortal predicates that purport to apply to *gruesome* composite *Fs*, i.e. to composite *Fs* that intuitively do not exist, such as 'is a nose-Eiffel-Tower'. Such predicates are introduced (usually by philosophers) with the same sort of descriptive content as predicates like 'is a giraffe', in that they are introduced as names for putative *single, unified things*. Hence it's plausible that, just as with predicates like 'is a giraffe', the descriptive content associated with predicates that purport to apply to gruesome composite *Fs* includes the requirement that the sample used to fix the extension of the predicate be a single object. But in the case of predicates that purport to apply to gruesome composite *Fs*, this requirement is not satisfied. Consider as an arbitrary example the predicate 'is a nose-Eiffel-Tower'. Given MICROPHYSICAL NIHILISM, the sample used to fix the reference of 'is a nose-Eiffel-Tower' can only have been some simples* arranged nose-Eiffel-Tower-wise, but the simples* arranged nose-Eiffel-Tower-wise do not satisfy 'is an object' (see §2.4.2). So the attempt to fix the reference of 'is a nose-Eiffel-Tower' has failed because the only candidates for being in its extension, i.e. pluralities of simples* arranged nose-Eiffel-Tower-wise, do not satisfy 'is an object'. Since 'is a nose-Eiffel-Tower' was an arbitrary example of a predicate that purports to apply to gruesome composites, it's plausible that this generalises, and that predicates that purport to apply to gruesome composites have no extension because the samples used to fix their extension do not satisfy one element of the descriptive content associated with them, namely the requirement that they be single objects.⁸¹

Taken together, I think all of this makes plausible the thought that, given the sort of causal theory we've been considering here (and on the assumption of microphysical nihilism), sortal predicates that purport to apply to single composite *Fs* apply to pluralities of simples* arranged *F*-wise, as long as those pluralities each count as an object. That vindicates the biconditional in (RN4).

Next, consider again (RN5):

(RN5) Predicates that putatively ascribe properties to single things express plural properties that are satisfiable by proper pluralities.

Consider the predicate 'is red', for example. According to the causal theory, this predicate was introduced by ostension of one or more red samples, and associated with some descriptive content that determines in what respect other things must be similar to the samples in order to fall under the

⁸⁰ One might doubt that pluralities of simples* can jointly display substructures. The idea seems comprehensible and plausible enough to me, but at any rate, in §8.2.1 I say more about the microphysical nihilist account of the substructures of putative composite objects.

⁸¹ One might prefer to think that predicates that purport to apply to gruesome composites have their extensions fixed via *description* rather than ostension. In that case the reasoning here becomes even more straightforward: since the predicates purport to apply to single objects, it's plausible that one element of the descriptive content used to fix their extension requires that for any *xx* to be in the extension of the predicate that *xx* must be a single object. Since any *xx* that are arranged gruesome-composite-wise will fail to be a single object, predicates that purport to apply to gruesome composite *Fs* will have empty extensions.

extension of ‘is red’. If microphysical nihilism is true, then the samples are really pluralities of simples* that are each *jointly* red. The descriptive content associated with ‘is red’ will then determine that the extension of ‘is red’ is all and only those pluralities of simples* that are similar to the sampled pluralities of simples* with respect to their colour.⁸² Thus the extension of ‘is red’ will be all and only those pluralities of simples* that are jointly red. In realist terminology, we can thus say that the predicate ‘is red’ expresses the plural property of *jointly being red*.

Further, I take this result to generalise nicely to other predicates that putatively ascribe properties to single things. Thus I think (RN5) is well-supported by the causal theory of reference, if microphysical nihilism is true.

One class of predicates requires a little more discussion. Thus I turn to (RN3):

(RN3) Mereological terms such as ‘parthood’ and ‘composition’ are defined by PARTHOOD, COMPOSITION, etc.

Consider ‘is a part of’, for example. According to the causal theory, this predicate was (like any other predicate) introduced by ostension of one or more samples. Or more accurately, since ‘is a part of’ is a two-place predicate, it was introduced by ostension of one or more *ordered pairs* of samples: in effect (and in realist phraseology), ‘is a part of’ was introduced as picking out the relation that *this* bears to *that*. The descriptive content the introducer of ‘is a part of’ associated with the predicate then served to specify exactly which of the many relations that *this* bears to *that* is to be the referent of ‘is a part of’, by specifying in what respect other ordered pairs must be similar to the sampled ordered pair to fall under the extension of ‘is a part of’. Suppose that in this case our predicate-introducer ostended (what they considered to be) a table-leg and a table when they introduced ‘is a part of’. Given microphysical nihilism, what they really ostended was a plurality of simples* arranged table-leg-wise and a plurality of simples* arranged table-wise. They thus determined that ‘is a part of’ applies to those ordered pairs whose members stand in a similar relation to the relation borne by the plurality of simples* arranged table-leg-wise to the plurality of simples* arranged table-wise. Together with the descriptive content our predicate-introducer associated with ‘is a part of’, that determines that ‘is a part of’ will apply to: the ordered pair containing a particular plurality of simples* arranged wheel-wise and a particular plurality of simples* arranged car-wise; the ordered pair containing a certain plurality of simples* arranged flipper-wise and a certain plurality of simples* arranged penguin-wise, and so on. Indeed, it determines that the extension of ‘is a part of’ will be all and only those ordered pairs such that the first member is/are *among* the second member. In realist terminology, then, ‘is a part of’ expresses the *amonghood* relation. This is what PARTHOOD says.

I don’t see any reason why this result wouldn’t generalise nicely to other mereological predicates. Given microphysical nihilism, then, the causal theory seems to vindicate (RN3).

Finally, I turn to (RN1):

(RN1) ‘Existence’ and its cognates express the plural existential quantifier, $\exists xx$.

At the end of §4.2.1 I conceded that I’m not sure what the right causal theoretic account of how ‘existence’ is associated with a meaning is. But I noted that any plausible account (causal theoretic or not) of the meaning of ‘existence’ must validate the following conditional:

⁸² Of course, the descriptive content will have to be specific enough to allow that e.g. pluralities of simples* that are jointly a different *shade* of red to the sampled pluralities still fall under the extension of ‘is red’.

If '*F*' applies or refers, then '*F*s exist' is true.

Now, given what I've said so far, the causal theorist should accept that for most sortal terms, '*F*', '*F*' applies to (proper) *pluralities* of simples*. For example, above I said that, on the assumption of microphysical nihilism, 'giraffe' (which we can take here to be shorthand for the sortal predicate 'is a giraffe') has in its extension pluralities of simples* arranged giraffe-wise. Thus, if microphysical nihilism is true, then 'giraffe' applies to pluralities of simples*. Now, suppose that 'existence' expresses the *singular* existential quantifier. Then the above conditional is *false*, given microphysical nihilism, for on this assumption 'giraffe' applies (i.e. 'giraffe' has a non-empty extension) but 'giraffes exist' would be false, as there is no *single entity* that is a giraffe. Since we said that any acceptable account of the meaning of 'existence' must validate the above conditional, this tells us that 'existence' can't express the singular existential quantifier, at least on the assumption of microphysical nihilism.⁸³ On the other hand, if we take 'existence' to express the *plural* existential quantifier, i.e. $\exists xx$, then the conditional comes out true, because 'giraffes exist' would then express the claim that there is at least one (proper) *plurality* in the extension of 'giraffe', which is of course true, given what I've said so far. More generally, since the causal theorist is (as I've argued) committed to the idea that sortal terms of the form '*F*' can apply to (proper) pluralities, in order to preserve the truth of the above conditional, it seems that she should accept that if microphysical nihilism is true then, as (RN1) says, 'existence' expresses the *plural* existential quantifier, rather than the singular one.

Of course, nothing that I've said rules out the idea that, if microphysical nihilism is true, then 'existence' expresses something different to either the singular existential quantifier *or* the plural existential quantifier. But, to the extent that the only real plausible candidates for being the meaning of 'existence' are these two quantifiers (perhaps because they are the only two quantifiers that have the correct inferential role), my argument above seems to provide good evidence via disjunctive syllogism for the truth of (RN1), on the assumption of microphysical nihilism.

That concludes my positive argument using the causal theory for (RN1)-(RN5), based on the assumption of microphysical nihilism. I proceed now to arguing for (RN6).

4.3 – Names

Here again is (RN6):

- (RN6) Referring expressions of the form '*N*' that purport to refer to a single composite *F* refer plurally to *xx* iff *xx* are arranged *N*-wise and *xx* are an object.

In order to use the causal theory to argue that (RN6) follows from the assumption of microphysical nihilism, we need first to get the causal theoretic account of how (proper) names come to have a reference on the table. I'll first (§4.3.1) introduce the causal theory; then (§4.3.2) I'll argue that the causal theory motivates (RN6).

⁸³ Of course, it remains open to philosophers to *stipulate* that by 'existence' they mean the singular existential quantifier. But in so doing, they'd be using 'existence' in a way that invalidates the conditional in the main text, and thus cannot claim to be using 'existence' in the standard way. It is the standard notion of existence, rather than a gerrymandered philosophical term of art, that I am interested in here.

4.3.1 – The Causal Theory of the Reference of Names

I'll begin by outlining a simple causal theoretic account (based on Kripke [1980]), before discussing a couple of ways in which it must be made somewhat more complex.

As in the case of the causal theoretic account of predicates, the causal theoretic account of names involves two components: an account of reference-fixing, and an account of reference-borrowing. On a simple causal theoretic account, the reference of a name is fixed by ostension, and that name then simply refers to the ostended entity for as long as that entity exists. When a name is introduced to refer to an entity, we say that that entity has been *baptised*. Subsequent speakers can then acquire or *borrow* the name from the baptiser (or from someone else who has borrowed the name from the baptiser, and so on), and their utterances of that name refer to the same entity that the baptiser baptised.⁸⁴ Speakers who borrow names in this way are said to be appropriately causally connected to the baptism of the entity that their utterance of a name refers to.

With this in mind, then, we can summarise a simple causal theoretic account of names with the following set of necessary and sufficient conditions for an utterance of a name, *n*, to refer to a particular entity, *e*:⁸⁵

An utterance of a name, *n*, refers to an entity, *e* iff:

- (1) *e** was baptised *n*, and
- (2) The utterance of *n* is appropriately causally connected to the baptism of *e**, and
- (3) *e* is identical to *e**

For example, suppose that some scientist in the Antarctic once came across a particular entity, *e**, gestured at it, and said or thought something to effect of, "I'm going to call you 'Pingu'". This amounts to the satisfaction of (1). The scientist then told her friends about Pingu, and her friends mentioned it in conversation with their friends, and the name is spread further and further in this sort of way, until someone tells *me* about Pingu. I then utter a sentence using the term 'Pingu'. Since the name 'Pingu' was passed to me via a person-to-person chain originating at the initial baptism, (2) is satisfied. Now if there's some entity *e* which is identical to *e**, then my utterance of 'Pingu' refers to *e*.

However, this simple causal theoretic account cannot be right, because the *qua* problem that we encountered in the previous section arises again here (see e.g. Papineau 1979: 158-168, Devitt 1981, and Devitt and Sterelny 1999: 79-81). The problem is that the act of ostension at the moment of baptism typically doesn't determine which of many candidate referents is being baptised. For instance, in the example just given, it is plausibly indeterminate whether the gesture that accompanied my introduction of the name 'Pingu' thereby baptised a certain *penguin* (as was my intention), or a certain cloud of air particles in between me and the penguin, or just the penguin's *flipper*, etc. It seems objectionable to say that 'Pingu' is to name (perhaps the mereological fusion of) *all* of these things; it

⁸⁴ Plausibly, we should also require that the reference-borrower must intend to use the name in the same way as the person from whom they acquired it used it (see Kripke 1980: 96). We can safely leave this detail implicit here.

⁸⁵ Kripke himself shied away from presenting his causal theory as a set of necessary and jointly sufficient conditions for reference (1980: 64; 93), preferring to give a mere 'characterisation [... i.e.] just a better picture than the picture presented by received views' (93). But the three conditions in the text capture the spirit, at least, of the picture he was painting.

seems equally objectionable to say that the reference of 'Pingu' is indeterminate between a penguin, a cloud of particles, a flipper, etc., or that 'Pingu' fails to refer. Rather, given that I gestured at an entity that I took to be a penguin and intended to baptise that entity 'Pingu', the intuitively correct thing to say is that I thereby successfully introduced 'Pingu' to refer to that entity. The simple causal theory outlined above apparently cannot secure this result.

As in the case of sortals and predicates, it seems that to solve this problem the causal theorist should allow that, in the moment of baptism, the introducer of a name not only gestures at some entity but also associates the name with some descriptive content that serves to determine which of the many candidate things they may have been gesturing at (the penguin, the cloud of particles, the flipper, etc.) is to be the genuine baptisee (e.g. Devitt and Sterelny 1999: 79-81, Thomasson 2007: 39-44). One common idea, for example, is that the introducer of a name conceives of the thing they intend to baptise as falling under a certain sortal predicate, and that that disambiguates which of the several possible candidates they baptise (e.g. Devitt and Sterelny 1999: 80). For example, when I introduced 'Pingu' with a gesture at either a penguin, or a cloud of particles, or a flipper, I conceived of the object of my gesture as falling under the sortal predicate 'is a penguin', which determined that I was in fact introducing 'Pingu' to refer to the *penguin*, rather than the cloud of particles or the flipper.

But we need not take a stand here on how exactly the causal theorist's account of how descriptive content enters into the account of reference-fixing should go. Rather, it's dialectically appropriate here to simply assume that the causal theorist has some plausible story to tell according to which descriptive content of some sort on the part of the introducer of a new name serves to successfully disambiguate what is being baptised (c.f. the discussion of the dialectic at the end of §4.1).

Finally, on this point, a side-effect of this sort of account that we should take note of, again as in the case of sortals, is that if *none* of the candidates for being the thing that I ostended during the introduction of a term satisfy the descriptive content associated with that term, then the term fails to refer in the first place (e.g. Devitt and Sterelny 1999: 79-80). For if we are to say that the descriptive content of *is a penguin* serves to rule out that anything other than the penguin that I am gesturing at is to be the referent of the name I am introducing, then we ought similarly to say that some descriptive content associated with a name that is not met by *any* of the candidate referents rules out any of those candidates being the referent of the name.

Before turning to deploying the causal theory of names in support of (RN6), we need to consider one further complication, which is that the causal theory will work slightly differently depending on what theory of *persistence* we assume. I'll now elaborate on this with regard to the three most prominent accounts of persistence: endurantism, perdurantism, and stage theory.

First, suppose we endorse *endurantism* (see e.g. van Inwagen 1990b). On this view, objects are three-dimensional (extended in space, not in time), and persist simply by moving through time. If a certain object, *o*, persists from t_1 to t_2 , then that object is *wholly present* at t_1 and t_2 ; the *o* that exists now, at t_2 , is numerically identical to the *o* that existed at t_1 . If endurantism is true, then the causal theory as I've outlined it so far doesn't need any special modification (at least as it pertains to *singular* reference – below I argue that the causal does need to be modified in a certain way to handle cases of *plural* reference).

Second, suppose we endorse a version of four-dimensionalism known as *perdurantism* (see e.g. Lewis 1986). On this view, objects are *four-dimensional*, in that they are not only extended in space, but also in time. Thus, just as the vast majority of objects do not wholly exist at any single point in space, but

rather occupy multiple points in space in virtue of having spatial parts at those points, so the vast majority of objects do not wholly exist at any single moment in time, and rather occupy multiple points in time in virtue of having *temporal* parts at those times. If a certain object, *o*, persists from t_1 to t_2 , then it does so by having a temporal part at t_1 and another temporal part at t_2 (and at any time in between). These temporal parts are not numerically identical to one another, though they of course both underlap (i.e. are both parts of) the same object. Objects that are temporally extended in this way can be called *spacetime worms*, to signify how they are spread out in both space and time. Now, importantly (for the point I want to make at least) it's usually supposed that any given temporal part is a part of a vast number of distinct spacetime worms. For perdurantism is usually combined with universalism, and as such the usual version of the view is committed to there being as many spacetime worms as there are ways of combining temporal parts (e.g. Hawley 2020: §8).⁸⁶ Thus consider a simple world that has two moments, t_1 and t_2 , and four temporal parts, b_1 , b_2 , e_1 , and e_2 , where b_1 and e_1 exist at t_1 , and b_2 and e_2 exist at t_2 . This world contains Big Ben, which is a spacetime worm composed of its temporal parts, b_1 and b_2 , and the Eiffel Tower, which is a spacetime worm composed of its temporal parts, e_1 and e_2 . Given universalism, this world also contains two further spacetime worms: the one composed of b_1 and e_2 , and the one composed of e_1 and b_2 . Thus e_1 , for example, is a temporal part not only of the Eiffel Tower, but also of the gruesome spacetime worm composed of e_1 and b_2 .

The upshot of this for the causal theory is as follows. On perdurantism, when we ostend an object in order to introduce a name for it, we are in almost all cases naming the *spacetime worm* of which that object is a temporal part.⁸⁷ Thus the causal theorist's account needs to explain both how a given name comes to refer to a certain spacetime worm rather than to the temporal part that was ostended, and how the given name comes to refer to the *right* spacetime worm, given that the ostended temporal part is a part of a vast number of spacetime worms. For example, if I gesture at a temporal part of a certain tower, e_1 , and baptise it 'the Eiffel Tower', then the causal theorist's account needs to explain how 'the Eiffel Tower' comes to refer to the Eiffel Tower, i.e. the spacetime worm that has e_1 and e_2 as temporal parts, rather than referring either to just e_1 or to the spacetime worm composed of e_1 and b_2 . The point here is *not* that perdurantism raises an insurmountable problem for the causal theorist. For, given what we've said so far, the causal theorist has an easy response to the suggestion that her theory will not work, given perdurantism: she can say that the descriptive content that the introducer of a given name associates with that name at the moment of baptism not only serves to specify exactly which presently-existing *temporal part* was gestured at (e.g. the present temporal part of a penguin, or of a flipper, or of a cloud of particles), but also which of the many spacetime worms of which the ostended temporal part is a part is to be the referent of the name. Again, I'll simply assume that, if perdurantism is true, then this sort of causal theoretic account will be workable.

The third major view of persistence is another kind of four-dimensionalist view. It is known as stage theory (see Sider 1996, Sider 2001, Hawley 2001). This view accepts the same metaphysical picture as that of perdurantism, complete with its ontology of spacetime worms, but denies that those spacetime worms are ordinary objects like mugs and penguins. Rather, ordinary objects are identified with instantaneous temporal parts, or *stages*, of those spacetime worms. Though these stages only

⁸⁶ To be precise, by 'universalism' here I mean what Sider (2001) calls *unrestricted diachronic composition*, which is essentially a form of universalism (i.e. the view that all disjoint objects compose) that is explicitly stipulated to apply to objects existing at different times to one another.

⁸⁷ The exceptions are cases in which we baptise an object that only exists for an instant, and thus is not a spacetime worm. For example, it seems at least possible to name an instantaneous temporal part, rather than the spacetime worm of which it is a part.

exist for an instant, they may nevertheless bear what we can call the *temporal counterpart relation* to stages that exist at other times. As with the ordinary (modal) counterpart relation, the temporal counterpart relation can be understood in terms of similarity, or perhaps continuity. Thus if a stage existing at t_2 is sufficiently *continuous* in the right way with a stage that existed at t_1 , then the earlier stage bears the temporal counterpart relation to the later stage. Stage theorists then analyse persistence in terms of this temporal counterpart relation. Consider the Eiffel Tower at t_1 , which on stage theory is an instantaneous stage, e_1 . At t_2 , e_1 no longer exists, but an extremely similar stage, e_2 , is there in its place. e_1 in this case seems to bear the temporal counterpart relation to e_2 , and this is sufficient, say stage theorists, for it to be true that the Eiffel Tower has persisted from t_1 to t_2 .

An important feature of stage theory is that, as is apparently true in the modal case (see Lewis 1971), there is a multiplicity of temporal counterpart relations (see e.g. Sider 1996: §5). That is, there are a number of different ways of spelling out what it is for a later stage to be continuous with an earlier one, and each such way corresponds to a different temporal counterpart relation. Which temporal counterpart relation is at issue in a claim concerning persistence is contextually determined. For example, in ordinary claims concerning the persistence of persons, it seems that the relevant type of continuity is psychological continuity, such that a person-stage existing at t_1 bears the relevant temporal counterpart relation to the person-stage existing at t_2 iff the second person-stage is sufficiently psychologically continuous with the first person-stage. In ordinary claims concerning the persistence of the Eiffel Tower, however, this standard of continuity seems inapt, and a different one is relevant: perhaps continuity of spatial location, form, cultural significance, etc.

How does stage theory interact with the causal theory of reference? On perdurantism we said that when the name ‘the Eiffel Tower’ (for example) is introduced via ostension of e_1 , it names the spacetime worm composed of e_1 and e_2 ; on stage theory, when ‘the Eiffel Tower’ is introduced in this way it names only the presently existing e_1 . Now, given the *simple* causal theory that I outlined above via (1)-(3) (together with the additions to that theory that were made to avoid the *qua* problem), as time moves on from t_1 to t_2 and e_1 ceases to exist, ‘the Eiffel Tower’ will not refer to *anything* any more. For condition (3) of the simple causal theory requires that the present referent of ‘the Eiffel Tower’ be *identical* to the entity that was baptised ‘the Eiffel Tower’, and at t_2 that entity doesn’t exist any more. More generally, since baptised entities are said on the stage theory to be momentary *stages* that go out of existence as the next instant arrives, on the simple causal theory outlined above, names only refer in the moment that they are introduced, and then cease to refer after that. This is of course extremely counterintuitive.

So much the worse for the simple causal theory: it seems that we should reject that (1)-(3) (plus the additions to that causal theory that were made to avoid the *qua* problem) adequately characterise the causal theory of reference if stage theory is true. Rather, since stage theorists deny strict and literal identity over time, and in its place posit *temporal counterpart relations* holding between non-identical stages, it seems thematic, on the assumption that stage theory is true, to modify the simple causal theory so that it reads as follows:

An utterance of a name, n , refers to an entity, e iff:

- (1) e^* was baptised n , and
- (2) The utterance of n is appropriately causally connected to the baptism of e^* , and
- (3*) e^* bears the salient temporal counterpart relation to e (which is a presently existing stage).

Thus suppose again that ‘the Eiffel Tower’ was introduced as the name for e_1 . As time moves on, e_1 ceases to exist, but e_2 comes into existence in its place. e_2 is the stage to which e_1 bears the salient temporal counterpart relation. So now, on stage theory, ‘the Eiffel Tower’ refers not to e_1 , but to e_2 .

This introduces a complication for the causal theorist: what is it that determines what the salient temporal counterpart relation is? Suppose that I succeed in introducing ‘the Eiffel Tower’ as the name for e_1 . When e_1 goes out of existence, and when e_2 comes into existence, what is it that determines that the salient temporal counterpart relation is the one that e_1 bears to e_2 , rather than e.g. to b_2 , or perhaps to nothing at all? Again, the point here isn’t that the causal theorist cannot meet this challenge, and again, her answer to the *qua* problem seems to provide her with a way to do so. For the causal theorist can insist that the descriptive content that the introducer of a name associates with that name at the moment of baptism serves not only to determine exactly which object is being baptised, but *also* to make salient a particular temporal counterpart relation. In this way, causal theorists can apparently render their theory of reference compatible with stage theory.

That concludes my exegesis of the causal theory of reference as it pertains to names. As we’ve seen, the simple causal theory that I presented at the beginning of the section via (1)-(3) is inadequate, as it is subject to the *qua* problem. I’ve additionally discussed how the causal theory must be extended or modified so as to be consistent with different theories of persistence.

4.3.2 – Applying the Causal Theory to Names on the Assumption of Microphysical Nihilism

I turn now to arguing that, on the assumption of microphysical nihilism, the causal theory as I presented it in the previous section supports (RN6):

- (RN6) Referring expressions of the form ‘N’ that purport to refer to a single composite F refer plurally to xx iff xx are arranged N -wise and xx are an object.

Given my remarks in the previous section concerning the interaction between the causal theory and different theories of persistence, exactly how my argument must be stated will depend on facts about whether endurantism is true, or whether one of the two forms of four-dimensionalism is true. To be as general as possible, I’ll consider all three possibilities.

4.3.2.1 – On the Assumption of Perdurantism

First, let’s assume that perdurantism is true. Now, one might think that this assumption is incompatible with our starting assumption of microphysical nihilism. For perdurantism is committed to the existence of temporal parts, whereas microphysical nihilism is usually supposed to be inconsistent with there being parts of any kind, temporal or not (see e.g. Long 2019: §3.2). However, whilst microphysical nihilism is by definition inconsistent with there being *parts** of any kind, it is not inconsistent with there existing *parts*, including temporal parts. For recall that ‘part’, as I’m understanding it here at least, is defined as follows:

PARTHOOD: xx is/are a part of yy =_{df} xx are among yy

Parthood, so defined, is fully compatible with microphysical nihilism. Thus when I say that we should assume that perdurantism is true, we should understand that assumption as committing the nihilist only to the existence of temporal parts, and *not* to the existence of temporal *parts**. On this understanding of perdurantism, there exist temporally extended pluralities of simples* (which we

might helpfully refer to as pluralities of simples* arranged spacetime-worm-wise), but no composite* spacetime worms.⁸⁸

To make things even clearer, it may be helpful to provide the definition of temporal parts that the reductive nihilist is employing here. Just as the reductive nihilist takes parthood to equate to subplurality-hood (see §2.2.2), she takes temporal parthood to equate to temporal subplurality-hood:

TEMPORAL SUBPLURALITY: xx are a temporal subplurality of $yy =_{df} (1) xx$ is the largest plurality containing only entities that exist at a particular time and that are among yy .

Now, as we've seen, if perdurantism is true then causal theorists say that names are introduced via ostension of a present temporal part, and that the descriptive content the introducer of the name in question associates with it determines that (in typical cases at least) the name is to refer to a particular spacetime worm of which the ostended object is a temporal part. If microphysical nihilism is true, then we can redescribe this as: what is ostended is as a matter of fact a certain presently existing plurality of simples*, and the descriptive content the introducer of the name in question associates with it determines that (in typical cases at least) the name is to refer to a particular plurality of simples* arranged spacetime-worm-wise of which the ostended plurality is a temporal subplurality.

For example, then, consider a simple universe containing two instants, t_1 and t_2 . At t_1 , there exists a plurality of simples* arranged Eiffel-Tower-wise (call them ee_1) and a plurality of simples* arranged Big-Ben-wise (bb_1); at t_2 there also exists a plurality of simples* arranged Eiffel-Tower-wise (ee_2) and a plurality of simples arranged Big-Ben-wise (bb_2). At t_1 , the name 'the Eiffel Tower' was introduced with a gesture towards ee_1 .⁸⁹ Now, ee_1 is a temporal subplurality of a variety of pluralities of simples* arranged spacetime-worm-wise: notably, it is a temporal subplurality of the plurality of simples* composed of ee_1 and ee_2 (call it ee_1+ee_2) and of the plurality of simples* composed of ee_1 and bb_2 (call it ee_1+bb_2). According to the causal theory, the descriptive content associated by the introducer of 'the Eiffel Tower' then serves to determine to which plurality of simples* arranged spacetime-worm-wise 'the Eiffel Tower' is to refer: i.e. it serves to determine in this case (presumably) that the name is to refer to ee_1+ee_2 rather than to ee_1+bb_2 or any other plurality of simples* arranged spacetime-worm-wise. We can call ee_1+ee_2 the plurality of simples* arranged Eiffel-Tower-spacetime-worm-wise, as it is the plurality formed out of all of the temporal subpluralities of simples* arranged Eiffel-Tower-wise.

I think this example generalises in obvious ways to most other instances of the introduction of a name that purports to refer to a single composite object, given microphysical nihilism. That is, in most other

⁸⁸ Perhaps it will be objected that perdurantism must by definition be understood as implying the existence of temporal parts*, rather than just temporal parts. I need not contest the point here. Those who are sympathetic to this objection may take my use of 'perdurantism' to pick out a theory that is closely related to perdurantism proper, and that differs from perdurantism proper only in that it denies the existence of temporal parts* (and thus of composite* spacetime worms, etc.) and accepts in their place temporal parts (and composite spacetime worms, etc.).

⁸⁹ Recall that we are assuming that the descriptive content the name-introducer associated with 'the Eiffel Tower' served to rule out that she was instead ostending just e.g. the plurality of simples* arranged bottom-half-of-the-Eiffel-Tower-wise, or a plurality of simples* arranged cloud-of-air-particles-wise, and so on. One might be sceptical that descriptive content could determine that the name-introducer was ostending ee_1 as opposed to e.g. *the same plurality minus one simple**. But this problem is not unique to nihilism (e.g. given universalism, we might be sceptical about whether descriptive content can determine that it is a particular composite* object that was gestured at during the introduction of 'the Eiffel Tower', as opposed the object composed* of the original composite* object plus one more simple*, etc.), so it's dialectically appropriate to ignore it here (c.f. my discussion of Unger's *Problem of the Many* at the end of §3.4.3).

cases, a name, 'N', that purports to refer to a single composite F will refer to the temporally extended plurality of simples arranged N -spacetime-worm-wise.

However, there are some cases to which the above worked example does *not* seem to generalise. These are cases in which a speaker has attempted to introduce a name for a single composite object, but in which the plurality of simples* that they have in fact ostended in the moment of baptism do not qualify as a single object. In such cases, I submit, the attempt to introduce the name *fails*. The reason for this has to do with the fact that, as I noted in §4.3.1, on the version of the causal theory that is modified so as to avoid the *qua* problem, if the descriptive content associated with a name at the moment of baptism is not satisfied by any of the candidates for being the thing that I am gesturing at, then the attempt to specify a referent for the name has failed. In cases in which speakers are attempting to introduce names for a putative single composite objects, it seems that one element of the descriptive content that those speakers associate with those names is that their referents are each single objects: after all, I've stipulated that we're here discussing names that purport to refer to single composite F s, which suggests that the introducers of these names regard them as being single composite F s of some variety or other, i.e. single composite *objects* that fall under some sortal term F or other. So it seems plausible that a given plurality of simples* arranged spacetime-worm-wise can only be successfully baptised with a name that purports to refer to a single composite F if the temporal subplurality of that plurality of simples* arranged spacetime-worm-wise counts as a single object: if the relevant temporal subplurality *doesn't* count as a single object, then the attempt to fix a reference to the simples* arranged spacetime-worm-wise will fail.⁹⁰

To illustrate this, consider again the example I introduced in §4.2.2 of the attempt to introduce a sortal predicate via ostending what appeared to be a giraffe, but was instead a nearby horse's body and a distinct crane on a building site. This time, let's suppose that our speaker attempts to introduce a name, 'Jill', for what they take to be the giraffe by gesturing at what is in fact the nearby plurality of simples* arranged horse's-body-wise and the distant plurality of simples* arranged crane-wise. Since our speaker intended to name a single *giraffe*, one element of the descriptive content they associated with 'Jill' was that it is to name a single object that falls under the extension of the sortal predicate 'is a giraffe'.⁹¹ Since the nearby plurality of simples* arranged horse's-body-wise and the distant plurality of simples* arranged crane-wise are not jointly a single object, they do not satisfy this element of the descriptive content, and so cannot be the referents of 'Jill'. On the plausible assumption that the descriptive content associated with 'Jill' also serves to rule out all other candidate referents (such as the nearby horse, or the distant crane), we should conclude that the attempt to fix the reference of 'Jill' has failed in this case.

Here's a final plausible consequence of the fact that one element of the descriptive content associated with names that purport to name single composite F s is that those names are to refer only to single objects: if those names refer, then every temporal subplurality of the plurality of simples* arranged spacetime-worm-wise to which they refer must count as an object. Consider again the case of 'the

⁹⁰ One might think that the descriptive content the introducers of names associate with those names is instead such that the name in question can only refer to a single *entity*, rather than only to a single *object*. But I think this is implausible, because 'entity' is a technical term of metaphysics that ordinary speakers do not typically employ. See fn 79 of this chapter as well as §3.5 in the previous chapter for more discussion.

In §4.4 I'll consider and rebut the objection that name-introducers instead associate descriptive content with the names they introduce that prevent them from referring to anything other than single *entities*.

⁹¹ Indeed, the descriptive content only needs to include the claim that 'Jill' is to refer to a (single) *giraffe*, as I've already argued in §4.2.2 that some xx can only be in the extension of a sortal predicate such as 'is a giraffe' if xx is a single object.

Eiffel Tower' above, which we said referred to ee_1+ee_2 in a universe containing only two instants. Suppose we add a third instant to that universe, t_3 , and stipulate that by t_3 all of the simples* in that universe have been scattered so each simple* is located a great distance from all of the others. Let the variable xx range over the pluralities of simples* existing at t_3 . Intuitively, the descriptive content associated with 'the Eiffel Tower' by its introducer at t_1 determines that at t_3 'the Eiffel Tower' doesn't refer to ee_1+ee_2+xx , for *any* value of xx : that is, intuitively the descriptive content associated with 'the Eiffel Tower' determines that the plurality referred to by 'the Eiffel Tower' has no temporal subplurality at t_3 . A plausible explanation of this is that the descriptive content associated with 'the Eiffel Tower' requires that *each* of its temporal subpluralities must count as a single object. And again, I think this generalises to other names that purport to name single composite *F*s.

Given all of this, then it seems plausible to conclude the following: if microphysical nihilism is true, and if perdurantism is true, then names of the form 'N' that purport to refer to a single composite *F* refer plurally to xx iff (i) xx are arranged *N*-wise, and (ii) each temporal subplurality of xx counts as a single object. If we stipulate that, on the assumption of perdurantism, 'arranged *N*-wise' applies to the plurality of simples* arranged *N*-spacetime-worm-wise, and that a plurality of simples* arranged spacetime-worm-wise counts as an object as long as each of its temporal subpluralities does, then we arrive at (RN6): referring expressions of the form 'N that purport to refer to a single composite *F* refer plurally to xx iff xx are arranged *N*-wise and xx are an object.

4.3.2.2 – On the Assumption of Stage Theory

Now suppose that stage theory is true. As we've seen, if stage theory is true, then causal theorists say that names are introduced via ostension of a present stage, which is thereby fixed as the present reference of the name. At later times, the name refers to the stage existing at that time to which the original stage (i.e. the one that was ostended at the moment of baptism) bears the salient temporal counterpart relation. Which temporal counterpart relation, and thus which later stage is to be the referent of the name, is determined by the descriptive content associated with that name by its introducer. If microphysical nihilism is true, then the stage that is ostended during the introduction of a name will typically be a plurality of simples*, and the descriptive content associated with that name will determine that at other times that name refers to other pluralities of simples* to which the baptised plurality bears a certain salient temporal counterpart relation.⁹²

With this in mind, consider again the case of the two-instant universe containing ee_1 (i.e. the simples* arranged Eiffel-Tower-wise) and bb_1 (the simples* arranged Big-Ben-wise) at t_1 , and ee_2 and bb_2 at t_2 . When our name-introducer introduces the name 'the Eiffel Tower' at t_1 , they do so by ostending ee_1 . Together with the descriptive content that serves to disambiguate what they are gesturing at, they thereby name what they conceive of as being a certain tower, which is in fact a certain plurality of simples* arranged Eiffel-Tower-wise, i.e. ee_1 . When time ticks over from t_1 to t_2 , ee_1 goes out of existence (as does bb_1), and ee_2 and bb_2 come into existence. Since the introducer of the name 'the Eiffel Tower' conceived of what they baptised at t_1 as being a *tower*, it's plausible that they thereby made salient the *tower*-temporal-counterpart-relation. What's more, it's plausible that ee_1 bears that tower-temporal-counterpart-relation to ee_2 , since ee_2 are arranged tower-wise, are spatiotemporally continuous with ee_1 , etc. Thus the descriptive content the introducer of 'the Eiffel Tower' associated

⁹² It's worth noting that the temporal counterpart relation in this sort of case will be a *plural* one that is non-distributive in both places. That is, it is a temporal counterpart relation that relates one plurality, xx , to another plurality, yy , without implying that each member of xx also bears that temporal counterpart relation to each member of yy .

with the name at t_1 determines that at t_2 ‘the Eiffel Tower’ names ee_2 , i.e. the simples* arranged Eiffel-Tower-wise.

Furthermore, as discussed in the previous section, I think it’s plausible that the descriptive content our name-introducer associated with ‘the Eiffel Tower’ determined that ‘the Eiffel Tower’ may only refer to a single *object*, both at the moment of baptism and at later times.

I see no reason to doubt that all of this generalises: in general, it’s plausible that, given all three of microphysical nihilism, the stage theory, and the causal theory of reference, names of the form ‘N’ that purport to refer to single composite *F*s refer plurally to xx iff xx are arranged *N*-wise and xx are an object. That is (RN6).

4.3.2.3 – On the Assumption of Endurantism

Now suppose that endurantism is true. Given endurantism, causal theorists say that the introducer of a name ostends an entity (perhaps with the help of some disambiguating descriptive content), and that name then refers to that very entity throughout its existence. In other words, causal theorists say that the introducer of a name ostends an x , and then at all subsequent times the name refers to y iff y is identical to x .

One might think, then, that if what was ostended during the introduction of a name was (due to the truth of microphysical nihilism) really a plurality of simples*, xx , then at all subsequent times that name refers to yy iff yy are identical to xx . If this is right, though, then it would have very counterintuitive consequences for the microphysical nihilist: more to the point, perhaps, if this is right then it would invalidate, rather than vindicate, (RN6).

To see why, first recall from §2.2.2 that the following definition of plural identity:

PLURAL IDENTITY: $xx = yy =_{df}$ for all z , z is one of xx iff z is one of yy .⁹³

Given PLURAL IDENTITY, two pluralities are identical to one another iff they have exactly the same members as one another. Now suppose that at t_1 I have introduced the name ‘the Eiffel Tower’ by ostending the plurality of simples* arranged Eiffel-Tower-wise, thereby determining that ‘the Eiffel Tower’ is the name for that plurality of simples*. Suppose that this plurality has as its members 100 simples*, s_1, \dots, s_{100} . As time ticks over to t_2 , s_1, \dots, s_{100} still exist (as endurantism is true) but are no longer arranged Eiffel-Tower-wise: given the nature of quantum reality, a certain number of s_1, \dots, s_{100} will have shot off in various directions, and their roles in the Eiffel-Tower-wise arrangement will have been occupied by numerically *different* simples*. That is, at t_2 there still *is* a plurality of simples* arranged Eiffel-Tower-wise, but it is not the plurality made up of s_1, \dots, s_{100} : rather, supposing for example that *ten* of the original simples* have shot off, the plurality of simples* that are arranged Eiffel-Tower-wise at t_2 might be s_{11}, \dots, s_{110} . However, since ‘the Eiffel Tower’ was introduced as the name for s_1, \dots, s_{100} , and since s_1, \dots, s_{100} still exist at t_2 , it appears that we have to say that at t_2 ‘the Eiffel Tower’ still refers to s_1, \dots, s_{100} . For given what we’ve said, at t_2 ‘the Eiffel Tower’ refers to the yy that is identical to s_1, \dots, s_{100} , and (given PLURAL IDENTITY) that must be s_1, \dots, s_{100} . This is not only counterintuitive, but it violates (RN6): at t_2 , s_1, \dots, s_{100} are not arranged Eiffel-Tower-wise, but ‘the Eiffel Tower’ refers to them; furthermore, at t_2 , s_{11}, \dots, s_{110} are arranged Eiffel-Tower-wise (and are jointly an object), but ‘the Eiffel Tower’ does not refer to them.

⁹³ In §2.2.2 I formulated the definition of plural identity differently, but I noted that this formulation is equivalent to the one I gave.

It is premature, however, to conclude that, if endurantism is true, then the causal theory implies the falsity of (RN6). For the above line of reasoning only goes through if the version of the causal theory that we should adopt if endurantism is true functions in the same way for *singularly referring* names as it does for names that *plurally* refer. In the singular case, I said in §4.3.1 that, given endurantism and the causal theory, a name is introduced by ostension, and then at future times refers to whatever it is identical to the ostended entity. On the assumption that the causal theory works in the same way when what is ostended is in fact a *plurality* of entities, the above reasoning goes through. But there's good independent reason for thinking that, given endurantism, the causal theory works somewhat differently in cases where what is fixed by ostension to be the reference of a name is a plurality.

This independent reason can be appreciated by considering how an ordinary case of plural reference works, given endurantism, and given the causal theory (in this example, then, I drop the assumption of microphysical nihilism).

Thus consider, for example, the term 'Oasis'. Suppose that this term was introduced in 1991 by Noel Gallagher, who gestured at himself and the four other musicians he was performing music with and said something to the effect of, 'from now on, we're called 'Oasis', yeah?'. Thus Noel introduced a plurally referring term, 'Oasis', that had as its five referents Noel and Liam Gallagher, Bonehead, Guigsy, and Tony McCarroll.⁹⁴

Now, fast-forward to 2005. Given the nature of endurantism, each of the original members of Oasis still exists, but given the nature of rock and roll bands, several of those original members have left the band. In particular, Bonehead, Guigsy, and Tony McCarroll have now left the band, and have been replaced by Gem and Andy Bell. Intuitively, when I utter the name 'Oasis' now, in 2005, what does it refer to? One option is to say that it still refers to the original five members who were baptised 'Oasis' in 1991. But that doesn't seem like the intuitively correct thing to say: rather, the intuitive thing to say is that, in 2005, 'Oasis' now refers to Noel, Liam, Gem, and Andy Bell.

What conclusion should we draw from this? I think we should conclude that, given endurantism, it is *not* the case that, for a plural name introduced at t_1 to refer to a certain plurality, xx , that plural name refers at t_n to yy iff yy are identical to xx . For, as we've just seen, that would get the intuitively wrong result in the example of 'Oasis'.⁹⁵

How should we do this? Happily, I think the causal theorist has a ready-made answer. For we have already seen in the previous subsection how the causal theorist's account of naming works given *stage theory*. And by simply endorsing that same account here, on the assumption of *endurantism*, it seems

⁹⁴ One might dispute the idea that 'Oasis' is in this case a plurally referring term. Perhaps 'Oasis' instead names a single entity, namely a *band*. I can't hope to refute this suggestion here. But here are three considerations that I think warrant my treating 'Oasis' as a plural name. First, it's common (in British English at least) to match band names with plural verb forms, and seems somewhat ungrammatical to match them with singular verb forms. For example, the sentence 'Oasis are playing at Knebworth' seems perfectly grammatical, whereas 'Oasis is playing at Knebworth' does not. Second, the sentence 'Bonehead is one of Oasis' sounds grammatical (to me at least: if it doesn't to you, then try 'Paul is one of The Beatles', and feel free to insert this example in place of mine throughout my discussion), and the 'is one of' predicate standardly takes a plurality in its second argument place. Third, other philosophers have treated band names as legitimate examples of plural names, and infer conclusions about the nature of plural semantics on that basis (see e.g. Hewitt 2012: 866, inc. fn24), so it seems warranted for me to do the same here.

⁹⁵ I hope it's clear that further examples that are structurally similar to that of 'Oasis' could be given in further support of this point (think for example of the name we might give to a crowd of people taking part in a protest, which seems to shift its reference as people leave and join the protest).

that the causal theorist can preserve our intuitions about what plural names refer to in cases such as that of ‘Oasis’.

First, then, given stage theory, causal theorists say (or should say) the following about the case of ‘Oasis’. In 1991, ‘Oasis’ was introduced as the plural name for five people: Noel, Liam, Bonehead, Guigsy, and Tony McCarroll. The name was also associated with some descriptive content that made salient a particular temporal counterpart relation, such that, in 2005, the name ‘Oasis’ refers to whatever it is that the original five people jointly bear this salient temporal counterpart relation to. Plausibly, when Noel introduced ‘Oasis’ in 1991, he associated with it descriptive content that makes salient what we might call the *band*-temporal-counterpart-relation; plausibly, in 2005, the plurality to which the original five members of Oasis jointly bear this band-temporal-counterpart-relation is the plurality containing Noel, Liam, Gem, and Andy Bell. Thus, in 2005, ‘Oasis’ refers to those four people, in line with our intuitions.

Now, although endurantism is a three-dimensionalist view whilst stage theory is a four-dimensionalist one, endurantist causal theorists can preserve the important features of this intuition-preserving account. In 1991, they can say, ‘Oasis’ was introduced as the plural name for the original five band members. It was also associated with some descriptive content that plays a role in determining who the future referents of ‘Oasis’ will be at later times. There’s nothing stopping us phrasing this as follows: ‘Oasis’ was at the moment of baptism associated with some descriptive content that made salient a particular temporal counterpart relation, such that, in 2005, the name ‘Oasis’ refers to whatever it is that the original five people jointly bear this salient temporal counterpart relation to.⁹⁶ When we fast-forward to 2005 and ask what ‘Oasis’ refers to now, there’s nothing to stop endurantist causal theorists from answering: ‘Oasis’ refers to whatever the original five band members from 1991 jointly bear the salient temporal counterpart relation to, which is the plurality containing Noel, Liam, Gem, and Andy Bell. And if they do say this, then endurantist causal theorists can agree with our intuition about what ‘Oasis’ should refer to in 2005.

In general, then, I think there’s independent reason, as displayed by consideration of the case of ‘Oasis’, for accepting the following claim: if endurantism is true, then the causal theoretic account of plural names should be the same as the causal theoretic account of names (both singular and plural) as it is given stage theory.

Finally, let’s apply all of this to the case of microphysical nihilism and ‘the Eiffel Tower’. So suppose, again, that at t_1 I have introduced the name ‘the Eiffel Tower’ by ostending the plurality of simples* arranged Eiffel-Tower-wise, thereby determining that ‘the Eiffel Tower’ is the name for that plurality of simples*. Suppose that this plurality has as its members 100 simples*, s_1, \dots, s_{100} . As time ticks over to t_2 , s_1, \dots, s_{100} still exist (as endurantism is true) but are no longer arranged Eiffel-Tower-wise: given the nature of quantum reality, a certain number of s_1, \dots, s_{100} will have shot off in various direction, and their roles in the Eiffel-Tower-wise arrangement will have been occupied by numerically *different* simples*. That is, at t_2 there still *is* a plurality of simples* arranged Eiffel-Tower-wise, but it is not the plurality made up of s_1, \dots, s_{100} : rather, supposing for example that *ten* of the original simples* have shot off, the plurality of simples* that are arranged Eiffel-Tower-wise at t_2 might be s_{11}, \dots, s_{110} . Which plurality of simples* does ‘the Eiffel Tower’ refer to now, at t_2 ? Given our new and improved version of the causal theory, the answer to that question depends on what

⁹⁶ Note that in saying this the endurantist does not compromise her endurantist analysis of persistence. For the endurantist causal theorist is *not* here saying that the temporal counterpart relation is involved in the analysis of *persistence*: rather, she’s just saying that the temporal counterpart relation is involved in the analysis of what a plural term that was introduced in the past currently refers to.

plurality it is at t_2 to which s_1, \dots, s_{100} jointly bear the salient plural temporal counterpart relation. As we said in the previous subsection in our discussion of the stage theory, since the introducer of ‘the Eiffel Tower’ conceived of it as picking out a *tower*, it’s plausible that the descriptive content they associated with the name made salient the *tower*-temporal-counterpart-relation; and again, as we said in the previous subsection, it’s plausible that the plurality to which s_1, \dots, s_{100} bear the tower-temporal-counterpart relation at t_2 is s_{11}, \dots, s_{110} . Thus, on this account, at t_2 ‘the Eiffel Tower’ refers to s_{11}, \dots, s_{110} . This is the intuitive result.

Since our endurantist causal theoretic account here mirrors the stage-theoretic-causal-theoretic account from the previous subsection, it is safe to conclude here, as we did there, that the descriptive content our name-introducer associated with ‘the Eiffel Tower’ determined that ‘the Eiffel Tower’ may only refer to a single *object*, both at the moment of baptism and at later times. It is also safe to conclude here, as we did in the previous subsection, that this result plausibly generalises: in general, it’s plausible that, given all three of microphysical nihilism, endurantist, and the causal theory of reference relevant to plural names, names of the form ‘N’ that purport to refer to single composite *F*s refer plurally to *xx* iff *xx* are arranged *N*-wise and *xx* are an object. That is (RN6).

4.4 – Conclusion

In this chapter I have argued that a broadly semantic externalist theory of meaning, and the causal theory of reference that it is naturally paired with, motivate the semantic claims of reductive nihilism, on the assumption that microphysical nihilism is true in the first place. Along the way I have argued for various modifications to the simplest versions of the causal theory, but in each such case my arguments have been made on grounds independent of microphysical nihilism. Thus I’ve argued that causal theorists should allow some role for descriptive content in the act of reference-fixing on the grounds that doing so will help to solve the *qua* problem; I’ve also argued for various modifications to the causal theory that must be made to accommodate complications arising from different theories of persistence; I’ve also argued that, on the assumption of endurantism, causal theorists should give a different account concerning the reference of *plural* names than the account they give in the case of *singularly referring* names, because otherwise they would get the intuitively wrong results in ordinary cases of plural reference. With these modified versions of the causal theory in hand, I have argued that (RN1)-(RN6) are justified.

Chapter 5

Defending the Semantics

5.1 - Introduction

Chapters 3 and 4 presented two independent arguments for the semantics of reductive nihilism. In this chapter I consider several objections to those semantics, and argue that those objections do not succeed.

My semantics have never been defended in their entirety in the literature before. Nevertheless, philosophers have formulated objections to the project of appealing to semantic considerations to attempt to render ostensibly counterintuitive metaphysical views compatible with common sense. I'll consider two such objections. The first (§5.2) explicitly targets *nihilist* attempts to render their view compatible with common sense in this way, arguing that such views are unmotivated; the second (§5.3) does not explicitly target nihilism, but may be thought to generalise to provide an argument from interpretive charity against reductive nihilism. After considering these, I'll anticipate two further objections that might plausibly be levelled against the semantics of reductive nihilism: that with those semantics in place, reductive nihilism is indistinguishable from a rival, non-nihilist answer to the SCQ, and therefore collapses into that rival view (§5.4); that reductive nihilism's semantics commit it to the objectionable principles of mereological essentialism and sufficiency (§5.5).⁹⁷

5.2 – A Lack of Motivation for Nihilist Compatibilism

Perhaps the leading critic of nihilist attempts to render their view compatible with common sense by adopting non-error-theoretic semantics is Daniel Korman (2009: §3; 2015b: ch. 5), in whose idiolect such nihilists are *compatibilists*. His objection to nihilist compatibilism is that nihilist compatibilist semantics are simply *unmotivated*, and that compatibilist nihilist semantics therefore amount to mere *wishful thinking* on the part of the nihilist, who *wishes* her view were compatible with common sense folk beliefs and utterances, but who provides no good reason for thinking so. In a review of a Horgan and Potrč's attempt to square their version of existence monism with folk beliefs and utterances via a compatibilist semantics, Korman summarises his general objection such a project neatly:

‘One argues at length for some surprising ontological thesis T (in this case, that tables do not exist) and then reasons as follows: the best conciliatory hypothesis on which T is not in conflict with ordinary discourse is H (e.g., that 'tables exist' is context-sensitive); therefore H. The missing premise is that T is not in conflict with ordinary discourse, and that is where the wishful thinking comes in. For there never seems to be any evidence whatsoever for H, which by hypothesis is the best conciliatory strategy on offer.’

(Korman 2008b)

Korman supports his objection that nihilist compatibilism is unmotivated in two ways. The first way is to consider two specific forms of nihilist compatibilism and to argue that the semantic claims that

⁹⁷ I won't consider here an objection endorsed by Merricks (2001: 163ff) and Korman (2015b: 50), according to which mereological nihilism, if true, intuitively *does* contradict our ordinary beliefs, and thus that any nihilist semantics that says otherwise is objectionably counterintuitive. I instead consider that objection in §6.7.3, as a part of a more general discussion of the relative intuitiveness of reductive nihilism.

these accounts make are unsupported by evidence. In particular, he suggests that there's none of the usual evidence for regarding sentences that apparently carry ontological commitment to composite *F*s as being instances of *loose talk* that, when interpreted as such, really only carry commitment to simples* arranged *F*-wise (2009: 247).⁹⁸ He similarly bemoans the lack of evidence for the attempt to render the claims of the folk and the nihilist compatible via the thought that the folk employ a 'lightweight' quantifier when they assert that composite objects exist whilst the nihilist instead employs a 'heavyweight' quantifier when she denies the existence of composite objects (247-8).

The second way that Korman supports his claim that nihilist compatibilist semantics are unmotivated is to consider three *general* considerations that might be taken to motivate nihilist compatibilism, and to argue that those considerations do not support nihilist compatibilism after all. First, he argues that metasemantic pressures such as considerations of interpretive charity and reference magnetism do not motivate nihilist compatibilism (2015b: 54-60).⁹⁹ Second, he objects to the idea that the surprising receptivity on the part of the folk to mereological nihilism lends support to the thought that ordinary folk assertions that are apparently in tension with nihilism should be given a compatibilist interpretation (60-63). Third, he dismisses the thought that evidence for nihilist compatibilism might be found in the fact that the appropriateness of ordinary utterances that apparently contain commitment to composite objects is insensitive to the truth or falsity of mereological nihilism (63-65).

We need not dwell on the details of what Korman has to say here. For our purposes, we can grant that the forms of nihilist compatibilism Korman cites are unsupported by the evidence, and that the general motivations for nihilist compatibilism he considers do not succeed. We can then simply reply to Korman that the compatibilist semantics of reductive nihilism are motivated in ways he doesn't consider. As I spent the previous two chapters demonstrating, on the assumption of microphysical nihilism these semantics follow from well-established and influential views of reference: in particular, the Ramsey-Lewis treatment of theoretical terms, and a broadly causal theory of reference. Only by undermining *these* arguments for reductive nihilism's semantics could it be shown that these semantics are unmotivated. But nothing Korman says engages with these arguments. For example, neither the argument for the semantics of reductive nihilism from the causal theory of reference nor that argument from the Ramsey-Lewis treatment of theoretical terms makes any appeal to interpretive charity, or to the receptivity of the folk to microphysical nihilism, etc.; similarly, the idea that we are engaging in loose talk when we commit ourselves to the existence of composite objects, or that we are using a 'lightweight' quantifier, is no part of the semantics of reductive nihilism.

Korman may be the most prominent critic of nihilist compatibilism, but reductive nihilism emerges unscathed from his attacks.¹⁰⁰

5.3 – *The Argument from Interpretive Charity Against Reductive Nihilist Semantics*

⁹⁸ See also Hawthorne and Michael (1996: 121-4).

⁹⁹ Thus Korman is here in disagreement with Hawthorne and Michael, who think that compatibilism is obviously favoured by consideration of charity (1996: 120), even if there are other reasons to reject compatibilism. See also below, where I explore whether consideration of interpretive charity might actually give us an *objection* to the compatibilist semantics of reductive nihilism.

¹⁰⁰ Korman also thinks that nihilist compatibilism is 'prima facie implausible' (2015b: 65), but doesn't say anything in defence of this claim. For what it's worth, I don't find the semantics of reductive nihilism at all 'prima facie implausible' – see §6.7.3 for further discussion.

David Liggins' (2020: §2) argument from charity against *hermeneutic fictionalism* might be thought to generalise to an argument from charity for the error-theoretic semantics of eliminative nihilism over that of reductive nihilism. I'll set out Liggins' argument and then consider how it might be generalised to provide an argument for accepting the error-theoretic semantics of eliminative nihilism over the non-error-theoretic ones of reductive nihilism. I'll then argue that this argument against the semantics of reductive nihilist fails.

Hermeneutic fictionalism is the view that participants in a given discourse 'do not believe or assert what is literally expressed by the sentences of the discourse' (2020: 81), but rather, on one version of the view at least, 'really believe something else – that according to some representation, things are thus and so – and this is what they are asserting' (82). If we limit our focus to cases in which what is literally expressed by the sentences of the discourse in question is *false*, hermeneutic fictionalism contrasts with what Liggins calls *error theory*, according to which participants in a discourse really do falsely believe or assert what is literally expressed by the sentences of the discourse. For example, suppose that there are no mathematical objects. Error theorists about mathematical discourse might regard the sentence '7 is a prime number' as false because it wrongly implies the existence of the number 7. Hermeneutic fictionalists might instead regard this sentence as true on the grounds that, contrary to appearances, it reports some fact that carries no commitment to the existence of numbers. Liggins takes the most influential argument for hermeneutic fictionalism to be from the metasemantic principle of interpretive charity (85), but argues (§2) that consideration of interpretive charity instead provides support for error theory.

Liggins' argument from interpretive charity for error theory over hermeneutic fictionalism takes as its starting point a conception of interpretive charity according to which there is metasemantic pressure to interpret speakers in such a way as to maximise the *reasonableness* of their utterances.¹⁰¹ On Liggins' gloss, this principle tells us 'not to ascribe inexplicable errors to people', where an inexplicable error is one that portrays us 'as less rational than we know ourselves to be' (86). The problem for hermeneutic fictionalism is that whilst error theorists ascribe widespread error to ordinary people with regard to e.g. mathematical discourse, hermeneutic fictionalists arguably ascribe widespread error to ordinary people with regard to their claims and beliefs about the *content* of their mathematical discourse.¹⁰² For example, when ordinary people assert '7 is a prime number', they arguably take themselves to be talking about the number 7 and saying of that number that it is prime. Not so, if hermeneutic fictionalism is true: if the folk were talking about the number 7 then they'd be speaking falsely because no such number exists, so the folk must instead be talking about something else. What's more, this error on the part of ordinary speakers with regard to the content of their mathematical discourse is much less reasonable than error on their part with regard to claims such as '7 is a prime number'. '7 is a prime number' is something that we were taught as children, and the arguments for its falsity due to the non-existence of mathematical objects are arcane and complicated, so falsely accepting this sentence as true is not at all unreasonable, irrational, or inexplicable. But to be wrong about the content of one's own beliefs and utterances is arguably much less reasonable or explicable: as Liggins says, 'common sense maintains that we have a powerful capacity to know our own minds; and philosophical argument has not overturned this' (90). If this is right, then interpretive

¹⁰¹ This conception of interpretive charity is typically preferred in the literature to the strict, Davidsonian requirement that we should instead interpret speakers in such a way as to maximise the *truth* of their utterances (see e.g. Davidson 1973: 19). Supporters of the weaker conception of charity *qua* reasonableness-maximisation include Lewis (1974: 336; 1983: 375) Hirsch (2011: 180), Korman (2015b: 36), and Grandy (1973: 440).

¹⁰² Liggins credits Stanley (2001: 46-7) with the point that hermeneutic fictionalism implies a lack of self-knowledge in this way.

charity favours error theory over hermeneutic fictionalism, since by adopting error theory the errors we ascribe to ordinary people are explicable and reasonable ones, whereas by adopting hermeneutic fictionalism we ascribe to ordinary people much less reasonable and explicable errors about the contents of their beliefs and utterances.

Applied to the case of eliminative and reductive nihilism, this sort of argument might run as follows. First, assume that microphysical nihilism is true. Now consider an ordinary person's utterance of 'the Eiffel Tower exists'. *Eliminative* nihilists say that this utterance asserts the existence of the Eiffel Tower. What's more, eliminative nihilists say that 'the Eiffel Tower' is the name for a putative single composite* entity, and that that single composite* entity does not exist. Thus, eliminative nihilists say, 'the Eiffel Tower exists' is false, and so the ordinary person is in error. But this is not an inexplicable or irrational error, because the reason for its falsity is the truth of microphysical nihilism, and there's no reason why ordinary people should be aware of the arcane and complicated arguments for microphysical nihilism. Moreover, the ordinary person is right to think that their utterance of 'the Eiffel Tower exists' asserts the existence of the Eiffel Tower.

Reductive nihilists, on the other hand, say that an utterance of 'the Eiffel Tower exists' *truly* asserts the existence of the simples* arranged Eiffel-Tower-wise. So whilst the ordinary person's utterance here is true, their belief about the content of their utterance is apparently *false*: they thought they were saying something about the existence of the Eiffel Tower, but really they were saying something about the existence (and arrangement) of some sub-atomic particles. What's more, this error on the part of the ordinary person is arguably unreasonable and inexplicable, given our capacity to know our own minds.

Thus, given the truth of microphysical nihilism, whilst both eliminative nihilism and reductive nihilism ascribe error to ordinary people, the errors that the ordinary folk make given *eliminative* nihilism's semantics are reasonable and explicable, whereas the errors that the ordinary folk make given *reductive* nihilism semantics are much less reasonable and explicable. So the principle of interpretive charity, understood as the metasemantic pressure to avoid ascribing inexplicable errors to the folk, apparently tells us to favour the error-theoretic semantics of eliminative nihilism.

Reductive nihilists can respond to this argument by rejecting the claim that they ascribe any error to ordinary people with regard to their beliefs about the content of their utterances in the first place. For whilst it's true that, given reductive nihilism, an ordinary utterance of 'the Eiffel Tower exists' says of some simples* arranged Eiffel-Tower-wise that they exist, it's *also* true that, given reductive nihilism, an ordinary utterance of 'the Eiffel Tower exists' says of the Eiffel Tower that *it* exists. It is wrong to characterise reductive nihilism as the view that 'the Eiffel Tower' *doesn't* refer the Eiffel Tower and *instead* refers to the simples* arranged Eiffel-Tower-wise: whatever else it might be true to say about what 'the Eiffel Tower' refers to, given reductive nihilism, by disquotation it refers to the Eiffel Tower. Thus when the folk utter 'the Eiffel Tower exists' they really are saying of the Eiffel Tower that it exists. The fact that reductive nihilists also say that 'the Eiffel Tower' refers to the simples* arranged Eiffel-Tower-wise doesn't change this.¹⁰³ The reductive nihilist simply doesn't ascribe any error to the folk with regard to their beliefs about the content of their utterances.¹⁰⁴

¹⁰³ Compare: an utterance of 'Mark Twain exists' says of Mark Twain that he exists, even though it's also true to say that this utterance says of Samuel Clemens that he exists (since 'Mark Twain' refers to Samuel Clemens).

¹⁰⁴ Hawthorne and Michael (1996: 122-3) similarly point out that van Inwagen's combination of organicism and compatibilist semantics isn't forced to deny that the predicate 'is a table' is true of tables.

One might object to this response by alleging that the folk believe that their utterances of ‘the Eiffel Tower’ exists concern the existence of a *single object* that is the Eiffel Tower, and that the reductive nihilist cannot accept that the folk are right in holding *this* belief, as reductive nihilists think that ‘the Eiffel Tower’ refers to *many* objects rather than one single one. But this is wrong: given the reductive nihilist’s semantics for the term ‘object’ (see (RN2) in §2.4.2), reductive nihilism does accept that ‘the Eiffel Tower’ refers to a single object. So, on reductive nihilism, the ordinary person is right to think that their utterance of ‘the Eiffel Tower exists’ says of a single object, i.e. the Eiffel Tower, that it exists.

One might rejoin by alleging that the folk don’t just believe that their utterances of ‘the Eiffel Tower exists’ say of a single *object* that it exists, but that they also believe that this utterance says of a single *entity* that it exists. Reductive nihilists cannot accept that the folk are right about *this* claim about the content of their beliefs, because reductive nihilists claim that ‘the Eiffel Tower’ refers to a many entities, i.e. many simples*. However, it’s not plausible that the folk believe that the content of an ordinary utterance of ‘the Eiffel Tower exists’ is that a single entity, i.e. the Eiffel Tower, exists. The notion of *being a single entity* is not a notion that features at all in ordinary discourse. Instead it’s a metaphysical term of art. To ascribe to the folk beliefs involving the term ‘entity’ is thus to ascribe to them an implausible amount of familiarity with metaphysics.¹⁰⁵

Thus, whilst considerations of interpretive charity might well show that there is no argument from charity *for* the semantics of reductive nihilism over those of eliminative nihilism (because the errors that the folk make given eliminative nihilism are perfectly explicable, and thus not in conflict with the principle of charity), it fails to show that there is an argument from charity *against* the semantics of reductive nihilism, because on reductive nihilism folk utterances and beliefs about not only the existence of composite objects, but also about the *content* of those utterances and beliefs, are true.

5.4 – Concerns about Collapse

Reductive nihilism affirms the existence of composite *Fs* such as tables, chairs, penguins, planets, etc. It agrees that those *Fs* each qualify as (single) *objects*; it agrees that they have (proper) parts and thus qualify as (single) *composite* objects. Indeed, reductive nihilism holds that composition (along with other mereological properties and relations) is a commonplace feature of reality. In accepting all of these things, reductive nihilism may sound very similar to rival answers to the SCQ. Given that reductive nihilists deny that gruesome composite objects exist, their view is perhaps most similar to a restrictivist answer to the SCQ according to which all of the composite objects of common sense exist but gruesome composites do not. Indeed, it might seem as though the reductive nihilist says *everything* that such a restrictivist says, and that the two views are indistinguishable. One might then object to the semantics of reductive nihilism on the grounds that their semantics make the view *collapse* into such a restrictivist answer to the SCQ.

But reductive nihilism does not collapse into any rival answer to the SCQ. This point is most simply appreciated by recalling that reductive nihilists by definition accept

MEREOLOGICAL NIHILISM: for all disjoint* *xx*, there is no *y* such that *xx* compose* *y*.

¹⁰⁵ See §3.5 and §6.4.3 for more on this.

whilst non-nihilist views by definition do not accept this. This is true no matter what semantic claims concerning terms such as ‘object’, ‘composition’, ‘parthood’, ‘the Eiffel Tower’, etc., that reductive nihilists accept.

Relatedly, although the reductive nihilist *agrees* with the restrictivist on perhaps all of the first-order sentences of English concerning what composite objects exist, what is a part of what, and so on, she *disagrees* with the restrictivist on what those sentences mean. In the reductive nihilist’s mouth, they mean things consistent with MERELOGICAL NIHILISM; in the restrictivist’s mouth, they mean things inconsistent with MERELOGICAL NIHILISM. For example, when the reductive nihilist claims that tables exist and are composite, she means

T: $\exists xx (xx \text{ are a table} \wedge \exists yy (yy \text{ are properly among } xx))$

T is consistent with MERELOGICAL NIHILISM. On the other hand, when the restrictivist claims that tables exist she means

T*: $\exists x (x \text{ is a table} \wedge x \text{ is composite}^*)$

T* is inconsistent with MERELOGICAL NIHILISM. As such, reductive nihilists reject T*.

Here’s another example. When the reductive nihilist claims that the brick is a part of the wall, she means

B: $\exists xx \exists yy (xx \text{ are a brick} \wedge yy \text{ are a wall} \wedge xx \text{ are among } yy)$

whereas the restrictivist means

B*: $\exists x \exists y (x \text{ is a brick} \wedge y \text{ is a wall} \wedge x \text{ is a part}^* \text{ of } y)$

Again, B is consistent with MERELOGICAL NIHILISM whilst B* is not. As such, whilst reductive nihilists accept B, they reject B*.

Thus the apparent equivalence between reductive nihilism and (a certain version of) restrictivism is illusory. Though the two views may accept the truth of the same claims about what exists, and what is a part of what, those claims mean different things in the mouths of the reductive nihilist and restrictivist respectively. Further, the reductive nihilist denies the truth of what the restrictivist means by these claims. So the two views are distinct.

5.5 – *An Objectionable Commitment to Mereological Essentialism and Mereological Sufficiency?*

One might object to the semantics of reductive nihilism on the grounds that they objectionably imply mereological essentialism and mereological sufficiency:¹⁰⁶

MERELOGICAL ESSENTIALISM: a whole’s parts are essential to it: a whole could not have different parts to the ones it actually has.

MERELOGICAL SUFFICIENCY: a whole’s parts suffice for it: in every world in which a whole’s actual parts exist, the whole exists too.

¹⁰⁶ This objection is analogous to one that has been levelled against CAI (see Cameron 2014: 93, Merricks 1999).

Mereological essentialism is objectionable.¹⁰⁷ It implies, for example, that the destruction of one of the atoms that is a part of the Eiffel Tower would amount to the Eiffel Tower's destruction, and that the Eiffel Tower cannot lose or gain parts.¹⁰⁸

Mereological sufficiency is likewise objectionable.¹⁰⁹ It forces us to say that we can only destroy the Eiffel Tower by destroying at least one of its parts. No matter how much we rearrange or scatter the Eiffel Tower's parts, as long as they exist, the Eiffel Tower does too.

Reductive nihilism's semantics might be thought to commit it to these objectionable principles. That's because, as I noted in §2.4.3.3, those semantics imply the identity of composite wholes with their parts. For example, the reductive nihilist's claim that 'the Eiffel Tower' refers plurally to the simples* arranged Eiffel-Tower-wise commits her to saying that the Eiffel Tower is identical to the simples* arranged Eiffel-Tower-wise. Now, call that plurality of simples* *ss*. Intuitively, *ss*'s members are essential to it, in that *ss* could not have had different members to the ones it actually has: for example, if we were to destroy one of *ss*'s members and replace it with a new simple*, then intuitively the resulting plurality would not be *ss*.¹¹⁰ Furthermore, intuitively, *ss*'s members suffice for it, in that in any world in which *ss*'s (actual) members exist, *ss* does too: intuitively, we can only destroy *ss* by destroying one or more of its members, and we cannot destroy *ss* by simply *rearranging* or *scattering* its members. What's more, it follows from these intuitive thoughts that *ss*'s *subpluralities* are both essential to it and suffice for it.¹¹¹ It then apparently follows from the identity of *ss* with the Eiffel Tower and Leibniz's Law that the Eiffel Tower's subpluralities are both essential to it and suffice for it; given that the reductive nihilist equates subpluralityhood and parthood, it follows from this that the Eiffel Tower's parts are both essential to it and suffice for it.

Thus the objection to the reductive nihilist is that, since her semantics require her to identify composite wholes with pluralities of parts, she is objectionably committed to mereological essentialism and mereological sufficiency.

The reductive nihilist, however, is *not* committed to these objectionable principles. She has (at least) two strategies available to her that allow her to agree that e.g. the Eiffel Tower could have had different parts to the ones it actually has, and that e.g. the Eiffel Tower would not exist any more were

¹⁰⁷ E.g. van Inwagen (1981, 1990a); van Cleve (1985); Cameron (2008: 8, 2014: 93); Merricks (1999).

¹⁰⁸ This latter consequence only follows if mereological essentialism is construed so as to entail mereological eternalism, the claim that a whole has its parts permanently. This construal seems standard, so I'll assume it.

¹⁰⁹ Merricks (1999: fn 4); Cameron (2008: 8, 2014: 93); Wallace (2014a: 114).

¹¹⁰ This corresponds to the conjunction of Uzquiano's principle (3) and (4) (2011: 220):

- (3) If one object is one of some objects, then it is necessarily one of them.
- (4) If one object is not one of some objects, then it is necessarily not one of them.

(Uzquiano's formulations here assume the necessity of existence, but clearly analogous principles to (3) and (4) could be formulated given contingentism). Uzquiano's (2011) argues that (3) and (4) are widely accepted and assumed, but don't actually follow from the usual axioms of plural quantification and propositional modal logic. He suggests supplementing the axioms of modal plural logic to render (3) and (4) theorems. If you *don't* find (3) and (4) intuitive, then so much the better for the reductive nihilist, as the objection to her that we are considering rests on them.

¹¹¹ This is because the following definitions hold: *xx* are a subplurality of *yy* =_{df} *xx* are among *yy*; *xx* are among *yy* =_{df} for all *z*, if *z* is one of *xx* then *z* is one of *yy*; *z* is an individual member of *xx* =_{df} *z* is one of *xx*. Together these definitions guarantee that facts about what subpluralities a plurality has are determined by facts about what individual members it has. Thus if in every world in which *xx* exists it has the same individual members, then in every world in which *xx* exists it has the same subpluralities; similarly, if in every world in which *xx*'s actual individual members exist, *xx* exists, then in every world in which *xx*'s actual subpluralities exist, *xx* exists.

its parts to be scattered. These strategies are adapted from strategies that proponents of *CAI* have deployed to try to demonstrate that their view is not committed to mereological essentialism and mereological sufficiency either. In the following two sections I'll outline two different strategies for avoiding mereological essentialism and mereological sufficiency that have been suggested on behalf of *CAI*, and show how the reductive nihilist can utilise closely related strategies to the same end.¹¹²

5.5.1 – Modal Subplurality Theory

We begin with perhaps the more extreme of the two strategies. Meg Wallace has argued for conjoining *CAI* with what I'll call *modal parts theory* (2014a). Modal parts theory is the view that, just as ordinary objects are apparently spatially extended in virtue of having spatial parts, so they are *modally* extended in virtue of having modal parts at other worlds.¹¹³ It is the modal analogue of *perdurantism*, the view that ordinary objects are *temporally* extended in virtue of having temporal parts. Indeed, Wallace's definition of modal parts is adapted from Sider's definition (2001: 59) of temporal parts:

MODAL PART: x is a world-bound modal part of y at a world $w =_{df}$ (i) x exists at, but only at, w ; (ii) x is a part of y at w ; and (iii) x overlaps at w everything that is part of y at w .

(Wallace 2014a: 117)

(Note that Wallace's use of singular variables for the relata of the parthood relation in MODAL PART suggests that she is taking 'parthood' to mean parthood* [see §2.3], as is standard)

Ordinary objects are thus conceived as modally extended mereological sums of their modal parts. The view is easiest to grasp if Lewisian genuine modal realism (GMR) is assumed, and possibilities thought to be as real and concrete as actualia. Wallace argues that modal parts theory is nonetheless consistent with any form of modal realism, including views that take possible worlds to be abstract entities of some sort, as abstracta can plausibly enter into mereological relations with concreta (2014a: 123-6; 2019: §4.1). I'll grant this point for now, and return to it briefly at the end of the section.

Modal parts theory can be used to provide an analysis of *de re* modal predication. To say that the Eiffel Tower is actually 324m tall is to say that the Eiffel Tower's @-modal-part is 324m tall; to say that the Eiffel Tower could have been 400m tall is to say that the Eiffel Tower has a 400m tall w_n -modal-part, for some possible world w_n .

When conjoined with modal parts theory, *CAI* is the view that ordinary objects are modally extended composites that are identical to the *modally extended* pluralities of parts that compose them. *CAI*, so conceived, can avoid mereological essentialism and mereological sufficiency.

First, assume *CAI* and modal parts theory, and suppose that, at @, a certain molecule, m , is a part of the Eiffel Tower. That is, suppose that m 's @-modal-part is a part of the Eiffel Tower's @-modal-part. Suppose then, as is perfectly consistent with this, that although the Eiffel Tower has a w_l -modal-

¹¹² I take my discussion here to generalise in relatively straightforward ways to the *temporal* analogue of the objection from mereological essentialism and sufficiency (namely that reductive nihilists cannot accept that, as time passes, some objects will gain and lose parts, and some parts will begin to or cease composing the object that they currently compose). But I won't spell out how this works here.

¹¹³ Wallace calls her view 'Lump Theory'. She defends it independently of *CAI* in her (2014b, 2019). As she notes, theories related to lump theory / modal parts theory are discussed, though not defended, in Weatherson (n.d.), Kaplan (1979), Varzi (2001) and Benovsky (2006). See also Quine (1976) and Lewis (1986: §4.3). The latter accepts that there are transworld fusions, but denies that they are ordinary objects.

part, m 's w_I -modal-part is not a part of the Eiffel Tower's w_I -modal-part. In that case, at w_I , m is not a part of the Eiffel Tower, despite the Eiffel Tower existing at w_I . This is a case in which the Eiffel Tower differs with respect to its parts from @ to w_I , and in which it's therefore true that the Eiffel Tower could have had different parts. Thus CAI escapes mereological essentialism.¹¹⁴

Second, suppose that, at @, the Eiffel Tower is composed of molecules m_1, \dots, m_{100} . That is, suppose that the @-modal-parts of m_1, \dots, m_{100} are each parts of the @-modal-part of the Eiffel Tower (and that no further disjoint object is part of the Eiffel Tower's @-modal-part). Suppose then, as is perfectly consistent with this, that each of m_1, \dots, m_{100} have w_I -modal-parts but the Eiffel Tower has no w_I -modal-part. This is a case in which the Eiffel Tower's (actual) parts do not suffice for it: its (actual) parts exist at w_I , but the Eiffel Tower doesn't. It's a case in which it's true that the Eiffel Tower's parts could exist without the Eiffel Tower existing. Thus CAI escapes mereological sufficiency.¹¹⁵

Reductive nihilists can escape mereological essentialism and sufficiency in an analogous way. They can endorse what I'll call *modal subplurality theory*. Modal subplurality theory holds that terms that purport to refer to ordinary objects, such as 'the Eiffel Tower', refer to *modally extended pluralities of simples**. Additionally, locutions such as 'the simples* arranged Eiffel-Tower-wise' also refer to these modally extended pluralities (thus it's still true that 'the Eiffel Tower' refers to the simples* arranged Eiffel-Tower-wise). These modally extended pluralities of simples* have *modal subpluralities*, to be understood as analogous to modal parts:

MODAL SUBPLURALITY: xx are a modal subplurality of yy at a world $w =_{df} xx$ are the largest plurality that exist only at w and that are among yy .

Again, the view is easiest to grasp if GMR is assumed. But it's plausible that if, as Wallace argues (see above), modal parts theory is consistent with abstractionist conceptions of possible world, then so is modal subplurality theory.

Modal subplurality theory provides the reductive nihilist with an analysis of *de re* modal predication. Suppose that 'the Eiffel Tower' refers to the modally extended plurality of simples* arranged Eiffel-Tower-wise. Call these simples* ee . To say that the Eiffel Tower is actually 324m tall is to say that ee 's @-modal-subplurality is jointly 324m tall; to say that the Eiffel Tower could have been 400m tall is to say that ee has a w_n -modal-subplurality that are jointly 400m tall, for some possible world w_n .

By endorsing modal subplurality theory, the reductive nihilist can also avoid mereological essentialism and mereological sufficiency.

First, suppose that, at @, some simples* arranged molecule-wise, mm , are among ee (note that mm , like ee , are a modally extended plurality of simples*). That is, suppose that mm 's @-modal-subplurality are among ee 's @-modal-subplurality. Suppose then, as is perfectly consistent with this, that although ee have a w_I -modal-subplurality, mm 's w_I -modal-subplurality are not among ee 's w_I -

¹¹⁴ Wallace agrees with the gist of this, but denies that it constitutes a counterexample to mereological essentialism (2014a: 118-9, 121-3). And there's clearly a sense in which the Eiffel Tower has its parts essentially even on modal parts theory: the Eiffel Tower *qua* transworld fusion trivally retains its parts (which include its modal parts) no matter what world we're in. But by the modal parts theorist's own semantics for *de re* modal predication, there's a sense in which 'the Eiffel Tower could have had different parts' is true on modal parts theory. And this result seems to be enough to escape mereological essentialism on its intended reading.

¹¹⁵ Again (see previous footnote), I'll regard being able to accept the truth of 'the Eiffel Tower's parts could exist without the Eiffel Tower' as enough to escape mereological sufficiency.

modal-subplurality. This is equivalent to: at w_1 , mm are not among ee , despite the fact that ee exist at w_1 . Now recall the reductive nihilist's semantics: 'the Eiffel Tower' refers to ee , and 'parthood' expresses *amonghood*. Additionally, the reductive nihilist can claim that 'molecule₁' refers to mm . So the reductive nihilist can describe this scenario as: at @, molecule₁ is a part of the Eiffel Tower, but at w_1 it is not, despite the Eiffel Tower existing at w_1 . In this way the reductive nihilist escapes mereological essentialism.

Second, suppose that, at @, ee 's subpluralities are 100 pluralities of simples* arranged molecule-wise, mm_1, \dots, mm_{100} (and that, at @, no plurality that is disjoint from all of mm_1, \dots, mm_{100} is a subplurality of ee ¹¹⁶). That is, suppose that ee 's @-modal-subplurality has as subpluralities the @-modal-subpluralities of each of mm_1, \dots, mm_{100} (and that no plurality that is disjoint from all of mm_1, \dots, mm_{100} is a subplurality of ee 's @-modal-subplurality). Suppose then, as is perfectly consistent with this, that each of mm_1, \dots, mm_{100} has a w_1 -modal-subplurality, but that ee has no w_1 -modal-subplurality. That is, at w_1 , each of mm_1, \dots, mm_{100} exist but ee does not. Recall again the reductive nihilist's semantics: 'the Eiffel Tower' refers to ee ; given that 'parthood' expresses *amonghood*, 'part' refers to *subplurality*; we can additionally assume that 'molecule₁' refers to mm_1 , and so on. Thus the reductive nihilist can describe this scenario as: at @, molecule₁, ..., molecule₁₀₀ are the parts of the Eiffel Tower, but at w_1 , although molecule₁, ..., molecule₁₀₀ all exist, the Eiffel Tower does not. In this way the reductive nihilist escapes mereological sufficiency.

Now, I concede that modal parts theory, and by extension modal subplurality theory, is controversial. Perhaps its most worrying consequence is that it commits us either to GMR or (if non-actual possible worlds are instead thought to be abstract objects) to the claim that ordinary objects have as parts many abstract objects. Both of these options may seem like a high price to pay.¹¹⁷ Thus I also offer a second strategy for avoiding commitment to mereological essentialism and mereological sufficiency on behalf of the reductive nihilist.

5.5.2 – (Nihilist) Counterpart Theory

The second reductive nihilist strategy for avoiding objectionable commitment to mereological essentialism and mereological sufficiency also co-opts a strategy that has been suggested on behalf of CAI for doing the same. Again, I'll present the strategy as it pertains to CAI, and then show how reductive nihilists can co-opt it.

Instead of endorsing modal parts theory, then, it has been suggested that proponents of CAI can endorse a counterpart theoretic analysis of *de re* modality to attempt to avoid commitment to mereological essentialism and sufficiency (Merricks 1999: 194; Cameron 2014: 93). According to this view, ordinary objects are not modally extended entities, but rather worldbound ones.¹¹⁸ Nonetheless, actual objects bear *counterpart* relations to objects that exist at other worlds (see §4.3.1 for mention of the related notion of *temporal* counterpart relations). Counterpart theorists then offer the following analysis of *de re* modal predication: an object, o , is possibly F iff at some possible world o has a

¹¹⁶ Two pluralities are disjoint iff they do not plurally overlap. That is, given the definition of plural overlap (see §2.2.2), xx and yy are disjoint =_{df} there is no z such that z is one of both xx and yy .

¹¹⁷ Objections to GMR (including the notorious incredulous stare) are well known; van Inwagen (1987: 35), for example, objects to fusions of concreta and abstracta.

¹¹⁸ Counterpart theorists can and often do (e.g. Lewis 1986: §4.3) accept the existence of transworld fusions, but deny that such objects are ordinary objects. If perdurantism is accepted, ordinary objects turn out to be spatially and temporally extended mereological fusions; if it isn't, then they're merely spatially extended mereological fusions.

counterpart that is F .¹¹⁹ The counterpart relation is a *similarity* relation; what it is for some object o^* that exists at w_m to be a counterpart of o (that exists at w_n) is simply for o^* to be sufficiently similar to o in some salient respect. Following Lewis (1971), there are a multiplicity of counterpart relations, each corresponding to a different respect in which two objects can be similar. Which counterpart relation (and thus which dimension of similarity) is salient in any given instance of a *de re* modal predication is determined by context: specifically, it is determined by the subject term to which the modal predicate (e.g. ‘is possibly F ’) is attached (Lewis 1971: 209; Noonan 1991: 188).¹²⁰ For example (adapted from Lewis [1971: 207-10]), ‘I am possibly F ’, selects the *personal* counterpart relation, such that this sentence is true iff I have at some world a counterpart that is sufficiently similar to me with respect to things like my personality, memories, etc. and that is F ; ‘my body is possibly F ’ instead selects the *bodily* counterpart relation, such that this sentence is true iff my body has at some world a counterpart that is sufficiently similar to it with respect to its physical characteristics (height, weight, appearance, etc.), and is F .

Given counterpart theory, CAI can avoid both mereological essentialism and mereological sufficiency. Consider the Eiffel Tower, and suppose that it is composed of (and thus identical to) 100 molecules, m_1, \dots, m_{100} . Nothing forces the proponent of CAI to think that all of the Eiffel Tower’s counterparts are composed of counterparts of m_1, \dots, m_{100} . Plausibly, in *de re* modal predications of the Eiffel Tower, a counterpart relation that corresponds to similarity with regard to *appearance*, and/or *cultural significance*, and/or *functional role*, (but not *exact mereological composition*), etc. is contextually salient. It’s plausible that there is some world, w_n , in which an object has the same appearance, cultural significance, functional role, etc. as the Eiffel Tower, and is thus the Eiffel Tower’s counterpart at w_n , but which is composed of the counterparts of e.g. m_{10}, \dots, m_{110} , rather than the counterparts of m_1, \dots, m_{100} .¹²¹ On the counterpart theoretic analysis of *de re* modal predication, it’s therefore true that the Eiffel Tower could have been composed of different parts to the ones it is actually composed of, and mereological essentialism is false.

Similarly, proponents of CAI can say that there is a world in which each of m_1, \dots, m_{100} have counterparts, but in which the Eiffel Tower has no counterparts. This might be a world in which the counterparts of m_1, \dots, m_{100} are scattered, and nothing is similar to Eiffel Tower with regard to appearance, cultural significance, functional role, etc. Thus proponents of CAI can say that the parts of the Eiffel Tower could exist without the Eiffel Tower existing, and deny mereological sufficiency.

Reductive nihilists can avoid mereological essentialism and sufficiency in a broadly analogous way.

Reductive nihilists extend counterpart theory by countenancing irreducibly *plural* counterpart relations, each corresponding to a way in which two pluralities of entities can be similar. The irreducibility of plural counterpart relations to ordinary singular ones is reflected in the reductive nihilist’s rejection of the following principle:

REDUCTIVE PLURAL COUNTERPART: xx at w_n are the plural counterparts of yy at w_m iff (for all z , z is one of xx iff z ’s singular counterpart at w_m is one of yy).

¹¹⁹ See Lewis (1968; 1971; 1973: 39-43; 1986).

¹²⁰ Modal predicates thus turn out to be *Abelardian*: they are predicates whose denotations are affected by the subject terms to which they are attached (see Noonan 1991: 188).

¹²¹ I don’t attempt here to supply an account of the counterpart relation that is contextually salient in *de re* modal predications of molecules: rather, I hope that it’s plausible that *any* acceptable way of filling in the details of this account will render the claim to which this footnote is attached true.

Rather, REDUCTIVE PLURAL COUNTERPART characterises just *one* plural counterpart relation, where (as in the singular case) we should accept a multiplicity of them.¹²² For whilst two pluralities can be similar in that they have all their members in common, they can also be similar in other ways despite differing with regard to their membership. For example, two pluralities with different memberships can be *arranged* in the same way as one another, or they can have the same functional role, or the same cultural significance. So reductive nihilists should insist that there are other plural counterpart relations to the one characterised by REDUCTIVE PLURAL COUNTERPART, each one corresponding to a different way in which two pluralities can be similar. We can then analyse *de re* modal predications of pluralities of the form ‘*xx* could have been *F*’ as being true iff *xx* bear the contextually salient plural counterpart relation to some *yy* at some other world w_n , and *yy* are *F*.

With this in mind, the reductive nihilist can avoid mereological essentialism. Suppose that, in @, there is a plurality of simples* arranged Eiffel-Tower-wise, *ee*; suppose further that *ee* has its subpluralities 100 pluralities of simples* arranged molecule-wise, mm_1, \dots, mm_{100} . Given reductive nihilism, ‘the Eiffel Tower’ refers to *ee*, ‘part’ expresses *subplurality*, and ‘molecule₁’ refers to mm_1 (etc.). Thus we can describe our scenario as: in @, the Eiffel Tower has as its parts 100 molecules, molecule₁, ..., molecule₁₀₀. Reductive nihilists claim that when we refer to *ee* using ‘the Eiffel Tower’, the contextually salient plural counterpart relation corresponds to similarity with respect to (*joint*) appearance, and/or (*joint*) cultural significance, and/or (*joint*) functional role, etc. It’s plausible that there’s some world, w_n , in which some plurality, *ff*, are exactly alike *ee* with regard to their (*joint*) appearance, (*joint*) cultural significance, etc., and are thus *ee*’s plural counterparts, but which have as their subpluralities the plural counterparts of each of mm_{10}, \dots, mm_{110} , rather than the plural counterparts of each of mm_1, \dots, mm_{100} . Given the reductive nihilist semantics, this is a case in which the Eiffel Tower has a counterpart in w_n that differs from the Eiffel Tower with respect to its molecular parts. Given the counterpart theoretic semantics, it’s therefore true that the Eiffel Tower could have had different parts to the ones it actually has, and mereological essentialism is false.

Similarly, reductive nihilists can say that there is a world in which each of mm_1, \dots, mm_{100} have counterparts, but in which there’s no plurality of simples* that are jointly sufficiently similar to *ee* with respect to appearance, cultural significance, etc. to be *ee*’s plural counterparts. For example, this might be a world in which the plurals counterparts of each of mm_1, \dots, mm_{100} are scattered. Given the reductive nihilist’s semantics, this is a case in which the molecular parts of the Eiffel Tower can exist without the Eiffel Tower existing. Thus mereological sufficiency is false.

5.6 – Conclusion

Thus the semantics of reductive nihilism escape four important objections: that they are unmotivated; that they violate a principle of interpretive charity; that they entail the collapse of reductive nihilism into a rival answer to the SCQ; that they entail the objectionable principles of mereological essentialism and mereological sufficiency. These semantics are thus not only well-supported by the consideration adduced in the previous two chapters, but are also robust and defensible against objections, both extant and anticipated.

¹²² Or rather, REDUCTIVE PLURAL COUNTERPART characterises just one *type* of plural counterpart relation. For there is a distinct *reductive* plural counterpart relation for each of the many singular counterpart relations that can be plugged in on the right-hand-side of this principle’s ‘iff’.

Chapter 6

An Intuitive Answer to the Special Composition Question

6.1 – Introduction

This chapter marks a shift in the thesis so far. Up until now, I have focused solely on arguing for the semantics of reductive nihilism (on the assumption that the metaphysical claims of microphysical nihilism are true). I turn now to arguing for reductive nihilism as a whole.

In this chapter, I'm going to argue that reductive nihilism is an intuitive answer to the SCQ. In fact, I'm going to argue that it's the *most* intuitive answer to the SCQ. Making this argument will require me to establish, broadly speaking, two things: first, what the relevant intuitions are; second, that reductive nihilism satisfies these intuitions better than any extant rival view does. So that is what this chapter will do.

I'm going to understand 'answer to the SCQ' in such a way as to allow for comparison between theories that, strictly speaking, answer the SCQ in the same way, but that differ in other important respects. Thus, for example, even though eliminative nihilism, reductive nihilism, emergence nihilism, and existence monism all give the nihilist answer to the SCQ (i.e. they all endorse MERELOGICAL NIHILISM – see §2.3), each of these views differs from the others with respect to its views on whether e.g. tables exist, or whether the universe exists, or even on whether composition (properly understood) occurs, etc., and these differences have consequences for the extent to which of these views can be said to accord with intuition.

The views that I'll be comparing in this chapter, then, are

- Reductive nihilism
- Eliminative nihilism
- Emergence nihilism
- Existence monism
- Various forms of restrictivism
- Brutal Composition
- Universalism
- CAI

Three entries on this list require further comment. First, since CAI is a view solely about the nature of composition (in that it is the view that composition is [many-one] identity), it is apparently consistent with any view about when composition occurs (c.f. Cameron 2012), about the spatial extension of simples, and about whether compositional facts are brute or not. So CAI is apparently combinable with any of the other views listed. Throughout this chapter, though, the views I discuss should be assumed *not* to accept CAI unless I make explicit otherwise.

Second, I'll understand Brutal Composition as accepting all of our intuitive judgements about when composition occurs and when it doesn't. Given Markosian's description of the view, Brutal Composition is consistent with any sets of claims about when composition does and doesn't occur

(see in particular Markosian 1998: 240). But characterising it as agreeing with our intuitions on this matter will give the view the best chance of prevailing in the intuition stakes.

Third, perhaps the most prominent defence of existence monism in the literature combines the view with a non-standard semantics that supposedly mitigates its apparently counterintuitive consequences (Horgan and Potrč 2008). Horgan and Potrč's treatment of this semantics is book-length, and I can't hope to engage with it fully here. I note, though, that Korman in a review (2008b) has argued that this semantics is unmotivated and amounts to little more than 'wishful thinking' on the part of Horgan and Potrč (note that, as I argued in §5.2, reductive nihilism's semantics are not similarly unmotivated [see also chapters 3 and 4]), and that the semantics is committed to a revisionary and 'potent form of relativism' with regard to the truth value of the contents of sentences. It also seems to me that the version or adaptation of Liggins' argument from charity against views that attempt to render their revisionary metaphysical claims compatible with common sense via non-standard semantics will bite against Horgan and Potrč's version of existence monism (note that I discussed this sort of argument from charity as it pertained to reductive nihilism in §5.3, where I argued that reductive nihilism was immune to it). Briefly, the Liggins-esque complaint here would be that Horgan and Potrč's semantics avoids ascribing errors to the folk regarding the existence of ordinary objects only at the cost of ascribing errors to them regarding the content of their own beliefs about ordinary objects, where the latter errors are far less reasonable ones, meaning that consideration of interpretive charity should pressure us to reject Horgan and Potrč's semantics in favour of a straightforward error theoretic semantics that ascribe errors to the folk with regard to the existence of ordinary objects but that vindicate folk beliefs about the contents of their own beliefs. So – though I cannot argue for this comprehensively here – there appear to be some good reasons for finding Horgan and Potrč's semantics objectionable. For these reasons, I'm *not* going to consider existence monism in conjunction with Horgan and Potrč's semantics: rather, I'll assume that existence monism has the counterintuitive consequences that it appears to.

Now, it's worth making a perhaps obvious point: a quick look at the literature on the SCQ – as well as a moment's reflection – reveals that the intuitions relevant to evaluating the various answers to the SCQ take the form not only of simple ontological intuitions (that is, intuitions concerning the existence or non-existence or particular objects), but also of intuitions concerning the nature of material objects, the nature of composition and parthood, and certain semantic intuitions about which I'll say more later. I will take into account all of these intuitions when arguing that reductive nihilism is the most intuition answer to the SCQ.

My discussion will be complicated by two further considerations. The first is that many of the intuitions that we seem to have concerning composition are typically stated in vocabulary that the reductive nihilist understands differently to other theorists. This opens to the door to the worry that, by interpreting statements of intuitions in line with her non-standard semantics, the reductive nihilist has misinterpreted the intuition. I'll deal with this worry as follows. As a default, I'll allow each theory to interpret the intuition in question in accordance with its own semantics. Thus, for example, as a default I'll take it that if one intuition we have is that the table is an object, and that given the reductive nihilist's metaphysical views together with her semantic claims about what 'table' and 'is an object' mean she accepts that the table is an object, then she has satisfied that intuition. But where it seems possible that in interpreting a given intuition in accordance with her semantics the reductive nihilist has thereby misinterpreted the intuition, I'll consider the putative alternative interpretation of the intuition and explicitly address both whether it is a fair interpretation of the intuition and, if it is, whether the reductive nihilist can satisfy it. So, for example, in the course of arguing that the reductive nihilist can satisfy our intuitions about tables being objects, I'll consider the objection that

the proper statement of that intuition is that tables are object-*qua*-entities, i.e. that tables are entities (see §6.5.3).

The second complicating factor is that even once we've determined that, for example, it's intuitive that there are tables, there's still a further question about what exactly it takes to satisfy this intuition. Has a theory done enough if it says that 'there are tables' is strictly speaking false but in some yet-to-be-determined sense *correct* or *assertable*? Or is it rather that a theory must say that 'there are tables' is strictly and literally true? Or is the focus on the correctness or truth of sentences or propositions beside the point, and it is only if our theory says that, simply, *there are tables* that it satisfies our existential intuitions about tables?

With this in mind, this chapter will be structured as follows. After a short discussion of why it's reasonable to think that we should *care* about which answer to the SCQ is the most intuitive (§6.2), I'll present what I take to be the most obvious intuitions concerning the SCQ, namely simple ontological intuitions about composite objects (§6.3). That will give us enough subject matter to meaningfully address the question of whether these intuitions are properly interpreted as concerning the truth of sentences, or the nature of reality, etc. I'll do this in §6.4, and argue that reductive nihilism satisfies existential intuitions, properly interpreted, better than (or at least as well as) any other rival theory. I'll then (§6.5) discuss intuitions concerning the nature of material objects and of composition (§6.6), arguing that reductive nihilism again outperforms almost all its rivals with respect to these intuitions. At this stage just one rival to reductive nihilism's claim to be the most intuitive answer to the SCQ will emerge: a form of CAI. I'll argue (§6.7) that this form of CAI is nonetheless significantly less intuitive than reductive nihilism because the former's claims about identity and number are wildly counterintuitive, whilst the latter's semantic claims are not counterintuitive at all.

6.2 – *Why care about intuition?*

Demonstrating that reductive nihilism is more intuitive than rival views would be a fairly pointless exercise here if doing so wouldn't allow us to conclude anything about the likelihood that reductive nihilism is true. But there does seem to be some good reason for taking intuitiveness to be a theoretical virtue – at least when it comes to theories of composition – such that if theory A is more intuitive than theory B then, *ceteris paribus*, we should have more credence in A than in B.

Indeed, in the composition debate, appeal to intuition is arguably the most common dialectical move made in attempts to settle the SCQ, or at least to rule out certain candidate answers.¹²³ And the alleged counterintuitive consequences of nihilism in particular are often regarded as sufficient to reject the view outright.¹²⁴ There is thus a presumption among metaphysicians working on the SCQ that intuitiveness is a relevant consideration in assessing the epistemic likelihood of competing theories.

Several different justifications have been offered in support of this broad presumption, though these various justifications support differing conceptions of just how weighty a theoretical virtue intuitiveness should be regarded as.

On one end of the spectrum, philosophers such as Zimmerman (2008) and Kriegel (2008) subscribe to the view that, in Zimmerman's words, 'something's being commonsensical must be allowed to count very strongly in its favor, other things being equal' (2008: 222). For Zimmerman this seems to be

¹²³ See for example Carmichael (2015), Contessa (2014), Korman (2015b), Markosian (1998; 2008).

¹²⁴ See for example Korman (2015b: 27), Markosian (1998: 220-1; 2008: §4), Noonan (2013: endnote 1), Rea (1998: 348). These sources would appear to have *eliminative* nihilism in particular in mind.

motivated by the thought that, in many cases, intuition is the *only thing we have to go on* when assessing philosophical views, and so to deny it any epistemic role would amount to ‘courting extreme skepticism’ (222) with regard to all sorts of questions, such as whether the external world exists, whether induction can provide us with knowledge, whether there are minds other than my own, and so on. Assuming that scepticism in such cases is unwarranted, we have reason to take intuition to be a valuable epistemic tool.¹²⁵ Kriegel instead suggests that our intuitions, or at least intuitions of a certain type,¹²⁶ constitute the *data* of metaphysical theories, in the same way that experience constitutes the data of scientific theories. Thus ‘a metaphysical theory faces the tribunal of intuition in the same way a scientific theory faces the tribunal of experience’ (2008: 362), and just as there’s a strong epistemic presumption in favour of a scientific theory that accounts for the empirical data better than all rivals, so the superiority of a metaphysical view with respect to intuitiveness ‘should create an overwhelming, if ultimately controvertible, presumption in [that view’s] favor’ (362).

Other philosophers, though more circumspect, suggest further (though related) reasons for taking intuitiveness to count defeasibly in favour of metaphysical theories. Contessa, channelling Moore, suggests that ‘we might think that our beliefs about cats, tables, and apples [etc.] are far more epistemically justified than some of the premises on which the arguments for nihilism [or any metaphysical view that would contradict such beliefs] rely’ (2014: 200), and continues, ‘or we might think that systematically ascribing false beliefs to ordinary people runs afoul of principles of interpretive charity, or is ultimately self-defeating’ (200, footnote citing Hirsch (2002) for the point about interpretive charity omitted). This latter point he attributes to Korman (2009) (who in turn cites van Inwagen [1990a: 103]), who points out that if our intuitions about the world are systematically false, then it’s reasonable to conclude that we are almost if not completely epistemically cut off from reality, in which case we have no reason to believe that our intuition-contradicting theories are correct. Korman additionally argues that regarding intuitively-held beliefs as being by and large true is the best explanation of why we find it reasonable to have these beliefs in the first place. For it’s hard to see, claims Korman, how can this explanatory burden be met, if not by appeal to the truth of such beliefs (2009: 243-4).

Some philosophers, however, play down the role of intuition in metaphysical theorising. Sider, for example, concedes that ‘match with ordinary usage and belief sometimes plays a role in this assessment [of competing metaphysical theories]’, but insists that it’s ‘typically not a dominant one’ (2009: 385). I suspect that many metaphysicians would agree (an obvious example would be Lewis, who was notoriously happy to trade off intuition for other theoretical benefits in his argument for Genuine Modal Realism [1986]). Rose and Schaffer go so far as to suggest a debunking argument against folk intuition, drawing from their claim that folk mereology is teleological the claim that folk intuition concerning composition, at least, is the product of non-truth-conducive biases and as such is not to be trusted (2017: §4). They furthermore take a sceptical stance on the suggestion that either folk intuition freed of its teleological biases or else the intuitions of philosophers might be truth-conducive. But their discussion on this point is brief, and rests on the claim that ‘it certain seems from the current literature as if there is little agreement among the metaphysicians, not just with respect to theory but with respect to intuitions reported’ (261). Even setting aside the question of whether the literature displays as much disagreement as Rose and Schaffer think it does (I’m dubious, and the rest of this

¹²⁵ See also Tallant (2011: 14) who argues that all scientific inquiry is built on intuitively-held foundations (such as that we’re not deceived by an evil demon, that induction is a legitimate method of knowledge acquisition, etc.), and the success of scientific inquiry suggests that those foundations (and thus their justification by intuition) are firm.

¹²⁶ That is, pre-philosophical intuitions of the form ‘*a* is *F*’ (2008: 362-3). See also Rosen and Dorr (2002: 152).

chapter contains many examples of completely or largely uncontroversial intuitions reported by metaphysicians in the literature¹²⁷) and granting for argument's sake that metaphysicians' intuitions sometimes or often conflict, it would not follow that intuitions give us no guide to truth whatsoever. That *empirical* evidence sometimes points in conflicting directions (the knife stained with the victim's blood found in Miss Scarlet's dustbin, the threatening letters bearing Professor Plum's fingerprints) does not mean that it ceases to be evidence, or that discovering evidence for p does not give us reason to raise our credence in p , all else being equal (and in particular, holding fixed evidence for $\neg p$); so why think that the existence of conflicting intuitions with regard to p implies that discovering that p is intuitive to some people does not give us reason to raise our credence in p , all else being equal (and in particular, holding fixed whether and the extent to which other people either find p unintuitive or find $\neg p$ intuitive)?

It seems to me, therefore, that there are good reasons for according intuition at least some role in epistemic assessment, at least when it comes to considering answers to the SCQ. Even if we, like Sider and perhaps Lewis, are tempted to downplay the role of intuition in theory choice, it still seems plausible to agree with Kriegel that 'it is unquestionable that, other things being equal, an account of composition that respects intuition is preferable to one that does not (2008: 361), with Rosen and Dorr that 'if unrestricted composition (or nihilism or some other restrictive principle) is at odds with settled opinion – and if there is nothing compelling to be said in favour of revision – then it is reasonable for us to reject it despite its consistency and its adequacy to the phenomena' (2002: 156), and with Merricks that a sensible methodology is to 'start off with all the common sense [we] can get, unloading it only when prompted to do so by the overall burden of argument' (2001: 24). By doing so we'll be staying consistent with the tacitly (and often explicitly) agreed methodology of most other philosophers working in the composite debate, and we'll be engaging with those would reject nihilism for its alleged counterintuitive consequences on their own terms.¹²⁸

Thus, in what follows, I'll assume that if reductive nihilism can be shown to be more intuitive than its rivals, then we have defeasible reason for endorsing it.

6.3 – *Ontological Intuitions*

The most obvious sorts of intuitions that we have concerning composition are intuitions concerning what material objects exist. Here's a list of just some of the things that philosophers intuitively hold to exist, in no particular order:

- People (Markosian 1998: 220, 2008: 351)
- Human beings (Carmichael 2015: 475; Rosen and Dorr 2002: 152; Markosian 1998: 211)

¹²⁷ The only evidence that Rose and Schaffer cite in favour of their claim that there is widespread disagreement among metaphysicians when it comes to intuitions about composition concerns disagreement in the literature as to whether it's counterintuitive or not to say that there are gruesome composites, such as the object putatively composed out of my nose and the Eiffel Tower (2017: 241). The disagreement here typically centres around whether the relevant intuition is satisfied as long as it's false in ordinary contexts that there are such gruesome composites, or whether intuition demands that there simply be no such objects. Note that even here there is widespread (though apparently not completely univocal – see e.g. Thomasson 2007: 184) agreement, with most thinkers granting that intuition disapproves of gruesome composites in *some* sense. See §6.4, and especially §6.4.4, for further discussion.

¹²⁸ Of course, I make no claim of having provided a conclusive defence of the idea that intuition should play a role in theory choice. But I take myself to have made the idea seem widespread and plausible (intuitive?) enough to justify this chapter.

- Me, i.e. the intuition-haver (Rosen and Dorr 2002: 159)
- Organisms (Carmichael 2015: 475)
- Cats (Contessa 2014: 199; Markosian 1998: 211)
- Dogs (Contessa 2014: 199)
- Apples (Contessa 2014: 199)
- Tables (Wallace 2011a: 804; Carmichael 2015: 475; Korman 2015b: 27; Contessa 2014: 199)
- Chairs (Wallace 2011a: 804; Carmichael 2015: 475; Markosian 1998: 211, 2008: 351)
- A model of the Salisbury Cathedral made out of ten thousand wooden blocks, on the assumption that we have arranged those blocks in an appropriate manner (van Inwagen 1990a: 35)
- Bicycles (Markosian 1998: 211)
- A house of cards (Rosen and Dorr 2002: 152)
- A stack of blocks, when the component blocks are stacked on top of each other (Dorr 2005: 247).
- Piles of sand (Carmichael 2015: 489)
- Rocks (Carmichael 2015: 475)
- Stars (Rosen and Dorr 2002: 152; Markosian 1998: 2011)
- Constellations (Carmichael 2015: 489)
- The Milky Way (Rosen and Dorr 2002: 152)

These intuitions strike me as eminently plausible. It is indeed intuitive that all of these things exist.

To this list some philosophers would add objects apparently studied by physicists or chemists: for example, atoms (e.g. Markosian 2008: 352; Rosen and Dorr 2002: 152). Now, it doesn't seem right to say that most of us have the intuition that atoms exist in the same way that we have the intuition that tables or even stars exist. I believe in atoms not because it's directly obvious or apparent that they exist, but rather because I know that best science says they do. So I think the most plausible thing to say here is that we have the intuition that, more or less *whatever the scientists univocally agree exists, exists*. Given that scientists univocally agree that atoms exist, it's then fair to characterise ourselves as having the *indirect* intuition that atoms exist. I suspect this intuition, so understood, is widespread.

Sometimes philosophers characterise the ontologically committing intuitions they have as not being of the simple form, '*F* exists', but nonetheless as straightforwardly implying such existential claims. Thus Tallant says 'I'm typing this on a computer – so there exists a computer' and Contessa (2014: 199) cites the intuitive truth of 'the cat is on the mat'. Clearly, then, computers, cats, mats could be added to our list if they weren't there already.

Further still, some philosophers appear to claim that they've identified more general intuitions concerning the conditions under which composition occurs (and thus what composite objects exist). Lewis, for example, claims that 'we are happy enough with mereological sums of things that contrast with their surroundings more than they do with one another; and that are adjacent, stick together, and act jointly' (1986: 211), whilst Rose and Schaffer allege that the folk, at least, believe that 'composition occurs in restricted circumstances, in which the question of whether the plurality has a purpose plays a significant role' (2017: 259). Plausibly, however, we again don't *directly* intuit such principles: rather, it seems right to regard them as indirectly intuitive insofar as they amount to extensionally correct generalisations of the firm intuitions we have on particular cases.¹²⁹ As such, we can safely ignore these alleged general intuitions in favour of the particular first-order ones.

Only slightly less obvious than intuitions concerning what composite objects there are are intuitions concerning what material objects there aren't. Intuitively, claim various philosophers, there is no composite object composed out of any of the following collections of objects:

- An unassembled set of parts for an Ikea desk (Kriegel 2008: 359)
- A tree trunk and a dog (Korman 2015b: 27)
- A little finger and the moon (Carmichael 2015: 475-6)
- A nose and the Eiffel Tower (Korman 2015b: 27)
- All of Ned Markosian's shirts and the Eiffel Tower (Markosian 1998: 288)
- Two people and the floor they are both in contact with (Markosian 2008: 352)
- Ned Markosian and his daughter, when Markosian puts his hand on her shoulder (Markosian 2008: 352)
- London Bridge, a certain sub-atomic particle located far beneath the surface of the moon, and Cal Ripken, Jr. (Markosian 1998: 228)
- Two people shaking hands (van Inwagen 1990a: 35; Korman 2015b: 27), no matter whether their hands are paralysed so they can't release their grips (van Inwagen 1990a: 57-8) or superglued together (58), or whether a surgeon has amputated their hands completely and sewn the two people seamlessly together at the forearm (59).

Again, these claims of non-existence are plausible: it is indeed intuitive that there is no such thing composed out of any of the collections of items on this list. And again, we can safely ignore attempts to unify these intuitions under general principles, since the extent to which a given theory will satisfy the intuitive pull of these principles will just be a function of the extent to which it satisfies the intuitive pull of the particular non-existence claims listed above.

Nonetheless, it *is* worth reflecting on the fact that philosophers have historically found it very hard to come up with any general principles that unify all of these intuitions. As a result, there are very few candidate answers to the SCQ that can reasonably claim to be consistent with all of our intuitions concerning what material objects exist. Van Inwagen (1990a: ch. 3, 56-71) demonstrates that several general principles concerning when composition occurs that at first blush appear to do the trick turn out on reflection to get the wrong results. Even Carmichael's proposed answer to the SCQ (2015), which was designed to avoid van Inwagen's counterexamples, doesn't escape counterexamples

¹²⁹ I've no reason to think that either Lewis or Rose and Schaffer would disagree here. Rose and Schaffer's discussion in their (2017) is in particular suggestive of just this sort of claim.

entirely. Carmichael himself concedes that ‘there are some composite material objects that are recognised by common sense, but that are not accommodated by the proposal’, including piles of sand and constellations (2015: 489).

The answers to the SCQ that have received the most attention in the literature also conflict with our existential intuitions. Both Merricks’ (2001) and van Inwagen’s views (1990a: ch. 9) deny the existence of all intuitively existent non-conscious and non-living objects respectively, whilst eliminative nihilism and existence monism deny the existence of *all* intuitively existent composite objects, whether they be conscious, living, or neither of these things. Universalists, on the other hand, appear to be committed to the existence of all of the intuitively non-existent material objects listed above. As I’ll touch on in the next section, there’s some debate as to whether universalism should really be regarded as counterintuitive in this regard (because universalists typically regard themselves as able to say that, for example, ‘there is a nose-Eiffel-Tower’ is false), but for now we can just note that universalism’s acceptance of the existence of nose-Eiffel-Towers, fusions of people shaking hands, etc. certainly *seems* to be counterintuitive.

To my knowledge, only four answers to the SCQ seem to be completely consistent with the deliverances of common sense with regard to what material objects there are and aren’t, and they do so more or less by fiat. This is most obvious in Kriegel’s *normal intuiter account* (NIA):

NIA: the *x*s compose some *y* iff the *x*s are disposed to elicit in most normal intuiters under normal forced-choice conditions the intuition that there is a *y*.

(2008: 367).

More or less: there is a composite *F* iff it’s intuitive that there’s a composite *F*. Clearly, this account will capture our ontological intuitions about material objects.

Brutal Composition will capture our ontological intuitions perfectly too. Or at least, it will if we understand it in the way I indicated I would at the beginning of this chapter. At the beginning of this chapter I said that I would understand Brutal Composition as claiming that (as a matter of brute fact) composition occurs in all and only those cases in which it is intuitive that composition occurs. So understood, Brutal Composition trivially entails that our ontological intuitions about material objects are correct.

A third view that can capture our ontological intuitions concerning material objects is emergence nihilism. This (recall from §2.3) is a form of mereological nihilism, and so accepts that all material objects are simples*, but allows that there can be macroscopic, spatially extended, and partially overlapping simples*. Emergence nihilism can simply stipulate that all and only the material objects of common sense exist *qua* extended simples*. Indeed, in this chapter I’ll assume that emergence nihilism does just this. So understood, emergence nihilism satisfies our intuitions concerning what material objects exist perfectly.

The fourth and final answer to the SCQ that captures our ontological intuitions in this way is, I contend, reductive nihilism. That is, given reductive nihilism, there exists a (composite) *F* iff it is intuitive that exists a (composite) *F*. Argument:¹³⁰

¹³⁰ See also §2.4.2, where I suggest that reductive nihilists can accept the existence of familiar composite objects without accepting the existence of gruesome composite objects. What I say in the main text here amounts to a more rigorous version of the same reasoning.

From right to left. Suppose it is intuitive that there is a table. Then the region of space, r , that is apparently exactly occupied by the table is intuitively exactly occupied by an object. Given the reductive nihilist's semantic claim about the meaning of 'object', (RN2) (see §2.4.2), it follows that the simples* arranged table-wise that the reductive nihilist claims in fact exactly occupy r are jointly an object. Given the reductive nihilist's semantic claim about the meaning of sortal predicates, (RN4), and given that the simples* arranged table-wise are jointly an object, the simples* arranged table-wise are jointly a table. Given the reductive nihilist's semantic claim (in (RN1)) that 'existence' expresses the plural existential quantifier, it follows that there exists a table. This example was arbitrary, so this reasoning generalises to all intuitively existing (composite) F s.

From left to right. I'll prove the contraposition. That is, I'll prove that if it is not intuitive that there exists a (composite) F then that (composite) F does not exist. Suppose, then, that it is not intuitive that there exists a nose-Eiffel-Tower (i.e. the putative composite object composed of a certain nose and the Eiffel Tower). Then the region of space, r , that is exactly occupied jointly by the nose and the Eiffel Tower is *not* intuitively exactly occupied by an object. Given the reductive nihilist's semantic claim about the meaning of 'object', (RN2), it follows that the simples* arranged nose-Eiffel-Tower-wise that the reductive nihilist claims in fact exactly occupy r do *not* jointly qualify as an object. Given the reductive nihilist's semantic claim about the meaning of sortal predicates, (RN4), some simples* can only qualify as a nose-Eiffel-Tower if they are both arranged nose-Eiffel-Tower-wise and are jointly an object. But we've already established that the simples* arranged nose-Eiffel-Tower-wise are not jointly an object. So it's false that there exists some xx such that xx are a nose-Eiffel-Tower. So the nose-Eiffel-Tower doesn't exist. This example was arbitrary, so this reasoning generalises to all (composite) F s that don't intuitively exist.

Reductive nihilism, then, along with three other answers to the SCQ, seems to maximally satisfy our intuitions concerning what composite objects exist.

6.4 – What does it take to satisfy our ontological intuitions?

At this point it's apposite to address a crucial objection to the claim that reductive nihilism can capture our ontological intuitions concerning composite objects. Two *prima facie* natural ways of making this objection are:

- (1) *Strictly speaking*, reductive nihilism denies the existence of all composite objects; similarly, it's not *strictly and literally true* on reductive nihilism that any composite objects exist.
- (2) Reductive nihilism can tell a semantic story that guarantees the truth of the right existential sentences, but that's not enough to satisfy our intuitions that some composite objects *really exist*, that there are composite objects *out there in the world*.

Either way, the thought is that some crucial intuition-satisfying ingredient is missing from the reductive nihilist recipe. The point of this section is to argue against this suggestion. I'll do so by first arguing that the reductive nihilist *can* in fact say the things that (1) and (2) allege she can't say concerning the existence of composite objects, before considering a more nuanced alternative to (2) that, at first pass at least, poses more serious problems for reductive nihilism.

6.4.1 – Strictly and Literally

As Sider puts it, “‘literal’ is opposed to things like metaphor and hyperbole’, and “‘strict’ casts off things like quantifier domain restriction and loose talk’ (2013a: 252). Thus the objection to reductive nihilism here is apparently a disjunctive one: when reductive nihilists affirm the existence of

composite objects such as tables, either reductive nihilists are engaging in metaphor or hyperbole, or they're employing quantifier domain restriction, or they're engaging in loose talk.¹³¹

Now, it's clear that the reductive nihilist is not engaging in metaphor or hyperbole when she affirms the existence of tables. Nor is she employing quantifier domain restriction. But might she be engaging in loose talk?

Some (e.g. O'Leary-Hawthorne and Cortens 1995: 156, citing Kripke 1976: 379-80) suggest that we can't make sense of what it would be to engage in loose talk when making existential claims. But we can give some examples to try to get a fix on the concept. Sider's own example is 'when people who live in Cherry Hill, New Jersey say they're "from Philadelphia"' (2013a: 252). Two less US-centric examples along the same lines: if someone from the Netherlands but *not* from either North or South Holland were to be said to be 'from Holland'; if there was a Dyson vacuum cleaner in the corner of the room and a British person (who uses 'hoover' where they mean to refer to vacuum cleaners of any make and model, not just those made by Hoover) were to say 'there's a hoover in the corner'. Such instances of loose talk can be characterised as being 'nearly as good as true' (Merricks 2001: 171), in order to denote their alethic status of falling short of truth but being nonetheless in some sense alethically better off than clear falsehoods (such as the claim that the Dutchman is from Mars, or that there's a dragon in the corner of my room).¹³²

Now, if the reductive nihilist really was engaging in loose talk when she affirmed the existence of things like tables, then it does seem as though she wouldn't really be capturing our ontological intuitions properly. It seems to me, as it does to many metaphysicians, that our intuition that there are tables requires that it be *true*, not just *nearly* (as good as) true that there are tables.¹³³ That's why I claimed above that eliminative nihilists, who can at best claim that our ontological intuitions are false but nearly as good as true, fail to satisfy our intuitions. But as it happens, it's *not* the case that reductive nihilists are engaging in loose talk when they affirm the existence of tables and the like: it falls out straightforwardly from the combination of metaphysics and semantics of reductive nihilism that claims of the existence of intuitively existing composite *F*s are simply *true* (thus nothing about my argument in §6.3 above for the claim that, given reductive nihilism, there exists a (composite) *F* iff it is intuitive that exists a (composite) *F* relied on weakening the alethic status of that claim to anything less than full-blown truth).

So it's just not the case that reductive nihilists can't accept that it's strictly and literally true that the composite objects of common sense exist.

6.4.2 – *Truth or the World?*

A different way to try to precisify the thought that reductive nihilists can't capture our ontological intuitions would be to say that merely securing the truth of e.g. 'tables exist' is inadequate: what's

¹³¹ Perhaps you think 'strict' and 'literal' should be interpreted in some other way instead. But I'm not sure what that other way could be. At any rate, my discussion throughout §6.4 covers all of the ways that I can think of of accusing the reductive nihilist of not properly capturing our ontological intuitions, so whatever is meant by 'contra the reductive nihilist, composite objects exist *strictly and literally*' should be addressed somewhere here.

¹³² You might still insist that the notion of loose talk so illustrated remains obscure, and that the examples given are simply false, or even true (as you might say if you thought that, in some contexts, 'Holland' now just refers to the whole of the Netherlands and 'hoover' (in the mouth of a Briton, at least) can now refer to any brand of vacuum cleaner). So much the worse for the critic of reductive nihilism here.

¹³³ E.g. Markosian (1998: 220-1, 2008: 350-1), Rosen and Dorr (2002: 158), Merricks (2001: 166ff). One dissenter is Sider (2013a: 237-8, 255-6).

really needed to satisfy our intuition here is simply the *existence of tables*. Occasionally this sort of thought seems to find expression in the literature. Here's Markosian:

'Speaking for myself, at least, the relevant intuition is not merely that sentences like (2) [i.e. 'There is a chair in the corner'] are sometimes, loosely speaking, true; rather, the relevant intuition is that *there really are* such composite objects as stars and chairs [...] That is, according to my intuitions, there simply are far more composite objects *in the world* than Nihilism allows.'

(1998: 220-1, emphasis mine).

'The objection is not merely that if we accept UC [i.e. universalism] then we shall have to say that many typically pronouncements of common sense – such as that there's no object composed of Mia Hamm and The Taj Mahal – turn out to be *false* [...] Rather, the objection is that UC commits us to *the existence* of many strange objects.'

(2008: 346, emphasis mine).

Or, for a non-Markosian example:

'I must say that the whole enterprise of "respecting common sense" by finding some unexpected truth in its utterances about chairs is a bit puzzling. At the end of the day, no one is going to be assuaged by being told that there is a way of understanding talk about chairs on which it comes out true [...] but there aren't REALLY any chairs. People think there really are chairs. (Really – just ask them.) *That* is the common-sense opinion.'

(Richard 2006: 174).

The truth or falsity of certain existential sentences is, you might think, neither here nor there. To satisfy our ontological intuitions tables need to *really be there, out there in the world*. Reductive nihilism, goes the thought, cannot accept this.

But this cannot be right. The disquotation principle (for all propositions/sentences p , ' p ' is true iff p) guarantees that no wedge can be driven in between a theory accepting a certain sentence as true and that theory claiming the world to be the way described by that sentence. In accepting the truth of e.g. 'there are tables', the reductive nihilist thereby claims that the world is such that it contains tables and commits herself to the existence of tables. As such, it's a red herring to focus on the distinction between claims about what sentences are true and claims about the way the world is when trying to think through why reductive nihilism might not be able to satisfy our ontological intuitions.^{134,135} There really are tables, out there in the world, given reductive nihilism.

6.4.3 – Another Try

Nevertheless, it's hard to shake the nagging suspicion that the fact that the reductive nihilist can secure the truth of e.g. 'there is a table', 'there are stars', 'atoms exist' etc. and the falsity of 'there exist nose-Eiffel-Towers' and 'there exist fusions of people shaking hands' etc. is somehow not enough to properly satisfy our ontological intuitions concerning composite objects.

¹³⁴ In fairness, Markosian does sometimes seem to recognise that getting the right truths suffices for satisfying our intuitions (e.g. 'the intuition is that it's literally true [...] that there's a chair in a corner' [2008: 351]).

¹³⁵ See also Sider (2013a: 253-7; 2013b: 761).

To see how we might make progress in making better sense of this objection to reductive nihilism, it's helpful to first consider a line of argument levelled against the *universalist* who purports to capture our ontological intuitions perfectly. Universalists typically employ the *quantifier restriction strategy* to demonstrate that, although she includes gruesome composites such as nose-Eiffel-Towers in her ontology, she can still agree that, in most contexts at least, e.g. 'there are no nose-Eiffel-Towers' is true, because 'there is' in most contexts expresses an existential quantifier that ranges not over the unrestricted domain but over a restricted *sub-domain* that excludes gruesome composites like nose-Eiffel-Towers (see e.g. Lewis 1986: 213; Sider 2001: 218). Universalists who adopt this strategy are widely regarded to nonetheless offend against intuition, because, allegedly, not only do we have the intuition that there are_{restrictedly} no nose-Eiffel-Towers, (where 'there are_{restrictedly}' expresses a quantifier that ranges over an appropriately restricted domain), we *also* have the intuition that there are_{unrestrictedly} no nose-Eiffel-Towers, a result with which the universalist cannot agree (Korman 2008a: 328-9; Markosian 1998: 228-9).

Perhaps we can raise an analogous argument against the reductive nihilist. Take the sentence, 'there are table_{entity}s', where 'is a table_{entity}' is just like the predicate 'is a table', except that 'table_{entity}' is stipulated to only apply to *single entities* that satisfy 'is a table'. 'There are table_{entity}s' is false on reductive nihilism (because the only single entities the reductive nihilist posits are individual quarks and leptons). But you might think that 'there are table_{entity}s' is like 'there are_{unrestrictedly} no nose-Eiffel-Towers': it's all very well that reductive nihilists can agree that there are tables (just as it's all very well that the universalist can agree that there are_{restrictedly} no nose-Eiffel-Towers), but it's *also* intuitive that there are table_{entity}s (just as it's also intuitive that there are_{unrestrictedly} no nose-Eiffel-Towers). Thus, the argument might go, the fact that reductive nihilism secures the truth of 'there is a composite *F*' when and only when *F* is an intuitively existing composite is not enough to fully capture our ontological intuitions, because to fully capture our ontological intuitions the reductive nihilist would also have to secure the truth of 'there is a composite *F_{entity}*' when and only when *F* is an intuitively existing composite.

Reductive nihilists should respond by simply denying that we have intuitions like 'there are table_{entity}s'. I think her denial of this is very plausible. First, recall the distinction reductive nihilists make between 'entity' and object': 'object' is to be understood as an ordinary term that almost all English speakers are familiar and competent with (§2.4.2); 'entity' is instead a technical term of metaphysics that is closely bound up with the notions of *being one in number according to the kind of counting that we do when we track reality's objective cardinality* and *being apt for being the value of a singular variable of first-order logic* (§2.2.1). Given that 'entity' is an obscure term of metaphysics, it simply doesn't seem plausible that we have intuitions one way or the other about whether there are table_{entity}s, as opposed to tables. Indeed, it seems to me that any temptation to agree that it is intuitive that there are table_{entity}s is rooted in the conflation of 'there are table_{entity}s' with 'there are table_{object}s' (where 'is a table_{entity}' only refers to tables that are additionally *entities*, 'is a table_{object}' instead only refers to tables that are additionally *objects*). Certainly, when I reflect on my intuitions about the existence of tables, I find that to the (meagre) extent that 'there are tables' feels like an inadequate expression of my intuition, 'there are table_{object}s' seems like the right alternative. If a theory can give us that there are tables, and that there are table_{object}s, it has surely done enough to capture our ontological intuitions about tables.¹³⁶

¹³⁶ The interaction between 'entity', 'object', and our intuitions pops up in multiple places throughout this thesis. In particular, see §3.5 and §6.5.3.

That accounts for one way in which the argument that reductive nihilists can't capture all of the intuitive ontological truths is much less forceful than the analogous anti-universalist argument. For it *does* seem plausible that, even after we've reflected on the possibility that our intuition concerning the non-existence of nose-Eiffel-Towers is best expressed by 'there are_{restrictedly} no nose-Eiffel-Towers', we still insist that there are_{unrestrictedly} no nose-Eiffel-Towers (Korman 2008a: 328-9; Markosian 1998: 228-9).¹³⁷

Thus, whereas it seems fairly reasonable to think that the universalist can't satisfy all of our intuitions concerning what composite objects there are and aren't, the objection that reductive nihilists cannot capture all of these intuitions should not convince us. It's highly plausible that the claims the reductive nihilist can accept, such as that there are tables and that there are table_{objectS}, suffice to capture our ontological intuitions fully.

6.4.4 – Summary

Reductive nihilists can agree that it's strictly and literally true that intuitively existing composite objects such as tables exist, and that these objects therefore really exist out there in the world. There are no further ontological intuitions (or disambiguations of our ontological intuitions) that the reductive nihilist cannot capture. I can't see any reason not to conclude, then, that reductive nihilism completely satisfies our ontological intuitions.

Currently, then, as I've said, reductive nihilism is tied for first place in the intuition-satisfying game with Brutal Composition, emergence nihilism, and NIA. Those who think that our intuitions concerning the non-existence of gruesome composites are fully captured by the claim that it's false that there are_{restrictedly} no gruesome composites may regard universalism as being equally intuitive so far too. But my discussion is not yet complete. Ontological intuitions do not exhaust all of the intuitions we have concerning composition, and so we now have to see how reductive nihilism stacks up against rival views when it comes to these further intuitions.

6.5 – Intuitions on the Nature of Material Objects

In addition to intuitions concerning what material objects there *are*, we also have intuitions concerning what material objects are *like*.

6.5.1 – Material Objects Have Parts

First, it's intuitive that at least some material objects, such as tables, human beings, bicycles, etc., are *composite*, i.e. they *have parts* (Korman 2015b: 28; Hewitt 2012: 829 fn 28; Liggins 2008: 190; Markosian 1998: 220, 2008: 349; McDaniel 2007: 243 fn5; Rosen and Dorr 2002: 159; Wallace 2011a: 804, 2014a: 115).

Most views on composition satisfy this intuition. In particular, reductive nihilism satisfies this intuition. Reductive nihilists accept that all material objects with the exception of quarks and leptons are composite. Given the reductive nihilist's semantics (§2.4.2), a material object has proper parts iff it is composite, and it is composite iff it is a proper plurality. Since the reductive nihilist takes

¹³⁷ There is additionally some empirical evidence to suggest that the folk similarly have the intuition that there are_{unrestrictedly} no gruesome composites such as nose-Eiffel-Towers (Rose and Schaffer 2017: §3.6). See Rosen and Dorr (2002: 156-7) and Thomasson (2007: 183-4) for the contrary claim that we do *not* have the intuition that there are_{unrestrictedly} no gruesome composites such as nose-Eiffel-Towers.

tables, human beings, bicycles, etc. to be material objects that are proper pluralities, they take all of these things to be composite, and thus to have proper parts.

Not all views on composition satisfy this intuition, though: emergence nihilism accepts the existence of material objects such as tables, human beings, and bicycles, but claims that all of these things are simple*.

6.5.2 – Mereological Essentialism and Sufficiency are False

Second, it's intuitive both that at least some material objects can survive the gain and loss of at least some of their parts (Cameron 2008: 8, 2014: 93; Merricks 1999; Wallace 2011a: 805; Wallace 2014a: 114¹³⁸) and that at least some material objects can exist without being a part of the whole of which they are actually and currently part (Cameron 2008: 8, 2014: 93; Wallace 2014a: 114¹³⁹). In other words, it's intuitive that mereological essentialism and mereological sufficiency are false.

All views on composition seem to be compatible with the falsity of these views. Though doubts have been raised about whether views that accept CAI (and that there are at least some composite objects) are compatible with the falsity of mereological essentialism and mereological sufficiency (Merricks 1999, Cameron 2014: 93), the supporter of CAI has defensive resources available to her to answer these doubts. I discussed these defensive resources in §5.5, and additionally argued there that the reductive nihilist can utilise (adapted versions of) them to similarly reject mereological essentialism and sufficiency. I direct the sceptical reader to that discussion. Here I proceed on the assumption that views that accept CAI, reductive nihilism, and indeed all other views under discussion can satisfy our intuitions about the falsity of mereological essentialism and mereological sufficiency.

6.5.3 – The Result of Composition is Always an Object

Third, it's intuitive that the result of composition – i.e. *what gets composed* – is in each case a (single) object (Korman 2015b: 48; Markosian 1998: 220). That is, in every case in which some objects compose something, the thing that they compose is a (single) object. For example, it is intuitive that composites such as tables, chairs, human beings, etc. are all objects.

All non-nihilist views seem to accept that when composition occurs, what gets composed is an object; at least, they are all able to accept this. Typically (perhaps ubiquitously) they will do so by equating 'object' with 'entity', and accepting that the result of composition is in each case a single entity.

Nihilist views that deny that composition ever occurs also arguably satisfy our intuition here: given such views the conditional 'if some objects compose then what they compose is an object' is trivially true, in virtue of having a false antecedent.

Further, reductive nihilism – the only nihilist view that accepts that composition occurs – also accepts that the result of composition is in each case an object. Indeed, this is built into the reductive nihilist's definition of composition (§2.4.3.3). We might also add that the reductive nihilist take composites such as tables, chairs, human beings, etc. to each be objects (see the discussion of (RN2) in §2.4.2).

¹³⁸ Wallace would only accept this claim with the clarification that 'material object' here is restricted so as to refer only to world-bound objects. Her (2019) (see also her 2014a, b) defends a view according to which ordinary material objects are transworld fusions of world-bound objects, in the same way that four-dimensionalists think that ordinary objects fusions of instantaneous time-slices.

¹³⁹ Again (see previous footnote) Wallace would only accept this claim if 'material object' here means 'world-bound material object'.

One might object that philosophers who report the intuition that the result of composition – tables, chairs, etc. – are objects really mean to report the intuition that in the reductive nihilist’s terms is expressed as: the results of composition are *entities*. The reductive nihilist can’t agree with *this* intuition. But I don’t think the reductive nihilist (nor anyone else for that matter) should agree that this is a genuine intuition. Her reasoning here should echo her reasoning, given above (§6.4.3), for her denial that ‘there is a table_{entity}’ is intuitive: first, ‘entity’ is a metaphysical term-of-art that seems too removed from ordinary language for us to have any strong intuitions about it; second, any intuitive pull that the insistence that tables are single *entities* may have had seems to evaporate when we are reminded that it has already been granted that tables are single *objects*.

So reductive nihilism, I submit, satisfies our intuition that the results of composition (tables, chairs, human beings, etc.) are single objects perfectly well.¹⁴⁰

6.5.4 – Co-location is Impossible

Fourth, it seems to me to be intuitive that no material entity can exactly co-locate with a plurality of entities taken together. That is, it seems to me to be intuitive that if some material entity *e* exactly occupies some region of space *r*, then it can’t be the case that some proper plurality of material objects *ee* also jointly exactly occupy *r*.

Why think that it is intuitive that that no material entity can exactly co-locate with a plurality of entities? Consider first what is often held to be counterintuitive about the idea of a material entity exactly co-locating with another *single* material entity. The thought here is that an entity that occupies a region of space intuitively leaves *no room* for another entity to occupy that region of space. Suppose, as almost all non-reductive-nihilists do, that tables are entities, and consider trying to get two tables that are identically-sized and -shaped (but that are nevertheless distinct) to occupy the same region of space. This task simply seems impossible. That feeling of impossibility is what underlies the intuitive dubiousness of exact co-location between one material entity with another. This anti-co-location intuition is additionally sometimes to be found in the literature.¹⁴¹ For example, Wasserman refers to ‘the common idea that impenetrability [i.e. the idea that material entities cannot be co-located with] is the mark of the material’ (2018: §2) (see also Zimmerman 1996: §2; Locke [1690] 1959: Bk 2 ch. 27, §§1-2).

Now, there seems to be no good reason why it would be counterintuitive to say that one material entity can exactly co-locate with another but *not* counterintuitive to say that one material entity can exactly co-locate with a plurality of other entities taken together. The same considerations seem to apply. We just considered the task of trying to get two identically-sized and -shaped tables to exactly

¹⁴⁰ Of course, it’s possible that there will be philosophers who accept that ‘entity’ is a metaphysical term-of-art, and who accept that it has already been granted that things like tables are single objects, but nevertheless insist that they find it intuitive that the results of composition are single objects. Short of claiming that such philosophers are mistaken about their own intuitions, the only thing that I can think of to say in response is that even if my arguments against the idea that it is intuitive that the results of composition are each single entities fail to convince entirely, they should at least make it plausible that the intuition that the results of composition are each single entities is not a particularly *strong* or *indispensable* one. And if this intuition is indeed fairly weak and dispensable, then my overall argument in this chapter should remain unaffected: reductive nihilism has enough of an intuitive advantage on its rivals that violating this weak intuition does not forfeit that advantage.

¹⁴¹ I note that discussions in the literature on this topic understandably do not seem to make the same distinction between ‘object’ and ‘entity’ that the reductive nihilist does, and so it seems appropriate to interpret claims made in the literature on co-location about material *objects* as claims about what the reductive nihilist would insist should be called material *entities*.

occupy the same region of space, and we suggested that it was impossible. Would our task be any easier if instead we were to try to get a table to co-locate exactly with a random proper plurality of entities that, taken together and arranged in the right way, would be jointly the same size and shape as the table? It seems obvious that the answer is no. So it seems that exact one-many co-location is intuitively objectionable if (and only if) exact one-one co-location is (for the same point see Cameron [2014: 90]; Lowe [1995: 177]; Wallace [2011a: 804-5]).

Perhaps it will be objected that the following caveat is in order: if a plurality of entities *compose** a material entity then the exact co-location of that plurality and that entity is intuitively unproblematic. But my intuition continues to find exactly co-location between a plurality and an entity problematic even in cases in which the plurality composes* the entity; there is furthermore at least some evidence of other philosophers sharing my view here – see especially Wallace (2011a: 804-5); see also Cameron (2014: 91). What’s more, we can perhaps offer the following reason to be sceptical of philosophers who insist that co-location is intuitively unproblematic when it is between a whole and the parts* that compose*. Given the widespread assumption among metaphysicians that composite objects are each *entities* that exactly co-locate with their parts*, saying that composite tables and chairs and human beings (etc.) exist requires accepting the exact co-location of each composite object and its parts* taken together. Maybe, then, philosophers have simply grown so used to the idea that they must accept co-location between parts* and whole in order to be allowed to posit the whole in the first place that they simply cease to notice or report that they find it counterintuitive. It is, after all, perhaps easy to confuse a claim that is somewhat counterintuitive but nevertheless apparently absolutely indispensable with one that wasn’t counterintuitive to begin with. Perhaps what has happened is that philosophers find the postulation of tables, chairs, etc. to be indispensable, and that they are under the impression that this postulation requires accepting co-location between an entity one the one hand and a plurality of entities on the other, that they have become subject to this very confusion.

What is the upshot of this discussion of co-location and intuitions? The first thing to say is that reductive nihilism avoids any kind of counterintuitive co-location. The only entities that reductive nihilists posit are simples*, and reductive nihilists can maintain that simples* do not co-locate. Though reductive nihilists additionally posit composite *objects* (that are not entities) that apparently co-locate with their parts taken together, since reductive nihilists take parts (taken together) and whole to be many-many identical (see §2.4.3.3), this is no more a genuine case of co-location than is the case of Bob Dylan and George Orwell ‘co-locating’ with Robert Zimmerman and Eric Blair.

Next, we should admit that some philosophers will maintain that there is nothing counterintuitive about the exact co-location of a material entity one the one hand and a plurality of material entities on the other. They will find my discussion of intuitions surrounding co-location to be completely unhelpful.

Others will maintain that it is counterintuitive that a material entity can exactly co-locate with a plurality of entities *unless* the plurality composes* the entity, in which case the co-location is not counterintuitive at all. These philosophers should at least grant that reductive nihilist can claim (another) intuitive advantage over *emergence nihilists*. For emergence nihilists accept that e.g. a certain table exactly co-locates with a plurality of simples* arranged table-wise, but they explicitly deny that those simples* compose* (or compose) the table.

It is perhaps unlikely that many will go further and grant that a material entity exactly co-locating with a plurality of entities is counterintuitive *even if the plurality composes* the entity*. But I find the

co-location of a material entity and its parts* counterintuitive, there's evidence that other philosophers do too (see the Wallace and Cameron citations above), and I've provided some reason to be sceptical of philosophers who claim not to. And if we accept that even the co-location of a material entity and its parts* is counterintuitive, then any view that accepts the existence of composite* entities without identifying those entities with their parts* is counterintuitive in a way that reductive nihilism is not.¹⁴² So existence monism and eliminative nihilism escape, as do views that accept CAI, but universalism and all forms of restrictivism would be counterintuitive in this way.

Ultimately, given the possibility that my intuitions about co-location are somewhat idiosyncratic, it would be a mistake to hang too much weight on them here. As we'll see, my argument for reductive nihilism's victory in the intuitions stakes would go through even without any appeal to intuitions concerning co-location. But it is nevertheless worthwhile to note that anyone who shares my intuition that there can be no exact co-location between a material entity on the one hand and a plurality of material entities on the other (even between an entity and its parts*) should take reductive nihilism's claim to be the most intuitive view of composition to be further strengthened.

6.5.5 – Material Objects Exist Mind-Independently

A fifth and final intuition concerning the nature of material objects is the intuition that material objects exist *mind-independently*. Intuition may allow some exceptions to this: for example, you might think that it offends against intuition not at all to say that the existence of a sculpture, depends partially on whether the lump of clay out of which it was made was shaped by an artist with the intention of creating a sculpture. But set aside such cases, and focus instead on mundane cases such as that of the table, or the apple, or the human body. Whilst it may be right to say that (for example) the fact that the table counts as a table (rather than, say, as a shelter, or a bed, etc.) depends on facts about the way we conceptualise it, the fact that the object that we think of as a table exists *in the first place* – and is thus available for being classified as a table, or a shelter, or a bed, etc. – is intuitively a mind-independent fact.¹⁴³

Almost all views on composition seems to be able to satisfy the intuition about the mind-independence of material objects. For most views, minds play no role in determining whether material objects exist or not. Instead, most views can say that simples* exist as a matter of brute fact (and so simples* don't depend for their existence on anything), and that composite objects depend for their existence on the existence of their parts and on certain mind-independent facts about those parts that determine whether they compose (for example, facts concerning how closely bonded those parts are to one another, or facts about whether those parts jointly constitute a life, etc.).¹⁴⁴

The exception to this is Kriegel's NIA. Here it is again:

NIA: the *x*s compose some *y* iff the *x*s are disposed to elicit in most normal intuiters under normal forced-choice conditions the intuition that there is a *y*.

(Kriegel 2008: 367).

¹⁴² Why think that identifying composite* entities with their parts* alleviates worries about the counterintuitiveness of their co-location? Because identity in general seems to preclude such worries. For example, the co-location of the Morning Star and the Evening Star is completely unproblematic.

¹⁴³ See Page (2006) for a helpful disambiguation of various sorts of mind-independence.

¹⁴⁴ For brevity I'll ignore the priority-monistic view that objects instead depend for their existence on the wholes of which they are part (see Schaffer 2010a) – priority monism doesn't raise any salient issues here.

Now, judging by this statement of NIA alone (which is taken verbatim from Kriegel), the view isn't committed to material objects existing mind-dependently. This statement of NIA is a simple biconditional, and doesn't make any claims about dependence at all. But despite the way he formulates his view, it's clear that Kriegel intends NIA to be committed to the claim that facts about whether some *x*s compose a *y* is a mind-dependent matter. For example, throughout his discussion of NIA he characterises it as the view that composition is a secondary quality, in the same sense that colours are sometimes thought to be secondary qualities (2008: 367-8, 372; see also the title of his paper, 'Composition as a Secondary Quality'); he then characterises secondary qualities as 'properties whose instantiation *depends* in some way on subjects' (368 fn19) (for further evidence that Kriegel takes NIA to be committed to mind-dependence in this way see e.g. [369-70]) And given this understanding of NIA, the view is committed to the existence of composite material objects being a (partially) mind-dependent matter. For on the view so understood, the fact that there is a composite object composed out of some parts depends on facts about the intuitions of normal intuiters.¹⁴⁵

Finally, what about reductive nihilism? On reductive nihilism, composite objects are just certain proper pluralities of simples*. There's no reason why reductive nihilists can't take these pluralities of simples* to exist mind-independently.

6.5.6 – Summary

In §6.4.5 I said that reductive nihilism was at that point tied for first place in the intuitiveness stakes with Brutal Composition, emergence nihilism, NIA, and *perhaps* universalism. It's clear from my discussion in this section that emergence nihilism and NIA have fallen behind the rest of the pack: emergence nihilism, because it counterintuitively denies that things such as tables and human beings are composite; NIA because it counterintuitively denies that material objects exist mind-independently. All three of reductive nihilism, Brutal Composition, and universalism can accept our intuitions on these matters. They can also deny mereological essentialism and sufficiency, and accept that the result of composition is in each case a single object. When Brutal Composition and universalism are combined with CAI, they can also join reductive nihilism in denying that it's ever the case that some entities collectively co-locate with the entity they compose.

To try to break the tie, and to be sure that no other view can claw back lost ground, we must look to the next class of intuitions.

6.6 – Intuitions Concerning the Nature of Composition

We seem to have two main intuitions concerning the nature of composition: that composition is not brute, and that composition is, in certain to-be-defined respects, an intimate relation.

6.6.1 – Composition is Not Brute

First, it's intuitive that compositional facts – i.e. facts concerning whether some objects compose or not – are not brute. Here is a typical statement of the intuition:

¹⁴⁵ I won't consider here whether NIA might instead be understood as repudiating the mind-dependency claim. Without that claim, NIA would be too underspecified to assess. It would give no account of what it is in virtue of which composition occurs, which means that it would give no account of how it is that compositional facts are exactly as intuition (or the intuition of a normal intuiter) takes them to be (and if it says that facts about composition are brute then it would collapse into Brutal Composition).

‘Even though explanation presumably must 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 general principles. Rather, if one bunch of physical simples compose a genuine physical objects, but another bunch of physical 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.’

(Horgan 1999: 695 emphasis original)

There’s plenty of evidence in the literature that this intuition is widespread (see e.g. Cameron 2014: 91; Korman 2015b: 7; Horgan and Potrč 2000: 11-2; Sider 2001: 122).¹⁴⁶

Now, one might think that Brutal Composition would by definition fail to satisfy the intuition that compositional facts are not brute. But that depends on how we interpret the ‘anti-brutal’ intuition. The way that Cameron (2014: 91) and Sider (2001: 122) formulate this intuition in particular suggests that we should understand it as the intuition that facts about whether some *x*s compose or not *supervene* on non-mereological facts such as the spatial proximity of the *x*s to one another, the extent to which the *x*s are fastened to one another, the extent to which the *x*s are causally integrated, the extent to which the *x*s are jointly distinct from their surroundings, and so on. But, as Markosian points out, Brutal Composition theorists can comfortably reject the principle he dubs ‘The Non-supervenience of Composition’ and maintain that ‘even though compositional facts are brute facts, it is nevertheless impossible for two worlds to be duplicates with respect to non-mereological universals but differ with respect to composition’ (1998: 216).

What this should demonstrate to us is that the formulation of the anti-brutal intuition in terms of supervenience is inadequate. Indeed, formulating the claim that some set of facts are brute in terms of supervenience should always have seemed dubious. Consider the claim that mathematical truths are brute. This claim seems perfectly consistent with the trivial supervenience of mathematical truths on all other truths. So brutality should not be understood as the denial of supervenience. Instead, the claim that mathematical truths are brute is much more aptly understood as the denial that mathematical truths *depend on*, or *hold in virtue of*, or *are explained by* any other truths.¹⁴⁷ We should thus understand the intuition that compositional facts are not brute as the claim that compositional facts depend on, hold in virtue of, or are explained by other non-compositional facts. That’s why, for example, in the quoted passage from Horgan above, he articulates his intuition that compositional facts are not brute in terms of it being incredible that specific compositional facts be ‘utterly *unexplainable*’ and ‘at the *explanatory bedrock*’ (1999: 695, emphasis mine).

So understood, Brutal Composition violates our intuition that compositional facts are not brute, since it explicitly denies that compositional facts hold in virtue of any other facts (Markosian 1998: 215).

¹⁴⁶ Kriegel, attributing the point to Kristie Miller, additionally suggests a reason *why* we have this intuition. He says that it is the result of the more basic intuition that ‘the man on the street can and often does know’ when there is a single object composed out of some smaller objects (2008: 364): if compositional facts were brute, then although the man on the street might have *true beliefs* concerning the nature of composition, it’s hard to see how those beliefs would be *justified* enough to count as knowledge, since there’d be no clear account of how we’re able to track those facts so well (Kriegel 2008: 365). That said, my discussion here doesn’t depend on any particular account of why it is intuitive that compositional facts are not brute.

¹⁴⁷ In a similar vein, Fahrbach defines brute facts as ‘facts that have no explanation’ (2005: 449).

All other views of composition are consistent with this intuition. In particular, reductive nihilism is consistent with this intuition. Reductive nihilists take facts about whether some objects compose anything or not to be a question of whether there is any xxx and yy that satisfy COMPOSITION, which I repeat here from §2.4.3.3:

COMPOSITION: xxx compose yy =_{df}

- (v) $\forall xx$ (if xx are superplurally among xxx then xx is/are a part of yy) AND
- (vi) $\forall xx \forall zz$ (if xx are superplurally among $xxx \wedge zz$ are superplurally among $xxx \wedge xx \neq zz$, then xx doesn't overlap zz) AND
- (vii) $\forall xx$ (if xx is/are a part of yy then $\exists zz$ (zz are superplurally among xxx and xx overlaps zz)) AND
- (viii) yy are an object

Thus the reductive nihilist takes compositional facts to depend on (/hold in virtue of / be explained by) the satisfaction of (i)-(iv). And though (i)-(iv) apparently make reference to the mereological notions of *parthood* and *overlap*, the reductive nihilist defines these notions in terms of the non-mereological notions of *amonghood* and *having a subplurality in common* respectively (see §2.4.2). So reductive nihilists do not take compositional or mereological facts to be brute: at the explanatory bedrock she requires no reference to compositional or mereological facts.

6.6.2 – Composition Is an Intimate Relation

It's also very intuitive that composition is an unusually *intimate* relation, in that the relationship between some parts and the whole they compose is a very close one (see especially Sider 2007: §2, §4; Cameron 2014). This claim is typically fleshed out in two main ways.

First, it seems intuitive that a whole inherits certain properties from its parts, most obviously *locational* properties (a whole is necessarily located wherever its parts are) (Cameron 2014: 90; Sider 2007: 70). All of the views that accept the existence of composites appear to be able to satisfy this intuition. All that seems to be required is to stipulate that the relevant property inheritance occurs as a matter of metaphysical necessity, and all views can do that.¹⁴⁸ Reductive nihilists, in particular, can accept this. Composite objects are, given reductive nihilism, identical to their parts taken (see §2.4.3.3), and so trivially inherit whatever properties that their parts, taken together, instantiate.

Second, it's intuitive that a given whole is, in some sense, *nothing over and above* its parts (Sider 2007: 54, 2015: 191; Cameron 2014: 91; Lewis 1991: 81; Hawley 2014: 70; Swinburne 1995: 395). To see why this second claim is supposed to be intuitive, consider the following case, due to Baxter:

‘Suppose a man owned some land which he divides into six parcels. [...] He might try to perpetrate the following scam. He sells of 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

¹⁴⁸ Emergence nihilism, though it cannot accept that wholes inherit properties from their parts, can nonetheless posit similar metaphysical laws that entail that e.g. a (simple*) table inherits its locational properties from some particular table-legs and table-top, etc.

the parcels argue that they jointly own the whole and the original owner now owns nothing. Their argument seems to be right.’

(1988: 579)

The *reason* the argument put forward by the six buyers seems to be right is that it’s intuitive that the whole land *just is*, in some sense, the six parcels of land that compose it. In other words, the whole land seems to be nothing over and above the six parcels of land that compose it.

The intuition that the whole is nothing over and above its parts is harder to respect. In fact, most views of composition seem to *rule out* the truth of this intuition. For most views regard wholes as entities that are not identical to their parts (taken collectively). And if the whole is not identical to the parts, then it’s false that the whole *just is* the parts; it’s false that the whole is nothing over and above the parts. That is, nothing-over-and-aboveness seems to require identity.¹⁴⁹ If a whole is not the same thing as one or all of the parts, then it is a further thing, and therefore something over and above those parts. At least, I think this is *intuitively* the case. Intuitively, that is, composite objects are nothing over and above their parts *in the sense of not being distinct from those parts taken together*. So whilst it is arguably open to theorists to simply *stipulate* that wholes are nothing over and above their parts, despite the distinctness of wholes and parts (as e.g. Lewis [1991: 81-7] in effect does), it seems to me that such views haven’t satisfied the salient understanding of the intuition that composite objects are nothing over and above their parts unless they deny that composite objects are distinct from their parts.

One might object to this: perhaps a whole can be nothing over and above its parts in the salient sense even without being identical to its parts. One way to try to flesh this out is to say that a whole can be distinct from its parts but nevertheless nothing over and above them in the salient sense as long as it is not an *extra ontological commitment* in addition to its parts. Grant this for the sake of argument. The challenge to those who would pursue this objection is then to say how it could be the case that a whole can fail to be an extra ontological commitment in addition to its parts. The only way that I can see of trying to meet this challenge is to appeal to the view that non-fundamental entities don’t count against the ontological parsimony of a theory (Schaffer 2015; see also chapter 7 of this thesis). Bennett has argued that this view makes sense of the idea that non-fundamental entities are nothing over and above fundamental ones (2017: 221-3). If composite objects are taken to be non-fundamental entities, and simples* are taken to be fundamental ones, then it could be maintained in this way that composite objects are nothing over and above their simple* parts in the salient way despite being distinct from them. The problem with this is that it simply isn’t plausible that non-fundamental entities don’t count against the ontological parsimony of a theory. I argue for this at length in chapter 7. So even if it were plausible that the intuitive sense in which wholes were nothing over and above their parts could be captured by saying that wholes are distinct from their parts without being extra ontological commitments to their parts, it isn’t the case that a whole can be distinct from its parts without being an extra ontological commitment to them, so the nothing-over-and-aboveness of wholes couldn’t be secured in this way.

Nor will it do here to complain, as e.g. Cameron has (2014: 99), that claiming nothing-over-and-aboveness to intuitively require identity relies on an objectionable commitment to a Quinean metaontology, such that if we drop the background commitment to Quineanism then we can say that say that wholes are nothing over and above yet distinct from their parts. For even non-philosophers –

¹⁴⁹ Some clarify talk of nothing-over-and-above in terms of identity, suggesting they agree with me here (e.g. Bricker 2016: 266; see also van Inwagen 1994: 213, Markosian 2008: 345).

who as a rule are unfamiliar with Quinean metaontology – surely find it intuitive that wholes are nothing over and above their parts in the sense of not being distinct from their parts.

For these reasons, I think it is plausible to maintain that the intuition that wholes are nothing over and above their parts is unsatisfiable by any view that claims parts and whole to be distinct. That is, this intuition is violated by all views that posit composite objects except for two. Those two are the two views that accept the identity of a composite object with its parts: CAI, and reductive nihilism (§2.4.3.3).

It seems to me that the intuition that composition is an intimate relation, and in particular that wholes are nothing over and above the parts that compose them, is one of the deepest intuitions we have about composition. That reductive nihilism and CAI can satisfy this intuition is surely a significant mark in favour of these views.

6.6.3 – Summary

After considering intuitions concerning the nature of composition, it's clear that competition to reductive nihilism has dwindled significantly. In particular, unlike reductive nihilism, Brutal Composition violates our intuition that compositional facts depend on non-mereological facts. And universalism, when *not* combined with CAI, cannot respect our deeply held intuition that wholes are nothing over and above their parts. Thus the only view that seems to be able to challenge reductive nihilism in the intuitiveness stakes is universalism combined with CAI. In the next section I argue that reductive nihilism is more intuitive than this view too.

6.7 – Reductive Nihilism versus Composition as Identity

I'll begin with a clarification regarding how I'm understanding CAI. Then I'll argue that, despite its intuitive virtues, the combination of universalism with CAI actually offends against intuition significantly. That result would not establish that reductive nihilism is more intuitive than such a combination if reductive nihilism had counterintuitive consequences that CAI doesn't have. So I'll then argue that reductive nihilism has no such analogous counterintuitive consequences, and conclude that reductive nihilism is more intuitive than universalism plus CAI.

6.7.1 – Composition as Identity

CAI, recall, is a candidate answer to van Inwagen's General Composition Question (GCQ) (1990a: ch. 4), which, asks, roughly, 'what is the nature of the composition relation that holds between a whole and its parts?'. CAI answers that composition is strict, Leibnizian identity (because of the revisionary things that CAI sometimes says about our concept of number [see below], I hesitate to use the phrase 'numerical identity', but that's the sort of thing I'm gesturing at here).¹⁵⁰ Thus it's the view that what it is for some objects, x_1, \dots, x_n , to jointly compose a single object, y , is just for x_1, \dots, x_n to be jointly identical to y . The 'jointly' here is important: it's not that each x_i is identical to y , but rather that the x s taken together are supposed to collectively be identical to y . The identity predicate used to express the identity that allegedly holds between x_1, \dots, x_n and y is, we can say, in this way non-distributive in its first place, much like the 'surround' predicate (when some protesters are said to

¹⁵⁰ Thus by 'CAI' I mean what is often called 'strong composition as identity' (e.g. Sider 2007: §3, Cotnoir 2014: 9), or the 'strong composition thesis' (Wallace 2011a: 806).

surround a building, we mean that they do so collectively, not that each protester herself surrounds the building).¹⁵¹

It's important to distinguish CAI from views that sound like CAI but which nonetheless deny that a whole and its parts are strictly identical to one another (e.g. Lewis 1991; Sider 2007). Typically this form of 'weak CAI' holds that composition is *analogous* enough to identity in enough respects to rightly be called a version of identity, but not the same thing as Leibnizian identity, and not governed by Leibniz's Law.¹⁵² But, whatever the merits of this sort of weak CAI might be, it isn't the view we're looking for here. As Wallace puts it,

[...] if composition does not obey the Indiscernibility of Identicals, then despite what other redeeming features this relation has, composition is simply *not* identity. Claiming that a relation is very similar to but not identity entails that such a relation is not identity.'

(2011a: 809)

And if weak CAI says that composition is not identity, then, given our discussion in §6.6.2, it cannot satisfy our intuition that wholes are nothing over and above their parts. Only the strong sort of CAI that I defined above can purport to challenge the reductive nihilist when it comes to intuitiveness, and so I'll only discuss this sort of CAI in what follows.

Now, as we've said above, it's only when we pair CAI with universalism that we get a view that can at least purport to satisfy our intuitions as much as reductive nihilism.¹⁵³ So from now on I'll only consider the combination of CAI and universalism, and I'll use 'CAI-U' to refer to this combination.¹⁵⁴

6.7.2 – The Counterintuitiveness of CAI-U

One reason you might think that CAI-U is counterintuitive is that by being a universalist view it is committed to the existence of intuitively non-existent gruesome composites, but by being a CAI-ist view it *can't* avail itself of the usual universalist quantifier domain restriction strategy to (allegedly) obviate the counterintuitiveness of this. That's because quantifier restriction is *extensional*, in that if a restricted sub-domain contains entity *e* or entities *ee* then it also contains every entity *f* or entities *ff* to which *e* or *ee* are identical: just as we can't restrict our quantifiers so that they range over the morning star but not the evening star, so we can't restrict them so that they range over the table and the donkey but not the table-donkey to which they're supposed to be identical. That means that if CAI-U insists that in ordinary contexts it's false that nose-Eiffel-Towers exist, on the grounds that in those contexts table-donkeys are excluded from the sub-domain over which our quantifiers range, then she also has to objectionably claim that in ordinary contexts it's false that noses and the Eiffel Tower exist, because noses and the Eiffel Tower must then also be excluded from the sub-domain over which our

¹⁵¹ Clearly, CAI requires some grammatical revisionism: usually we take an identity claim to be well-formed only if the terms that flank it are either both singular or both plural, whereas CAI requires us to countenance claims in which '=' takes a singular argument on one side and a plural one on the other. But we can assume here that these grammatical revisions are unproblematic.

¹⁵² Lewis, for example, conceives of ordinary, Leibnizian identity as 'the special, limiting case of identity in the broader sense' (1991: 85), with composition counting as another form of identity in this broader sense as well.

¹⁵³ For example, the combination of CAI and van Inwagen's answer to the SCQ violates the intuition that there are tables (§6.3), the combination of CAI and Brutal Composition violates the intuition that compositional facts are not brute (§6.6.1), and so on.

¹⁵⁴ Calosi (2016) and Loss (2018) have argued that CAI entails mereological nihilism, in which case CAI-U would be incoherent. If so, all the better for reductive nihilism.

quantifiers range. Quantifier restriction as a strategy for mitigating the counterintuitiveness of positing gruesome composite objects therefore doesn't appear to be available to CAI-U.

But if CAI takes away with one hand here, it gives with the other. For might seem to be the case that, given CAI, a commitment to a nose-Eiffel-Tower is no longer counterintuitive, because such a commitment is just a perfectly intuitive commitment to the nose and the Eiffel Tower to which the nose-Eiffel-Tower is identical (Merricks 2005: 630-1). Thus it's not clear that CAI-U fails to capture our ontological intuitions (though see Cameron [2012: 551-2]).

Despite this, CAI-U remains counterintuitive. It's just that its counterintuitiveness (or what is most *obviously* counterintuitive about it) stems not from the *U*, but from the *CAI*. There's a reason that almost no-one accepts CAI, understood as the claim that composition just is identity (rather than the weaker claim that composition is analogous to identity): it just seems intuitively absurd for many things to be identical to one thing. Sometimes this intuition is presented as an argument from Leibniz's Law against CAI, the idea being that the many have a property that the one doesn't have, namely *being many in number*, and the one has a property that the many doesn't have, namely *being one in number* (Lewis 1991: 87, McKay 2006: 38). Now, presumably most supporters of CAI will retort that this rests on the illicit assumption that cardinal properties exclude each other, and insist that parts and whole are each both many and one in number (Wallace 2011b: 820-1, McDaniel 2013: 217-8). The trouble is that this is just implausible, or at the very least, radically revisionary with regard to our usual conception of number, according to which cardinal properties do exclude one another (thus, for example, if the members of the Philosophy department are 15 in number at *t* then they can't also be 20 in number at *t*). On the other hand, if defenders of CAI try to hold on to the exclusivity of cardinal properties by allowing that the parts are many and not one, and that the whole is one and not many, and to nevertheless accommodate the identity of parts and whole by allowing the identicals can be discernible (as Baxter [2014] does), they still offend against intuition significantly, this time with regard to our conception of identity, a central component of which is that identity obeys the indiscernibility of identicals.

As Sider says, then, given CAI's commitment to many-one identities, it is a 'radical thesis' that 'challenges us to reorient our thinking about a wide domain of related concepts' (2007: 63). Since CAI-U accepts CAI, it does significant violence to our intuitions concerning identity and number.

6.7.3 – Reductive Nihilism's Counterintuitive Consequences?

Reductive nihilism does not join CAI-U in positing many-one identities. In particular, though it accepts the identity of a composite object with its parts, reductive nihilism takes this to be many-many, not many-one, identity. So reductive nihilism avoids the counterintuitive consequences of CAI-U just discussed.

The final step in arguing for the claim that reductive nihilism is more intuitive than CAI-U is to argue against the suggestion that reductive nihilism is in turn counterintuitive in a way that CAI-U is not. Now, we've already seen that reductive nihilism satisfies more or less every intuition about composition that we have. It additionally avoids positing many-one identities in order to achieve this; instead, it does so via its distinctive *semantic* claims. So it seems that the only option left open to those who would accuse reductive nihilism of counterintuitiveness is to allege that its distinctive *semantic* claims are somehow counterintuitive.

How might this be done? First, both Merricks (2001: 163ff) and Korman (2015b: 50) have argued that it's intuitive that mereological nihilism contradicts our ordinary beliefs concerning the existence of

composite objects. That is, they take it to be intuitive that if mereological nihilism is true, then many of our beliefs about composite material objects would turn out to be false. Thus Merricks argues (we can read ‘eliminativism’ as ‘nihilism’ here):

‘One reason to oppose [the thought that eliminativism doesn’t contradict our beliefs] stems from eliminativism’s undeniably striking and surprising nature. For – as yourself – *why* is eliminativism striking and surprising? It cannot be because of its revisionary practical or empirical consequences; it has no such consequences [...]. Instead, eliminativism is striking and surprising because – and this is the obvious answer – it contradicts what nearly all of us believe’.

(2001: 163)

The problem for reductive nihilism, then, is that it is a form of mereological nihilism whose semantics imply that our beliefs about composite objects are *true* (as I’ve spent this chapter arguing): because of its semantics, reductive nihilism is counterintuitively intuitive.

But I don’t think that we should agree that reductive nihilists are saying anything counterintuitive here. Consider the quoted passage from Merricks above. The argument contained therein is premised on the claim that eliminativism has an ‘undeniably striking and surprising nature’. But this is question-begging, because it is precisely what reductive nihilists reject. Indeed, the whole point of this chapter has been to demonstrate all of the ways in which reductive nihilism is *not* striking or surprising, and rather fits perfectly well with what we pre-theoretically believe. Furthermore, it seems obvious why some philosophers think that nihilism ought to be a surprising thesis: because they wrongly assume that mereological nihilism is equivalent to *eliminative* nihilism. After all, whenever mereological nihilism is discussed, whether it be in books, journal articles, introductory lectures, research talks, etc., it is invariably introduced as the view that there are no composite objects, or that tables don’t exist, or via variations on this theme. *Eliminative* nihilism is absolutely ‘striking and surprising in nature’. But of course, reductive nihilism is a distinct form of mereological nihilism to eliminative nihilism, and so the fact that philosophers find eliminative nihilism to be striking and surprising gives us no reason for thinking that mereological nihilism or reductive nihilism are striking and surprising. And if we reject that reductive nihilism is a striking or surprising thesis in the first place, then we have no reason to find it intuitive that it contradicts ordinary belief.

A second semantic intuition that might be thought to tell against reductive nihilism is suggested by an objection to reductive nihilism’s semantics that I discussed and rebutted in chapter 5 (see §5.3). There I formulated an objection to reductive nihilism’s semantics by adapting an objection that Liggins (2020) has raised against a different view, hermeneutic fictionalism. The objection to reductive nihilism’s semantics appealed to a principle of interpretive charity, and relied on the claim that reductive nihilism implies that we are wrong about the content of our claims and beliefs about composite objects, in that we mistakenly think these claims and beliefs are about tables, chairs, etc. when they are really about pluralities of simples* arranged table-wise, or chair-wise, etc. That is suggestive of the idea that reductive nihilism is counterintuitive with respect to what it says about the content of our claims and beliefs about composite objects: intuitively, such claims and beliefs are about tables, chairs, etc., but reductive nihilism says they are instead about pluralities of simples*.

In response, I make much the same response that I made in §5.3. Briefly (I direct the reader to §5.3 for a longer explanation): reductive nihilists do take claims and beliefs about tables (for example) to be about simples* arranged table-wise, but they *also* take it to be about tables. Reductive nihilists take

‘table’ to just refer to simples* arranged table-wise, so if some sentence is about some simples* arranged table-wise then it is automatically about tables. Another way of making the point is that, given reductive nihilism, any given plurality of simples* that compose a table are *identical* to that table (§2.4.3.3). So as long as some sentence is about those simples*, it is thereby about that table.

If someone wanted to insist that the reductive nihilist says something counterintuitive here, they’d have to maintain not that it’s intuitive that table-talk is about tables (for reductive nihilists can agree with this intuition), but rather that it’s intuitive that table-talk is *not* about simples* arranged table-wise. But I simply don’t see what reason we have for thinking that this is intuitive at all. Certainly, my intuition regarding the aboutness of sentences such as ‘there’s a table over there’ is entirely satisfied if such sentences turn out to be about tables, and I have no intuition to the effect that such sentences aren’t about sub-atomic goings on as well.

So I conclude that reductive nihilism faces no successful charge of failing to respect our semantic intuitions concerning what our discourse is about.

6.7.4 – Summary

CAI-U says radical, counterintuitive things about the way that identity and number function; reductive nihilism doesn’t accept these counterintuitive claims of CAI, and nor are the semantic claims that are distinctive of reductive nihilism intuitively objectionable. I conclude that reductive nihilism is a more intuitive view of composition than CAI-U, that is, than universalism when paired with Composition as Identity.

6.8 – Conclusion

It’s clear that we have a great many intuitions concerning composition. The most obvious, and most-cited, are those concerning what composite objects there are and aren’t, but these are by no means the only ones. We also have intuitions concerning the nature of composite objects and the nature of the composition relation, as well as intuitions concerning the functioning of identity, number, and semantics. Capturing all of these intuitions is an enormous ask. But it’s one met by reductive nihilism, and reductive nihilism alone:

Unlike eliminative nihilism, existence monism, universalism (unless the quantifier domain restriction strategy is successful), and all restricted views of composition apart from Brutal Composition and NIA, reductive nihilism satisfies our ontological intuitions concerning what composite objects exist.

Unlike emergence nihilism, reductive nihilism satisfies our intuition that ordinary material objects have parts.

Unlike every almost every other view that accepts the existence of composite objects (the exceptions being views that accept CAI), reductive nihilism satisfies what is arguably our intuition that there can be no exact co-location between a single entity on the one hand and a plurality of entities on the other.

Unlike NIA, reductive nihilism satisfies our intuition that composite objects exist mind-independently.

Unlike Brutal Composition, reductive nihilism satisfies our intuition that composition is not brute.

Unlike almost all other views (the exceptions again being views that accept CAI), reductive nihilism satisfies our intuition that composite objects are nothing over and above their parts.

Unlike views that accept CAI, reductive nihilism satisfies our intuition that identity, governed as it is by Leibniz's Law, can never be many-one, and that cardinal properties are exclusive.

It appears that there are additionally no other intuitions that reductive nihilism cannot capture. For example, it agrees that things like tables and human beings are single things or objects, and it agrees with our intuitions concerning what sentences that purport to be about composite objects are really about.

To the extent that I have discerned all of the relevant intuitions concerning composition, then, reductive nihilism emerges as clearly the most intuitive view of composition. To the extent that intuition is a theoretical virtue, we have strong defeasible reason to endorse reductive nihilism.

Chapter 7

A Parsimonious Answer to the Special Composition Question I: In Defence of the Razor

7.1 - Introduction

Over the course of the next two chapters, I'm going to argue that reductive nihilism is the (joint) most parsimonious theory concerning composition. Since parsimony is almost always thought to be a theoretical virtue, a reductive nihilist victory in the parsimony stakes gives us another defeasible reason to be reductive nihilists.¹⁵⁵

The kind of parsimony that I'll be concerned with here is *ontological* parsimony, i.e. the kind of parsimony that (roughly speaking) concerns the minimisation of the amount of ontological posits a theory makes. For space considerations I'm *not* going to explicitly discuss ideological parsimony, i.e. the kind of parsimony that (roughly speaking) concerns minimising the amount of primitive predicates (and perhaps other bits of conceptual machinery) employed by a theory. This omission will, however, be at least partially offset by the assumption of realism about properties, which will have the effect of, to a large extent, subsuming ideological parsimony under ontological parsimony. For if bits of ideology are assumed to correspond to ontological posits, then a measurement of much of a theory's ideology will be included in the measurement of that theory's ontology. This will not become relevant until the next chapter, however, so for now we can move on.

My argument for the reductive nihilist's victory in the ontological parsimony stakes has two steps. The first step concerns how we should measure ontological parsimony in the first place. Only once we have a settled answer to this question can we turn to considering which answer to the SCQ is the most parsimonious. Thus in this chapter, I defend the claim that we should endorse *Ockham's Razor*, and the measure of parsimony it suggests, over its rival, *the Laser*. In the next chapter, I'll make the argument from Ockham's Razor for reductive nihilism, arguing that reductive nihilism is the joint most parsimonious theory concerning composition.

7.2 – *The Razor and the Laser*¹⁵⁶

Up until fairly recently, ontological parsimony had been widely understood as being a feature of theories that posit relatively few entities, or perhaps relatively few kinds of entities.¹⁵⁷ The idea that ontological parsimony is a theoretical virtue is reflected in the widespread idea that we should accept Ockham's Razor as a good methodological principle:

OCKHAM'S RAZOR: Don't multiply entities beyond necessity.

¹⁵⁵ I do not here attempt to motivate why we should think that ontological parsimony is a theoretical virtue in the first place. For arguments for this claim, see e.g. Tallant (2013) and Brenner (2017).

¹⁵⁶ The rest of this chapter is heavily based on my forthcoming article 'There is No Reason to Replace the Razor with the Laser'.

¹⁵⁷ In this chapter, the difference between *qualitative* and *quantitative* ontological parsimony (that is, parsimony understood as the minimisation of *kinds* of entities and entities *simpliciter* respectively) won't be very important, as arguments for adopting the Laser over the Razor seem to be equally strong on either conception. So in this chapter I'll often simply ignore the distinction (in the next chapter the distinction becomes salient again).

In recent times, however, it has become more and more common to find philosophers recommending replacing OCKHAM'S RAZOR with an alternative principle that Schaffer (2015) dubs *the Laser*:

THE LASER: (i) Don't multiply *fundamental* entities without necessity, and
(ii) Multiply non-fundamental entities all you like.¹⁵⁸

The Razor (as I'll call it) and the Laser are associated with differing conceptions of what ontological parsimony (*qua* theoretical virtue) consists in: parsimony with respect to *total* ontology and parsimony with respect to *fundamental* ontology, respectively. Thus the Razor reflects the thought that it is theoretically advantageous to reduce the number of entities in one's ontology as much as possible. The Laser instead reflects the thought that it is only theoretically advantageous to minimise the number of *fundamental* entities in one's ontology, and that there is no theoretical advantage to be had in minimising the number of non-fundamental entities one posits.¹⁵⁹

I assume here that the Razor, and the measure of ontological parsimony it recommends, is the default position. For I take it that, on the assumption that ontological parsimony of *some* sort is a theoretical virtue, the onus is on those who would impose restrictions on *which* kinds of entity should be minimised to convince us that their restriction is apt (cf. Da Vee 2020: §4). So in the absence of some compelling reason to replace the Razor with the Laser, we should accept the Razor.¹⁶⁰

In an attempt to discharge the burden of proof that they shoulder, supporters of the Laser have brought to bear an impressive battery of arguments for replacing the Razor with the Laser. Broadly speaking, they divide into three categories. *Arguments from the nature of non-fundamentality* attempt to motivate the Laser by appeal to various observations about what it is to be non-fundamental. *Arguments from cases* describe hypothetical or actual cases, and allege that only the Laser accords with our intuitive judgements about them. *Arguments from analogy* claim that ontological parsimony is analogous to conceptual economy, and that this analogy recommends the Laser.

My aim in this chapter is to argue that all of these arguments fail.¹⁶¹ In doing so, I'll demonstrate that we have no reason to replace the Razor with the Laser, and that we should therefore accept the Razor.

7.3 – Arguments from the Nature of Non-Fundamentality

Three supposed features of the nature of non-fundamentality have been cited in support of the Laser. Those features are: that the existence of non-fundamental entities is *necessitated* by that of the fundamental entities; that non-fundamental entities are in an important sense *ontologically innocent*;

¹⁵⁸ Schaffer's original formulation of the Laser omits (ii) (2015: 647). But (i) by itself is an implication of, rather than a rival to, the Razor (Baron and Tallant 2018: 597-8). Adding (ii) (which is equivalent to Baron and Tallant's 'Free Derivatives' principle [598]) fixes this, and is clearly what supporters of the Laser intended (see e.g. Bennett 2017: 220).

¹⁵⁹ Philosophers sympathetic to the Laser include: Schaffer (2007: 189; 2009: 361; 2010a: 40; 2010b: 313; 2012: 87-9; 2015), Bennett (2017: §8.2.2), Cameron (2010: 262-3), Korman (2015a: §4; 2015b: ch. 6 §3.2), and Sider (2013a: 240).

¹⁶⁰ I also assume here that the Razor and the Laser are the only two principles concerning ontological parsimony and theory choice. This is because the only other such principles that have been endorsed in the literature are minor variants on the Laser that differ from that principle only in that they better accommodate the possibility of there being no fundamental level (see Bennett's *Taser* [2017: 224n10]; see also Schaffer's *Phaser* [2015: 663]). My objections to arguments for the Laser should therefore generalise to arguments for these variants.

¹⁶¹ Others have attempted to defend the Razor against some (though not all) of the arguments for the Laser that I consider here (see Baron and Tallant 2018, Fiddaman and Rodriguez-Pereyra 2018, Da Vee 2020). I here provide novel responses to every argument for the Laser that has been proposed in the literature. I'll also note along the way why I take my responses to be preferable to extant ones.

and that non-fundamental entities play no role in fundamental explanations. I'll consider each such way of motivating the Laser in turn.

7.3.1 – *The Argument from Necessitation*

Bennett proposes an argument for the Laser based on the observation that the existence of the non-fundamental entities is typically claimed to be necessitated by the existence (and nature) of the fundamental entities. She writes:

Let T- and T+ be two theories that agree on all fundamental matters. According to T- that's all there is; according to T+, there are also a variety of nonfundamental matters. My claim is that T+'s extra ontological commitments do not tell against its simplicity in a way that makes it less likely to be true [...]. The key point is that according to T+, its extra ontological commitments are *necessitated* by the fundamental matters. Thus T+'s statements about the non-fundamental matters NF are – by its lights – entailed by statements about the fundamental matters F. And the following is a theorem of the probability calculus:

if $A \vdash B$, $\Pr(A) = \Pr(A \& B)$

It follows that according to T+, the probability of F is the same as the probability of F and NF. This means that – again according to T+ – *T+ is exactly as likely as T-*.

(2017: 223, emphasis original)

Bennett concludes from this that adding non-fundamental posits to a theory does not make that theory any less likely to be true (as long as the theory says that these posits are necessitated by the fundamental entities already posited) (2017: 225). On the innocuous assumption that theoretical virtues are those the possession of which makes a theory *ceteris paribus* more likely to be true, Bennett's argument would, if successful, establish that we should take the relevant theoretical virtue to be parsimony *qua* minimising only fundamental entities, rather than parsimony *qua* minimising entities *simpliciter*. This is in line with the Laser, not the Razor.¹⁶²

There must be something wrong with Bennett's argument. Consider your favourite theory of reality, and imagine that I have an identical theory except that mine additionally posits an empirically undetectable teapot hovering above the Eiffel Tower. Furthermore, according to my theory, this invisible teapot is a necessary existent. Since my theory says the teapot is a necessary existent, the ontology of your theory trivially necessitates the existence of the teapot. Thus, by Bennett's logic, your theory's ontological claims entail my theory's claim about the existence of the teapot. And it

¹⁶² It's worth noting a minor quibble immediately. Bennett's argument moves from the claim that (by T+'s lights) F *necessitates* NF, to the claim that (by T+'s lights) sentences about F *logically entail* sentences about NF, in order to use the principle of the probability calculus that she identifies to argue that (by T+'s lights) the probability of F and is equivalent to that of F alone. But presumably theories of the non-fundamental, including T+, think that the kind of necessitation holding between F and NF is merely *metaphysical* rather than *logical* necessitation – few theories regard it as *logically contradictory* to accept F but deny NF. So Bennett's argument really relies on a different principle to the one she identifies, namely:

(*) If A metaphysically necessitates B, then $\Pr(A) = \Pr(A \& B)$

The reason that this is only a minor quibble is that (*) seems just as plausible as the principle Bennett identifies. So for ease of presentation I'll stick with the principle as Bennett formulates it, on the understanding that strictly speaking (*) ought to be substituted in its place.

follows then from the probability calculus that (according to my theory) my theory is no less likely to be true than yours. This is an absurd conclusion. Additionally positing a necessarily existing empirically undetectable teapot with no good reason surely makes a theory much less likely to be true. Denying this is especially absurd when we consider that my theory *would* be much less likely to be true than yours if my theory didn't claim the invisible teapot to be a necessary being but rather a contingent one. Somehow, by Bennett's logic, claiming that the teapot exists in all possible worlds in which your theory's posits exist makes my theory much more likely than claiming that it exists in just some of those worlds. This cannot be right.

The problem with Bennett's argument is that it misconstrues what the probability calculus tells us about how entailment claims affect probability.¹⁶³ Consider again the theorem that Bennett cites:

$$\text{if } A \vdash B, \text{ then } \Pr(A) = \Pr(A \& B)$$

In order to use this theorem in support of the Razor, we'd have to read it as telling us to *disregard* the effect of an entailed claim on the probability of the theory that entails it. For only then could we conclude from T+'s supposition that the existence of NF is entailed by the existence of F that the inclusion of NF in T+'s ontology has no negative impact on the probability that that ontology is correct. But the above principle does *not* tell us to disregard the effect of an entailed claim on the probability of the theory that entails it. If it did, then we wouldn't be able to use the fact that a theory entails a contraction as a reason to assign that theory a probability of 0, which is absurd. Rather, what the above principle tells us is that, if A entails B, then when a theory accepts A it is *thereby incurring* any negative impact on its probability associated with accepting B. Thus, for example, if A entails B, and B is a contradiction, then accepting A reduces the probability of a theory to 0 (even if that theory doesn't explicitly accept B).

In our case, the above principle does *not* allow us to reason that, since T+ already claims that F exists and that F's existence entails NF's existence, any negative impact on the probability of T+'s ontology that would otherwise be associated with positing NF can be disregarded. Rather, it tells us only that, given T+'s claim that the existence of F entails the existence of NF, any negative (or indeed positive) impact on the probability of T+'s ontology associated with positing NF *was already incurred when T+ posited F*. This is entirely consistent with there *being (ceteris paribus)* a negative impact on the probability of T+'s ontology associated with positing NF, which is all that supporters of the Razor must claim here.

So, specifically, here is what supporters of the Razor should say about Bennett's case. Initially, our evidence suggests that we should assign a certain probability, say 0.7, to the claim that F exists. If we then discover, in line with T+, that F necessitates NF, what we have discovered is that positing F brings with it commitment to further entities than we previously thought. Since, in line with the Razor, we think that (*ceteris paribus*) extra ontological commitments lower the probability of a theory's truth, we should now in turn think that (*ceteris paribus*) the claim that F exists is less

¹⁶³ Da Vee instead objects to how Bennett allegedly assesses all probabilities by the lights of T+ (2020: 3684-5). But Bennett doesn't assess all probabilities by the lights of T+: her argument relies only on assuming T+'s claim that F necessitates NF when calculating the probability of T+'s postulation of NF. That's why Bennett explicitly claims that the probabilities of T+ and T- are equal, rather than it being the case that the probability of T+ is 1 whilst that of T- is 0. Bennett instead seems to be doing the following: granting T+'s assumption that F necessitates NF when calculating the probability of T+'s ontological claims, whilst granting T-'s assumption that F *doesn't* necessitate NF when assessing the probability of T-'s ontological claims. I lack the space to pursue the point, but suffice it to say that it's not clear to me that Bennett is doing anything illicit here. At any rate, my objection to Bennett is distinct from Da Vee's.

probable than we previously thought. On the other hand, if we decide, in line with T-, that F does *not* necessitate NF, then we have no reason to revise down the probability that we originally assigned to the claim that F exists. That is, given the Razor, the claim that F exists receives a different probability assignment depending on whether it is taken to require positing NF or not. This is perfectly consistent with the probability calculus, since it doesn't involve any violation of the claim that if F necessitates NF then $\Pr(F) = \Pr(F \& NF)$: what supporters of the Razor are saying here is that if F doesn't necessitate NF then $\Pr(F) = 0.7$, but if F does necessitate NF then $\Pr(F) = \Pr(F \& NF) = n$ for some $n < 0.7$. In this way supporters of the Razor can maintain their claim that T+'s commitment to NF has a negative impact on the probability of its truth in comparison to that of T-.

In general, it is consistent with the probability calculus to regard extra non-fundamental posits as (*ceteris paribus*) having a negative impact on a theory's probability, as the Razor implies. This is true even if non-fundamental posits are necessitated by the fundamental ones. Thus Bennett's argument from necessitation does not succeed.

7.3.2 – *The Argument from Ontological Innocence*

Bennett also thinks the alleged *ontological innocence* of non-fundamental entities provides support for the Laser (2017: 221-3). Her thought is this. Whilst it's not unusual to hear philosophers (such as Lewis 1991: 81) claiming non-fundamental entities to be ontologically innocent (or, equivalently, an *ontological free lunch* [Armstrong 1989: 56], or *nothing over and above the fundamental entities*, etc.), it's not obvious how we can make sense of this idea, especially given that non-fundamental entities are not generally taken to be identical to fundamental entities. But we apparently *can* explain this if we accept the Laser: claiming that non-fundamental posits don't count against a theory's parsimony in the way that matters for theory choice arguably captures and explains the sense in which those posits are ontologically innocent. This is a point in favour of the Laser.

But the argument from ontological innocence is unsuccessful, because there's a way of capturing the thought that non-fundamental entities are ontologically innocent that doesn't require accepting the Laser.¹⁶⁴ For, as Hawley notes (2014: §2), another way to secure the ontological innocence of the non-fundamental is to say that commitment to the fundamental entities automatically carries with it commitment to the non-fundamental entities. Non-fundamental entities would then be ontologically innocent in the sense that explicitly committing oneself to them doesn't add to one's ontological commitments at all (as commitment to them was already implicit in prior commitment to the fundamental entities). As I argued in the previous section, this is consistent with the thought that commitment to non-fundamental entities nonetheless still counts against the probability of a theory's truth, and is thus consistent with the Razor. So the ontological innocence of the non-fundamental doesn't give us any reason to abandon the Razor.

7.3.3 *The Argument from Explanation*

Korman argues that the relationship between (non-)fundamentality and *explanation* provides us with an argument for the Laser. He presents his argument as follows:

The most parsimonious theory is the one that explains all that needs to be explained using the fewest resources. Since fundamental objects are those in terms of which

¹⁶⁴ Da Vee instead objects to Bennett's argument here by rejecting that non-fundamental entities are ontologically innocent in the first place (2020: 3684). My response to Bennett has the advantage of remaining neutral as to whether non-fundamental entities are ontologically innocent or not.

everything is explained, it only makes sense to measure ontological parsimony in terms of which items are taken to be fundamental.

(2015a: 306; repeated in 2015b: 75-6)

Now, taken at face value, this argument is invalid. If the most parsimonious theory is the one that explains all that needs to be explained using the fewest resources, and if we can explain everything that needs to be explained by appealing only to fundamental objects, then what follows is *not* that our measure of parsimony should be blind to non-fundamental entities, but rather that the most parsimonious theory is the one that eliminates all non-fundamental objects (for on those assumptions, any theory that posits non-fundamental entities in addition to fundamental ones will explain nothing more but will posit more resources than a theory that posits only the fundamental entities). This conclusion is consistent with the Razor.

But immediately after the passage quoted above, Korman writes, ‘the mere fact that a theory’s fundamental mode of being is enjoyed by a wide range of objects is no strike against the parsimoniousness of that theory, since one need not suppose that those objects themselves all enter into *fundamental* explanations’ (2015a: 306; 2015b: 76; emphasis mine). With this in mind, I think we should take Korman to be arguing as follows:

- (1) The most parsimonious theory is the one whose *fundamental* explanations explain all that needs to be explained using the fewest resources.
- (2) Fundamental objects are those in terms of which everything is *fundamentally* explained.
- (3) So we should measure ontological parsimony in terms of which items are taken to be fundamental.

(1)-(3) is a valid argument for the Laser. But are (1) and (2) true?

To answer, we need to know what a fundamental explanation is. Unfortunately, Korman doesn’t provide an explicit definition of this notion. But a clarification of what he means by it might be implicit in what he says about what it is for an object (or an ‘item’) to be fundamental:

‘The basic idea is that an item’s fundamentality should be a function of the way in which it features in metaphysical explanations. [...] We] could say that (i) A is fundamental simpliciter iff it features in facts that do not obtain in virtue of any other facts, and (ii) A is more fundamental than B if some B-involving facts obtain partly or wholly in virtue of A-involving facts and never vice versa.’

(2015a: 305; repeated in 2015b: 73)

Korman is here talking about *metaphysical* explanations. These are *acausal* explanations, and are usually invoked by the ‘in virtue of’ locution (see e.g. Schaffer 2009, 2016: 52-4, 2017: 303; Correia and Schneider 2012: 1). With that in mind, Korman’s remarks might be taken to be suggestive of the following definition of the notion of a fundamental (metaphysical) explanation:

FUNDAMENTAL EXPLANATION: p is a fundamental explanation =_{df} p is a fact in virtue of which at least one further fact obtains but which does not itself obtain in virtue of other facts.

The problem is, if (1) and (2) are read as employing the notion of a fundamental explanation at issue in FUNDAMENTAL EXPLANATION then Korman's argument becomes question-begging against supporters of the Razor.¹⁶⁵ For, as is clear from the above quotation, Korman takes it to be definitive of what it is for an item to be fundamental that that item features in facts that do not obtain in virtue of any other facts. But given FUNDAMENTAL EXPLANATION, every fundamental explanation is a fact that does not obtain in virtue of any other facts. It follows that, *by definition*, fundamental explanations only make reference to fundamental entities. So (1) in the argument above is equivalent to:

(1*) The most parsimonious theory is the one whose explanations that only make reference to fundamental entities explain all that needs to be explained using the fewest resources.

'Resources' as it appears in (1*) can only mean *fundamental resources*; only those antecedently sympathetic to the Laser will agree that parsimony is a matter of reducing the numbers of only *those* resources, as opposed to *all* resources, including non-fundamental ones.

Perhaps Korman could attempt to make his argument non-question-begging by continuing to accept FUNDAMENTAL EXPLANATION but divorcing the definition of a fundamental entity from that of a fundamental explanation.¹⁶⁶ That would prevent it from being analytic that a fundamental explanation is one that only features fundamental entities, and so would block the analytic equivalence between (1) and the question-begging (1*).

But it's still not clear this would prevent the argument from begging the question against those who don't antecedently accept the Laser; at the very least, (1) still seems unmotivated without a background commitment to the Laser. For what reasons have we been given for thinking that parsimony only concerns minimising the amount of resources required for *fundamental* explanations, rather than all explanations? Those attracted to the Razor will presumably be inclined to think that considerations of parsimony pressure us to instead *eliminate* non-fundamental explanations and the non-fundamental entities they make reference to. Nothing Korman has said motivates thinking otherwise. I conclude that his argument from explanation for the Laser does not succeed.

7.4 – Arguments from Cases

I turn now to two attempts to motivate the Laser by appealing to hypothetical or real-world cases.

7.4.1 – The Argument from the Case of Esther and Feng

Schaffer asks us to consider the following case. Suppose that Esther formulates a scientific theory according to which there are 100 types of fundamental particle, and that her theory is widely accepted. Then:

'Feng comes along and – in a moment of genius – builds on Esther's work to discover a deeper fundamental theory with 10 types of fundamental string, which in varying combinations make up Esther's 100 types of particle. This is intended to be a

¹⁶⁵ Da Vee shares my concern that Korman's argument is question-begging (though he doesn't spell out why) – see his (2020: 3690-1).

¹⁶⁶ For example, perhaps Korman could regard what it is to be a fundamental entity as *primitive*, à la Wilson (2014: 560), or alternatively take fundamental entities to be those that are not grounded by anything else, à la Schaffer (2009: 373, 2010a: 38) (though to avoid begging the question he would then need an account of grounding that didn't render it analytic that ungrounded entities are those that feature in fundamental explanations).

paradigm case of scientific progress in which a deeper, more unified, and more elegant theory ought to replace a shallower, less unified, and less elegant theory. Feng's theory is evidently better in every relevant methodological respect.'

(2015: 648)

This, says Schaffer, tells in favour of the Laser. For given the Razor, there is apparently at least *one* respect in which Feng's theory is not better than Esther's: since Feng posits everything Esther posits *plus* the additional strings, Feng's theory would apparently be disfavoured by the Razor, all else being equal. The Laser, on the other hand, correctly favours Feng's theory for its smaller fundamental ontology. As Schaffer has it:

'So, by the lights of *The Razor*, Feng's theory is an affront to ontological economy for positing these additional strings. It is to be strongly dispreferred, all else equal. This is obviously backwards, as far as sound methodological counsel is concerned.

Feng's theory is obviously no affront to ontological economy, but – when judged purely by the methodological virtues – is evidently a more economical, tighter, and more unified improvement. It is *The Laser* that gets this right.'

(648)

In response, others have pointed out that it's consistent with the Razor to think that Feng's theory is superior to Esther's *all things considered*, because Feng's theory secures other, weightier theoretical virtues such as explanatory power and theoretical unification (Baron and Tallant 2018: 599; Fiddaman and Rodriguez-Pereyra 2018: 343; Da Vee 2020: 3681).¹⁶⁷ This is all well and good, but it leaves the door open to Schaffer to rejoin by insisting again that his intuition in this case is that Feng's theory is 'obviously no affront to ontological parsimony', not merely that Feng's theory is *all things considered* superior to Esther's. Pointing to the greater explanatory power (etc.) of Feng's theory over Esther's does nothing to show how that Razor is consistent with *this* intuition. Absent some reason for thinking that Schaffer's intuition can safely be discounted here, it's not clear that this line of response to Schaffer's argument for the Laser from the case of Esther and Feng is successful.

Instead, supporters of the Razor should offer a way of explaining away the intuition that Feng's theory is 'no affront to ontological parsimony' that doesn't require accepting that ontological parsimony is blind to non-fundamental ontology. They should say that that intuition is instead generated by the

¹⁶⁷ Fiddamann and Rodriguez-Pereyra also advance a different objection to Schaffer here, which is that his argument based on the case of Esther and Feng is simply 'fallacious' (2018: 342). They concede that consideration of the case of Esther and Feng demonstrates that positing fewer fundamental entities is a good-making feature of a theory, but note that both the Laser and the Razor agree on this (342-3); they then complain that an argument for the Laser here would have to show that multiplying non-fundamental entities is not a bad-making feature of a theory, but that the case of Esther and Feng doesn't show this (342-3). But it seems to me that Fiddaman and Rodriguez-Pereyra fail to appreciate the structure of Schaffer's argument, which is as follows:

- (1) Esther's theory posits 100 types of entity, whilst Feng's posits 110 types of entity.
- (2) Feng's theory is better than Esther's in every relevant methodological respect.
- (3) So, from (2), Feng's theory is better than Esther's with respect to ontological parsimony.
- (4) So, from (1) and (3), the Razor is the wrong measure of ontological parsimony.

With the auxiliary assumption (that Fiddaman and Rodriguez-Pereyra do not seem to dispute) that the Laser is the only alternative to the Razor, it follows from (4) that the Laser is the right measure of ontological parsimony.

perception, created by Schaffer's description of the case, that Feng's theory is *true*, or at least correct in positing the extra layer of fundamental entities below the entities that Esther posits. Schaffer says a number of things to encourage this perception: he describes Feng's formulation of his theory as a 'moment of genius', and the theory itself as a 'deeper fundamental theory' in comparison to Esther's; most tellingly, he says that Feng's building upon Esther's theory to propose his own 'is intended to be a paradigm case of scientific progress' (2015: 648). To be clear, I don't say that Schaffer's description of the case of Esther and Feng is logically inconsistent with the idea that Esther's theory might be true after all (so e.g. it is probably consistent – though a bit odd – to describe Feng's replacement of Esther's theory with his own as a 'moment of genius' even though Esther's theory is true): my point is just that it's easy and natural to read what Schaffer says and assume that in the universe he describes Feng is right to say that there is an extra layer of entities below the ones that Esther posited. And if we do read Schaffer in this way, then we have a way of explaining the intuition that Feng's theory is no affront to ontological parsimony that doesn't require swapping the Razor for the Laser. For the Razor only tells us to avoid positing *unnecessary* entities. But if we assume that the extra entities that Feng posits exist, then we must also accept that it was necessary to posit them. Thus Feng's theory doesn't offend against the Razor at all, since it doesn't multiply entities *beyond necessity*. So on the assumption that, in the universe that Schaffer describes, Feng's extra entities really do exist, we can hold on to the Razor whilst agreeing that Feng's theory is no affront to ontological parsimony.

Schaffer might reply that the implication that Feng's theory is true is a red herring, and that the Razor would get the wrong result in an analogous case to that of Esther and Feng that doesn't smuggle in or encourage the assumption that Feng's theory is true. But it seems to me that the presumed truth of Feng's theory is *precisely* what drives the intuition that Feng's theory is no affront to ontological parsimony, at least to those not independently drawn to the Laser. Strip away all the admiring language with which Schaffer describes Feng's theory, and what we are left with is a case in which we have two theories, T1 and T2, where T1 posits 100 fundamental entities, and T2 posits 10 fundamental entities and 100 non-fundamental entities. Nothing about this case seems to add anything to dialectic: those who find the Laser independently plausible will judge that T2 is ontologically simpler in the way that matters, those who find the Razor independently plausible will judge that T1 is ontologically simpler in the way that matters, and undecided parties will continue to be undecided, having been given no reason to make up their minds one way or the other. What was apparently so compelling about Schaffer's original argument was that *even those not antecedently sympathetic to the Laser* (including those who find the Razor independently plausible) were likely to have the intuition that Feng's theory is no affront to ontological parsimony, and to the extent that this intuition was suggestive of the Laser, everyone therefore had a reason to be attracted to the Laser. But once we strip away the confounding assumption that Feng's theory is correct in positing its extra layer of fundamental entities below those of Esther's, those not antecedently sympathetic to the Laser are no likelier than they were before considering the case of Esther and Feng to have the pro-Laser intuition.

So, in sum: either the case of Esther and Feng is described in such a way as to encourage the assumption that Feng's extra entities really do exist, in which case the Razor agrees with Schaffer's intuition that Feng's extra entities are no affront to ontological parsimony, or it is described without this implication, in which case only those who already find the Laser plausible on independent grounds will think that the Razor is wrong to regard Feng's extra entities as genuine (though outweigh-able) costs of his theory. Either way, the argument from the case of Esther and Feng doesn't give us any reason to abandon the Razor for the Laser.

7.4.2 – *The Argument from a Bias Towards the Built*

Bennett claims that a commitment to the Laser is latent in actual scientific practice (2017: 220-1).¹⁶⁸ In particular, she argues that we are implicitly attracted to a scientific methodology that includes a *bias towards the built*. For Bennett, to be built is to be non-fundamental, and so a bias towards the built is the bias in favour of claiming entities in our ontology to be non-fundamental, rather than fundamental. Such a bias would suggest an implicit preference on our part for the Laser over the Razor, as only the Laser is consistent with the thought that there is theoretical advantage to be had in shifting ontological commitments away from our fundamental ontology – thereby minimising the amount of fundamental entities we posit – and into our non-fundamental ontology.

Is it plausible that a bias towards the built is implicit in scientific methodology? Bennett provides the following motivation for thinking so:

‘[...] We all think that things ought to be explained wherever possible. We don’t rest content believing in water; we want to know what water is made of, and how exactly those components come together to behave as water does. This is what drives science: we want to account for some things in terms of other things. All else equal, we prefer things to be built. Indeed, we prefer things to be built from components to which we are already committed. When scientists are faced with some interesting new phenomenon, they first try to explain it in terms of things they already believe in. Of course, they may eventually have to posit some new fundamental entity or force to explain it, or may even have to accept the phenomenon as itself fundamental. But that is a last resort [...] All this is to say that we have a *bias towards the built*.’

(2017: 221)

To evaluate this, we need to know what ‘phenomenon’ means. On one plausible interpretation, ‘phenomenon’ means something like *empirical datum*. On this reading it is plausible that scientists do indeed try to explain new phenomena ‘in terms of things they already believe in’, but this fact is suggestive only of a bias on the part of scientists against unnecessarily positing new entities, fundamental or not, rather than of a bias towards the built. This is consistent with the Razor. So Bennett must be using ‘phenomenon’ to mean something like *object or process or event*, such that accepting the phenomenon as genuine or real (as opposed to illusory) is equivalent to accepting ontological commitment to that phenomenon. Then a bias towards explaining new phenomena in terms of entities that are already in our ontology would amount to a bias towards the built (as long as the sort of explanation at issue is metaphysical explanation, since saying that a phenomenon is *built* or *non-fundamental* is to say that it is metaphysically explained by the entities of which it is derivative¹⁶⁹).

¹⁶⁸ In the same discussion, Bennett also describes the following hypothetical case. Suppose that T1 and T2 agree on what entities there are, but that whereas T1 regards all of those entities as fundamental, T2 claims that some are derivative of the others. ‘Surely’, says Bennett, ‘the second theory is simpler in the way that matters’ (2017: 221). But if this is to be treated as an argument for the Laser in its own right, then it is unconvincing. Suppose T1 and T2 both posit 100 quarks, that T1 says that all the quarks are fundamental, but that T2 says that 99 of the quarks are derivative of the remaining one. At the very least, it’s not obvious that T2 is the superior theory here; indeed, my intuition is that T1 is clearly preferable to T2, and that there’s no pressure from parsimony to say otherwise.

¹⁶⁹ See Bennett (2017: §3.4).

But on this reading of ‘phenomenon’, we don’t have any reason to think that scientists have a generalised preference for saying that new phenomena are (metaphysically) explained by other entities already in their ontology. First, whilst it’s true that scientists in general ‘don’t rest content’ simply believing in a given phenomenon, and rather try to find out what it is made of, this is evidence only of the fact that scientists want to *find out* whether a given phenomenon is built out of smaller components, rather than that they are *antecedently inclined to believe* that the phenomenon is built out of smaller components (compare: a scientist who attempts to discover whether some particle is charged or not need not have a ‘bias towards the charged’). Further, in cases in which it does seem plausible that scientists would intuitively prefer to regard the new phenomenon as built out of smaller components, we can explain this preference without appeal to a general bias towards the built. Whilst a scientist who (for example) discovers water for the first time may well prefer a theory that predicts that water is made up of smaller components, this seems likely to be because experience has taught her that macroscopic phenomena have in the past always turned out to be made from smaller particles (perhaps accompanied by the intuition that macroscopic extended simples are inherently dubious). Absent a case in which such considerations *don’t* plausibly explain a scientist’s preference for regarding the new phenomenon in question as probably built out of smaller components, we don’t have a good reason for agreeing with Bennett that scientists (or we) have an implicit bias towards the built.

So Bennett’s appeal to actual scientific practice doesn’t demonstrate that we have a bias towards the built, and so her argument from a bias towards the built for the Laser doesn’t succeed.

7.5 – Arguments from Analogy

Schaffer (2015: §§4-5) offers two arguments for the Laser based on an alleged analogy between ontological parsimony and *conceptual economy*, which concerns the minimisation of concepts invoked by a theory

7.5.1 – The Argument from the Conceptual Laser

Schaffer’s first argument from analogy with conceptual economy is the most direct. He begins (2015: 649) by asking us to consider two candidate principles concerning conceptual economy:

THE CONCEPTUAL RAZOR: Do not invoke concepts without necessity.

THE CONCEPTUAL LASER: (i) Do not invoke primitive concepts without necessity.
(ii) Multiply non-primitive concepts all you like.¹⁷⁰

Each principle is suggestive of a different way of measuring conceptual economy. The Conceptual Razor suggests that conceptual economy consists in minimising the total number of concepts; the Conceptual Laser suggests that conceptual economy consists in minimising the total number of *primitive* concepts, and that multiplying the number of defined (aka *derivative*) concepts invoked doesn’t count against a theory’s conceptual economy at all.

Schaffer thinks that it’s defeasibly reasonable to suppose that conceptual economy and ontological parsimony are analogous, and thus that ‘it is defeasibly reasonable to expect that the apt measures of economy will be parallel’ between conceptual economy and ontological parsimony (649). Since primitive concepts appear analogous to fundamental entities, and defined concepts appear analogous

¹⁷⁰ Though Schaffer’s formulation of the Conceptual Laser omits (ii), he clearly intends the Conceptual Laser to be committed to it.

to non-fundamental/derivative entities, the Conceptual Razor appears to be directly analogous to the (i.e. Ockham's) Razor whilst the Conceptual Laser appears to be directly analogous to the Laser. Further, Schaffer thinks that the Conceptual Laser is the right measure of conceptual economy, arguing that only the Conceptual Laser is consistent with our intuitions in various cases he cites (649-51). By analogy, then, he concludes that the Laser is the right measure of ontological parsimony.

Now, to be clear, if conceptual economy is to be analogous to ontological parsimony in a way that might support the inference from the Conceptual Laser to the Laser, then we must understand conceptual economy as being *truth conducive* (as opposed to e.g. merely making for more aesthetically pleasing or pragmatically useful theories). The Razor and the Laser are supposed to imply that theories that do not multiply entities (or fundamental entities) beyond necessity are, all else equal, more likely to be true than more ontologically profligate theories. If the Conceptual Razor and the Conceptual Laser are to be analogous to these principles, then we must understand them as implying that theories that do not multiply *concepts* (or primitive concepts) beyond necessity are, all else equal, more likely to be true than more conceptually profligate theories.

With this in mind, I think Schaffer is right to say that the Conceptual Laser is the correct measure of conceptual economy, but wrong to conclude from this that the Laser is by analogy the correct measure of ontological parsimony.¹⁷¹ Indeed, I think that examination of the *reason* why the Conceptual Laser is the correct measure of conceptual economy reveals a crucial disanalogy between conceptual economy and ontological parsimony, which suffices to block the conclusion that the Laser is the correct measure of ontological parsimony.¹⁷²

First, then, the reason we should endorse the Conceptual Laser has to do with the way in which defined concepts are *eliminable* from the theories that employ them, in the following sense: for any theory T_d that employs defined concepts d_1, \dots, d_n that it defines in terms of primitive concepts p_1, \dots, p_n , there is a theory T_p that is *equivalent* to T_d but that replaces each d_i of T_d with its definiens, and that thus employs only p_1, \dots, p_n . This much follows from the nature of a definition, which guarantees that we can always preserve meaning (and certainly truth) by replacing an instance of a definiendum with an instance of its definiens. For example, suppose that T_1 says that objects a and b *overlap*, and defines 'overlap' in terms of the primitive concept of *parthood*, in the usual way: then there is an equivalent theory, T_2 , that makes no mention of overlap, and that instead says only that a and b have a part in common.

Now assume for *reductio* that defined concepts count against a theory's conceptual economy. Then T_1 is less conceptually economical than T_2 , because T_1 employs both *parthood* and *overlap*, whilst T_2 employs only *parthood*. So T_1 is less likely to be true than T_2 . But T_1 and T_2 are equivalent. So T_1 and T_2 must be equally likely to be true. Contradiction. So defined concepts do not count against a

¹⁷¹ Thus I disagree with Baron and Tallant (2018: 600-1), who object to Schaffer's arguments for the Conceptual Laser and deny that this principle is the right measure of conceptual economy in the first place. At any rate, whatever the merits of their objections to Schaffer's arguments for the Conceptual Laser, they don't undermine the argument that I provide below in support of the Conceptual Laser.

¹⁷² Fiddaman and Rodriguez-Pereyra adopt a similar line of argument (2018: 346) but fill in the details wrong. They claim that the motivation for the Conceptual Laser is that theories with fewer primitive concepts are more *intelligible* to an ideally intelligent person (whilst the number of derivative concepts a theory contains doesn't affect this sort of intelligibility at all), and that this motivation doesn't carry across to the case of ontological parsimony. But on the assumption that conceptual economy is truth conducive, it's dubious that considerations of intelligibility can motivate any particular measure of conceptual economy, because it's dubious that theories that are more intelligible are more likely to be true. Suppose otherwise. Then, since equivalent theories can apparently differ with regard to intelligibility, equivalent theories can differ with regard to the probability that they are true, which is impossible.

theory's conceptual economy. That means that the Conceptual Laser is the right measure of conceptual economy (cf. Cowling 2013: 3893).

In general, since we can always reformulate a theory that employs defined concepts in addition to primitive concepts so that that it only employs its primitive concepts, *without changing the meaning or probability of truth of that theory*, it makes sense to ignore the defined concepts when measuring conceptual economy.

The reason that conceptual economy and ontological parsimony are crucially *disanalogous*, then, is that that same motivation doesn't carry across to the ontological case, because non-fundamental ontological commitments are not eliminable in the requisite sense.¹⁷³ That is, we cannot 'reformulate' a theory that contains ontological commitment to both fundamental and non-fundamental entities in such a way as to remove all ontological commitment to non-fundamental entities without thereby changing the meaning of the theory. For example, suppose we start with a theory that posits both fundamental simples and some non-fundamental mereological fusions of those simples; then suppose we strip away from that theory all ontological commitment to the non-fundamental mereological fusions, reformulating it so that any sentence implying the existence of a composite *F* is replaced with a sentence instead implying the existence only of simples arranged *F*-wise. The result of this would be a theory *distinct from* (i.e. non-equivalent to) the one we started with, for the theory we'd end up with would be consistent with mereological nihilism, whilst the one we started with was not. The non-fundamental fusions that our original theory posited, then, were not eliminable in the requisite sense, and so we cannot infer that they shouldn't count against a theory's ontological parsimony.

To sum up, then: the very consideration that warrants ignoring defined concepts when calculating a theory's conceptual economy, namely the eliminability of defined concepts, does not apply in the case ontological parsimony, for non-fundamental entities are not eliminable in the requisite sense. In this way, ontological parsimony is importantly *disanalogous* to conceptual economy. As such, there are no grounds for inferring from the fact that the Conceptual Laser is the right measure of conceptual economy that the Laser is the right measure of ontological parsimony.

7.5.2 – *The Argument from Bang for the Buck*

Schaffer (2015: 651-3) proposes a further argument for the Laser from analogy with conceptual economy. His argument takes as its starting point the thought that ontological parsimony is relevant to theory choice only insofar as it is relevant to the question of how much *bang for the buck* a theory secures. The best theories are those that, *ceteris paribus*, find the best balance between minimising their buck whilst maximising their bang. Schaffer thinks that the right formulation of the principle of 'bang for the buck' that concerns ontology is:

ONTOLOGICAL BANG FOR THE BUCK: Optimally balance minimisation of
fundamental entities with maximisation of derivative entities (especially useful ones).

ONTOLOGICAL BANG FOR THE BUCK plainly recommends the Laser over the Razor, as only the Laser is consistent with the injunction to minimise fundamental posits whilst maximising non-fundamental ones. Given this, it's clear that no-one who isn't antecedently attracted to the Laser will find ONTOLOGICAL BANG FOR THE BUCK independently plausible. But Schaffer thinks that this

¹⁷³ See also Da Vee (2020: 3682-3) for a different way of arguing that conceptual economy and ontological parsimony are *disanalogous*.

principle is well-motivated by the fact that conceptual economy and ontological parsimony are analogous, and the following principle concerning conceptual economy is correct:

CONCEPTUAL BANG FOR THE BUCK: Optimally balance minimisation of primitive concepts with the maximisation of defined concepts (especially useful ones).

This time we need not dispute that conceptual economy and ontological parsimony are analogous in the way that Schaffer needs them to be to support his argument. For CONCEPTUAL BANG FOR THE BUCK cannot be correct in the first place: it cannot be that, *ceteris paribus*, theories are more likely to be true if they optimally balance minimisation of primitive concepts with the maximisation of defined concepts (especially useful ones). The reason for this has to do again with the notion that defined concepts are eliminable from the theories that employ them. Consider T3, a theory that uses a mixture of primitive concepts and defined concepts, and T4, which is the result of reformulating T3 so as to replace its defined concepts with combinations of its primitive ones. As I argued in the previous section, the nature of definition means that T3 and T4 are equivalent. But CONCEPTUAL BANG FOR THE BUCK says that T3 is superior to – i.e. more likely to be true than – T4. This is a contradiction, so CONCEPTUAL BANG FOR THE BUCK must be false. Thus Schaffer’s argument from this principle for ONTOLOGICAL BANG FOR THE BUCK, and therefore for the Laser, fails.¹⁷⁴

7.6 – Concluding Remarks

I’ve examined all seven arguments for replacing the Razor with the Laser that appear in the literature: three that appeal to various facets of the nature of non-fundamentality, two that appeal to hypothetical or real-world cases, and two that appeal to an alleged analogy with conceptual economy. I’ve argued that none of these arguments for the Laser succeed.

In closing, it is worth considering a final way of trying to justify accepting the Laser over the Razor, one that doesn’t appeal to any argument that has the Laser as its conclusion: perhaps those who support the Laser can simply maintain that they find the idea that non-fundamental ontological posits do not count against the ontological simplicity of a theory directly intuitively plausible. Indeed, the very fact that there are so many philosophers offering arguments for the Laser may be taken to be suggestive of the prevalence of this intuition. It would help to shore up the defence of the Razor presented here if something could be said in response to this way of justifying the Laser.

To that end, I think supporters of the Razor can first legitimately raise doubts about the idea that the Laser is really as directly intuitive as it perhaps appears to be. Isn’t it possible that what supporters of the Laser *really* find intuitive is that non-fundamental entities are ontologically innocent, or that it is theoretically virtuous to explain things wherever possible, etc., and that they claim to find the Laser directly intuitive simply because they are (perhaps implicitly) convinced of the arguments from these sorts of claims to the Laser? If this is right, then what I’ve done here in showing that the Razor is consistent with the ontological innocence of non-fundamental entities (§7.3.2), and with the thought that it is theoretically virtuous to explain things wherever possible (§7.4.2), and more generally in finding fault with arguments for the Laser from apparently intuitively true premises, is to show that there is no intuitive grounding for the Laser after all.

¹⁷⁴ Da Vee (2020: 3683) also rejects *Conceptual Bang for the Buck*, but does so because he thinks that *useless* defined concepts shouldn’t be considered as part of the conceptual ‘bang’ of a theory. But the usefulness or otherwise of defined concepts is irrelevant, given my argument: defined concepts of any sort, whether useful or not, cannot be the conceptual bang of a theory.

That said, supporters of the Razor ultimately need not rely on the claim that no-one really finds the Laser directly intuitive (and so they need not worry about recalcitrant supporters of the Laser who insist that they find it to be an intuitive principle independently of considerations of ontological innocence, etc.). For they can retreat to the defence that, whilst some philosophers may well find the Laser to be directly intuitive, there's evidence that *other* philosophers think that it's the *Razor* that is the intuitively correct principle. After all, if the existence of arguments for the Laser counts as prima facie evidence that the authors of those arguments find the Laser to be directly intuitive, then the existence of defences of the Razor against those arguments (including that of this chapter), as well as of positive arguments for accepting the Razor over the Laser (see Baron and Tallant 2018: 603ff. Da Vee 2020: §3), should similarly count as prima facie evidence that the authors of *those* arguments find the Razor to be directly intuitive. We seem to have no reason, then, for thinking that there's any more intuitive support for the Laser than there is for the Razor.

So the extant arguments for the Laser fail, and consideration of the alleged bare intuitiveness of the Laser doesn't seem to tip the scales in its favour either. I conclude that we currently have no reason to replace the Razor with the Laser. Since we started out by assuming that the Razor (and the measure of ontological parsimony that it recommends) is the default principle (see §7.2), that means that we should accept the Razor.

Chapter 8

A Parsimonious Answer to the Special Composition Question II: Keeping it Simple

8.1 - Introduction

Having defended against the pressure to reject Ockham's Razor in favour of the Laser, I turn now in this chapter to presenting the argument from the Razor for reductive nihilism. In particular, I'll argue that, given that ontological parsimony *qua* theoretical virtue is measured by the total size of a theory's ontology, reductive nihilism is the joint most ontologically parsimonious theory concerning composition.

My argument takes as a starting point the observation that reductive nihilism is just as ontologically parsimonious as eliminative nihilism. After all, as I noted back in chapter 2, reductive nihilism and eliminative nihilism agree on the metaphysical picture of the world, in that they both accept MICROPHYSICAL NIHILISM. Admittedly, reductive nihilism has the *appearance* of being more ontologically profligate than eliminative nihilism, because the former accepts the existence of (e.g.) the Eiffel Tower as well as the simples* arranged Eiffel-Tower-wise, whereas the latter denies the existence of (e.g.) the Eiffel Tower and only accepts the existence of the simples* arranged Eiffel-Tower-wise. But this apparent ontological difference between reductive and eliminative doesn't stand up to scrutiny. Though the reductive nihilist accepts the existence of the Eiffel Tower as well as that of the simples* arranged Eiffel-Tower-wise, that's because she claims that 'the Eiffel Tower' refers to the simples* arranged Eiffel-Tower-wise. Thus the reductive nihilist's acceptance of the existence of the Eiffel Tower commits her only to the existence of the simples* arranged Eiffel-Tower-wise, which the eliminative nihilist also posits. In general: wherever the reductive nihilist appears to be committed to an 'extra' ontological posit in comparison to the eliminative nihilist, the reductive nihilist's semantics imply that her apparently 'extra' ontological posit was really something the eliminative nihilist was committed to all along. So, plausibly, the ontologies of reductive nihilism and eliminative nihilism are identical. The differences between the views lie solely in the semantic claims that they make.¹⁷⁵

I think that the equivalence between the ontologies of reductive nihilism and eliminative nihilism, together with the fact that the reductive nihilist's non-standard semantics give the appearance that its ontology is larger than it really is, is suggestive of a sensible way to proceed: drop the reductive nihilist's non-standard semantics in place of the eliminative nihilist's when comparing reductive nihilism's ontological parsimony with that of other theories. Equivalently: a sensible way forward is to argue that *eliminative* nihilism is the joint most ontologically parsimonious answer to the SCQ, and conclude from the success of this argument that reductive nihilism, which after all shares an ontology with eliminative nihilism, is also the joint most ontologically parsimonious answer to the SCQ. So that is what I will do.

¹⁷⁵ A similar line of argument suggests that reductive nihilism and eliminative nihilism are equally ideologically parsimonious as well. For example, though the reductive nihilist posits mereological concepts (such as *parthood*), she defines these in terms of non-mereological concepts that the eliminative nihilist is already committed to (such as *amonghood*, etc.) – see §2.4.2. Thus although the reductive nihilist accepts that certain objects stand in the *parthood* relation to one another, her semantics entail that in saying this she is committing herself only to the claim that certain pluralities stand in the *amonghood* relation to one another, which the eliminative nihilist also accepts.

My argument for eliminative nihilism's (joint) victory in the ontological parsimony stakes comes in two parts. In the first part (§8.2) I'll argue that eliminative nihilism is more ontologically parsimonious than all non-nihilist answers to the SCQ. In the second part (§8.3) I'll compare eliminative nihilism to two other versions of nihilism: emergence nihilism, and existence monism. Whilst it's obvious, I think, that eliminative nihilism is more ontologically parsimonious than emergence nihilism, it's dubious that eliminative nihilism is more ontologically parsimonious than existence monism, and indeed it has been argued that existence monism is more ontologically parsimonious than eliminative nihilism. I'll argue that there's a good case for thinking that eliminative nihilism and existence monism are rather equally ontologically parsimonious.

A final note before beginning in earnest. Throughout this chapter I'll ignore CAI. CAI, as we've seen, identifies composite entities with their parts. Plausibly, if an *F* is identical to some thing or things that are in your ontology, then your ontology doesn't become any less ontologically parsimonious if you explicitly accept the existence of the *F*. Thus suppose that the proponent of CAI starts out with the microphysical nihilist's ontology, i.e. an ontology of simple* quarks and leptons. Although she then additionally commits herself to the existence of composite *F*s, for each composite *F* she posits, she identifies that *F* with some simples* that were already in her ontology.¹⁷⁶ Thus, plausibly, the ontology of CAI becomes no less ontologically parsimonious for positing composite *F*s in addition to the simples*. Thus, plausibly, CAI is equally as ontologically parsimonious as eliminative nihilism. Or, at least, I won't dispute this here. So for most of the remainder of this chapter, I'll ignore CAI, mentioning it again only at the end of this chapter as a reminder that eliminative nihilism is only the *joint* most (rather than *the* most) ontologically parsimonious view of composition.

With that in mind, I turn now to arguing that eliminative nihilism is more ontologically parsimonious than non-nihilist answers to the SCQ.

8.2 – Nihilism versus Non-Nihilist Answers to the SCQ¹⁷⁷

I'll argue for the claim that eliminative nihilism prevails over non-nihilist answers to the SCQ in the ontological parsimony stakes by first establishing the presumptive case in favour of this claim, and then vindicating the presumptive case by responding to an important challenge to it.

Ontological parsimony comes in two varieties: *quantitative* (ontological) parsimony and *qualitative* (ontological) parsimony. We've already come across these notions informally, but it will be helpful to explicitly define them here:

QUANTITATIVE PARSIMONY: Theory A is more quantitatively parsimonious than theory B =_{df} A posits fewer entities than B.

QUALITATIVE PARSIMONY: Theory A is more qualitatively parsimonious than theory B =_{df} A posits fewer kinds of entities than B

Now, eliminative nihilism's ontology contains only simples*, and properties of the simples*, whilst the ontologies of rival views contain not only these things but also at least some composite entities

¹⁷⁶ Strictly speaking, CAI is not committed to any answer to the SCQ, so need not posit any composite entities at all. What I mean here is that even if CAI *does* posit composite entities, it remains no less ontologically parsimonious than eliminative nihilism for doing so.

¹⁷⁷ Much of §8.2 is based on my (2017) article 'Mereological Nihilism: Keeping it Simple'.

and properties of those composite entities.¹⁷⁸ Thus eliminative nihilism appears to be more *quantitatively* parsimonious than its rivals, in that it seems to posit fewer entities (tout court), and more *qualitatively* parsimonious because it seems to posit fewer distinct kinds of entity (it doesn't posit composite entities). Since quantitative and qualitative parsimony are the two kinds of ontological parsimony, we have a presumptive case in favour of the claim that eliminative nihilism is more ontologically parsimonious than almost all of its rivals.

This presumptive case has not gone unchallenged, however. Karen Bennett (2009) argues that, on closer inspection, eliminative nihilism is no more quantitatively parsimonious than other answers to the SCQ. What's more, her reasoning can be redirected to attack the claim that eliminative nihilism is more qualitatively parsimonious than its rivals. So I turn now to rebutting but Bennett's argument and the suggested redirection thereof, in order to vindicate the presumptive case in favour of eliminative nihilism's (joint) victory in the ontological parsimony stakes.

8.2.1 – Bennett's Challenge to the Presumptive Case

Since Bennett's argument is that eliminative nihilism is no less ontologically parsimonious than *any* rival answer to the SCQ, it won't hurt to pick a particular rival answer to compare to eliminative nihilism for the sake of presentational clarity. Thus I'll paint Bennett as arguing that eliminative nihilism is no more ontologically parsimonious than *universalism*. Further, Bennett hedges in the course of her argument and claims that eliminative nihilism and universalism merely 'come out *roughly* on a par' with regard to ontological parsimony (2009: 64, my emphasis), and that 'it is not obvious' that eliminative nihilism is more parsimonious than universalism (63). But her reasoning appears to support the unhedged conclusion that eliminative nihilism and universalism are equally ontologically parsimonious, so I'll characterise her as endorsing this conclusion.

Here's a brief reconstruction of Bennett's argument.

- (1) Theory A is more quantitatively parsimonious than theory B iff A posits fewer entities than B.
- (2) Properties are entities.

(2) is controversial, but Bennett claims that if it's false then a more complicated nominalist-friendly version of her argument invoking ideological parsimony can be made instead (2009: 64). A full discussion of ideological parsimony is outside the scope of this dissertation, but, in the interests both of granting as much to my opponents as possible and of heading off the suspicion that eliminative nihilism secures its victory in the ontological parsimony stakes only by countenancing an extremely complicated ideology, I'll grant (2) for argument's sake.

- (3) For every composite entity posited by universalism, there is a perplural property (I'll say more about these below) posited by eliminative nihilism.

This is the crucial premise. I'll elaborate on it and present Bennett's defence of it below. For now, moving on:

¹⁷⁸ Rival views to eliminative nihilism may not agree that the particles that the eliminative nihilist regards as simple* are really simple*, but must posit these particles nonetheless (the exception to this is *existence monism*, which I'll discuss in §8.3). Indeed, if rivals views *do* deny that particles such as quarks and leptons are simple, then they apparently bloat their ontology even further in comparison to the eliminative nihilist's, because their ontology must additionally contain the parts of quarks and leptons. So, to make my opponent's case as strong as possible, I'll assume in what follows that all sides agree that quarks and leptons are simple*.

(4) So, eliminative nihilism and universalism posit the same number of entities.

(5) So, eliminative nihilism and universalism are equally quantitatively parsimonious.

The argument as stated isn't valid. So Bennett must be assuming the following premise between (3) and (4), which validates the argument:

(3.5) If we ignore the universalist's composite entities and the eliminative nihilist's alleged perplural properties, universalism and eliminative nihilism posit the same number of entities.

An inventory of the eliminative nihilist and universalist ontologies demonstrates the plausibility of (3.5). First, both ontologies contain the simples* (i.e. the quarks and leptons). Second, both ontologies contain the singular properties of the simples*, such as *having +1/2 spin*. Third, both ontologies contain the plural arrangement-properties of the simples*. Plural properties are properties that are instantiated by multiple *xs*, without entailing the existence of any singular composite* entity somehow 'made up of' the *xs*. Plural *arrangement*-properties determine how some things are arranged.¹⁷⁹ Fourth, the eliminative nihilist ontology contains plural properties that don't concern the arrangement of the simples*, such as [*collectively*] *weighing 20kg*; the universalist ontology contains singular properties of composites such as *weighing 20kg*. There appears to be a one-to-one correspondence between these eliminative nihilist plural non-arrangement-properties and universalist singular properties of composites.

This inventory seems exhaustive, so I'll grant (3.5).

Bennett's argument now turns on (3). In the remainder of this section I present her defence of it.

Bennett argues (2009: 60, 64) that the eliminative nihilist cannot give an adequate account of the often *complexly structured nature* of reality with an ontology consisting only of simples*, singular properties, and plural properties. 'Complexly structured' here applies to any structure that has substructures. So consider a (putative) molecule, for example. Best science has it that molecules are made out of atoms, which (roughly speaking) are made out of quarks and leptons. Putative molecules thus have a complex structure: an overall molecular structure with atomic substructures. So putative molecules are a good example of the bits of reality that Bennett thinks the eliminative nihilist is going to have difficulty in adequately accounting for.

Now, the eliminative nihilist usually takes herself to be able to adequately account for putative molecules (and all other putative composite* entities for that matter) by adopting a strategy familiar from van Inwagen (1990a: 98-114) and Rosen & Dorr (2002). Where there appears to be a composite* molecule, the eliminative nihilist claims that there are instead just some simples* arranged molecule-wise. In other words, the eliminative nihilist takes herself to be able to use the simples* and the plural arrangement-properties (in particular, the property of *being arranged molecule-wise*) that are already in her ontology to adequately account for putative molecules. After all, the empirical data apparently underdetermines whether there are composite* molecules or just

¹⁷⁹ One might object that the universalist doesn't need to posit plural arrangement-properties. But I'll assume that she does, for three reasons: 1) Bennett agrees (2009: 64); 2) It's just implausible to suggest that the simples* aren't arranged: they stand in spatial relations to one another, which appears to be all that's required to be arranged; 3) To anticipate an argument I make in §8.2.3, the universalist needs to say that the simples* are arranged e.g. *F*-wise in order to explain why they compose an *F* (rather than a *G*).

simples* that instantiate *being arranged molecule-wise*, so the eliminative nihilist account appears to be consistent with best science and therefore adequate.

But Bennett argues that, to give an adequate account of reality, the eliminative nihilist must tell a more complicated story about the way the simples* are arranged than simply saying that they are arranged molecule-wise (2009: 59-60). Simply saying that the simples* are arranged molecule-wise *doesn't* give an adequate account of putative molecules because it doesn't capture their complex structure and therefore *isn't* consistent with best science. Accounting for putative molecules by claiming that there are simples* that instantiate *being arranged molecule-wise* only captures the overall molecular structure of the putative molecule. It says nothing about, and therefore fails to capture, the atomic substructures, and is analogous to some ill-informed scientist claiming that the quarks only compose molecules, without first composing atoms.¹⁸⁰

Given this, the eliminative nihilist might be tempted to say instead that the simples* instantiate *being arranged atom-wise*, and it's the *atom-wise arrangements* of simples* that instantiate *being arranged molecule-wise*. This would allow her to capture both the overall molecular structure and the atomic substructures of the putative molecule. But this seems to involve reifying 'atom-wise arrangements' (i.e., presumably, composite* atoms), which the eliminative nihilist cannot accept.

So the eliminative nihilist must find some other way of capturing, or accounting for, complex structures such as that of the putative molecule. In Bennett's words, the eliminative nihilist must 'preserve compositional structure' when she gives her composite*-free account of reality (60). Her ontology of simples* and plural arrangement-properties alone doesn't give her the tools she needs to do this (and the other plural properties and singular properties in her ontology aren't going to help).

Bennett concludes that the eliminative nihilist must inflate her ontology with new, complex arrangement-properties that the eliminative nihilist can say are instantiated by the simples*, thereby allowing her to capture complex structures in reality (2009: 60). I'll call these 'superplural arrangement-properties', because the bit of language that describes them is superplural predication (cf. Hazen 1997: 247). I've already given a short characterisation of superplural logic in §2.4.3.2, though I didn't mention superplural properties. Briefly: superplural properties are to plural properties what plural properties are to ordinary singular properties. Whilst plural properties tell us what properties groups of *x*s have, without entailing the existence of a composite* 'group of *x*s', superplural properties tell us what properties groups of groups of (groups of groups of... etc.) *x*s have, without entailing the existence of a singular composite* 'group of groups of *x*s'.¹⁸¹ Superplural *arrangement*-properties allow us to capture complex structures whilst only committing us to the existence of the basic building blocks of the structures. For example, to capture the complex structure of a putative molecule, the eliminative nihilist can say that the simples* instantiate the superplural property of *being arranged ((atom-wise) molecule-wise)*. This complex property captures how the simples* are arranged into atomic structures, which are themselves arranged into a molecular structure, but which doesn't entail the existence of atomic structures (i.e., atoms).

¹⁸⁰ A natural response here for the eliminative nihilist is that the atomic substructures are somehow *included* in what it is for the simples* to be arranged molecule-wise. But this is just to say that the predicate 'arranged molecule-wise' picks out a complicated property that captures both the atomic substructures and overall molecular structure, which is exactly the sort of property that, as we'll see, Bennett is trying to get the eliminative nihilist to posit here.

¹⁸¹ So the property expressed by 'are separated from each other' as it appears in P and Q in §2.4.3.2 is a *prima facie* good example of a superplural property.

So the eliminative nihilist apparently must posit superplural arrangement-properties to account for complex structures in reality. And, claims Bennett, the eliminative nihilist must posit one of these superplural properties for every composite* entity the universalist posits (2009: 65)¹⁸² (for example, where the universalist posits a composite molecule, the eliminative nihilist must posit *being arranged ((atom-wise) molecule-wise)*). If this is right, then it follows that (3) holds.

Given that I've granted the other premises of Bennett's argument, if her defence of (3) succeeds then her conclusion that eliminative nihilism and universalism are equally quantitatively parsimonious follows. In the next subsection, however, I'll say why (3) is false and should be replaced with a different premise that generates a different conclusion: that eliminative nihilism is more quantitatively parsimonious than universalism.

8.2.2 – A Recount

Premise (3) of Bennett's argument (as I've reconstructed it) is false. Though I concede that the eliminative nihilist must posit superplural arrangement-properties for the reasons Bennett gives, contra Bennett the eliminative nihilist posits *fewer* superplural arrangement-properties than the universalist posits composite* entities.

My argument here will make use of the notion of *first-level* and *higher-level* composites*, which I define in terms of being *directly* and *indirectly* composed by the simples* respectively. So let us say that:

x is *directly* composed* by the y s iff x is composed by the y s and there are no z s such that the z s are composed* by the y s and x is composed* by the z s;

x is *indirectly* composed* by the y s iff x is composed by the y s but there are some z s such that x is composed by the z s and the z s are composed by the y s;

x is a *first-level* composite* iff x is directly composed* by the simples*;

x is a *higher-level* composite* iff x is indirectly composed* by the simples*.

For example, according to the universalist, molecules are higher-level composites* because they're indirectly composed* by the simples*: a molecule is composed* by the simples*, but it is also composed* by the atoms which themselves are composed* by the simples*. In contrast, the universalist's *first-level* composites* are the simple*-pairs, where 'simple*-pair' refers to the mereological fusion of two simples*: for any simple*-pair, there can be no other entities that are composed by two simples* that themselves compose the simple*-pair.

We can now deploy these terms to undermine Bennett's argument. To establish (3), Bennett needs the eliminative nihilist to have to posit as many superplural properties as the universalist posits composite* entities, both higher-level *and* first-level. But the eliminative nihilist need only posit as many superplural properties as the universalist posits *higher-level* composite* entities. This is because Bennett's argument only compels the eliminative nihilist to posit superplural properties to account for complexly structured putative composites*, and the only complexly structured putative composites* are the higher-level ones. Putative first-level composites* have no internal substructure as they are composed* directly by the simples*: simple*-pairs only have an overall 'simple*-pair' structure, and since the simples* that (directly) compose* them are simple*, they have no internal substructure that

¹⁸² Note that Bennett uses 'perplural' where I use 'superplural'.

needs to be captured by the eliminative nihilist with a superplural property. As such, the eliminative nihilist doesn't need to posit superplural properties to account for putative first-level composites*. She can instead simply use simples* and merely *plural* arrangement-properties to account for them. Where the universalist posits a composite* simple*-pair, the eliminative nihilist need only posit the plural property of *being arranged simple*-pair-wise*.

So (3) is false and should be replaced with

- (3*) The eliminative nihilist posits *fewer* superplural properties than the universalist posits composite* entities.

(1), (2), (3*), (3.5) and (4) together entail my conclusion:

- (5*) Eliminative nihilism is more quantitatively parsimonious than universalism.

8.2.3 – Bennett's Argument Redirected

Bennett's argument can be redirected so as to attack the claim that eliminative nihilism is more qualitatively parsimonious than universalism. Whilst universalism countenances a *kind* of entity that the eliminative nihilist doesn't, i.e. composite* entities, Bennett's reasoning apparently demonstrates that the eliminative nihilist must in turn posit a kind of entity that the universalist doesn't, i.e. superplural properties. If this is right then the eliminative nihilist cannot claim victory in the qualitative parsimony stakes.

In response, I claim that the universalist must posit superplural properties too, thus tipping the balance of qualitative parsimony back into the eliminative nihilist's favour. I have three arguments for this.

First argument. Facts about the ways some simples* are arranged seem to depend solely on facts about the locations of the simples*, not on facts about whether the simples* compose anything or not (Brenner 2015: 1311). Since the eliminative nihilist and the universalist agree on the locations of the simples*, they ought to agree on the way the simples* are arranged (even though they disagree on whether the quarks also compose* something). For the sorts of reasons I elaborated on in §8.2.1, the eliminative nihilist must in some cases say that the simples* are arranged ((atom-wise) molecule-wise).¹⁸³ So the universalist must say so too. Thus the universalist must posit superplural properties such as *being arranged ((atom-wise) molecule-wise)*.

Second argument. Most believers in composites*, including universalists, accept that composites* are not fundamental, but rather exist in virtue of their parts*.¹⁸⁴ For example, the molecule is usually thought to exist in virtue of the atoms that compose* it. Plausibly, one of the features of the in-virtue-of relation is that if *x* exists in virtue of *y* then then facts about *y* fix (or determine) all of the facts

¹⁸³ Actually, given that the universalist's first-level composites are composite pairs of simples*, not atoms, she must actually say something closer to 'the simples* are arranged (((simple-pair-wise) atom-wise) molecule-wise)'. Nothing turns on this wrinkle, which I'll ignore in what follows.

¹⁸⁴ E.g. Cameron (2014), Korman (2016: §4), Koslicki (2008: 186), and Wilson (2014) accept this more or less explicitly. Even those who don't explicitly accept it are most naturally read to be tacitly endorsing it, and the claim seems required in order for believers in composites to be able to explain the intimate connection between composites* and their parts*, as demonstrated by e.g. the fact that, the parts* always go wherever the whole they compose* goes (see Cameron [2014] for elaboration of this point). Some believers in composites* might follow Schaffer (2010a) in accounting for this intimate connection by adopting priority monism, the view that composites* exist in virtue of the larger composites* of which they are parts*, but this view is non-standard (priority monists can ignore my second argument here and instead consider my first and third arguments instead).

about x , such that fixing the facts about y is *all God has to do* in order to fix the facts about x (e.g. Schaffer 2009: 351).¹⁸⁵ So facts about the atoms are supposed to fix all of the facts about the molecule. Furthermore, since atoms are supposed to exist in virtue of the simples* that compose* them, facts about the simples* are supposed to fix all of the facts about the atoms.¹⁸⁶

So far so unproblematic for the universalist. But there's a twist. The in-virtue-of relation is transitive (e.g. Audi 2012: §4): if x exists in virtue of y , and y exists in virtue of z , then x exists in virtue of z . If the molecule exists in virtue of the atoms, and the atoms exist in virtue of the simples*, then the molecule exists in virtue of the simples*. This implies that facts about the simples* must fix all of the facts about the molecule. Crucially, one of the facts about the molecule that must be fixed by facts about the simples* is the fact about the molecule's complex structure. What fact can this be? This is essentially the same question that the eliminative nihilist was faced with in §8.2.1. There we agreed with Bennett that the fact that the simples* instantiate the plural arrangement-property *being arranged molecule-wise* won't do the trick, as this doesn't fix the facts about the atomic substructures within the molecule; we can also agree that the fact that the simples* instantiate the plural arrangement-property *being arranged atom-wise* won't do either, because it fails to fix the overall molecular structure: each group of simples* that are arranged atom-wise could be distant to the others, thus not forming an overall molecular structure at all. The only fact about the simples* that seems to fix the fact of the molecule's complex structure is that they instantiate *being arranged ((atom-wise) molecule-wise)*. So the universalist must accept that the simples* instantiate *being arranged ((atom-wise) molecule-wise)*, and must therefore posit this superplural property.¹⁸⁷

Clearly the point generalises: for every higher-level composite* that she posits, the universalist must also posit a superplural arrangement-property in order to be able to account for what it is in virtue of which that particular composite* exists.

Third argument. Perhaps some will take issue with the idea that composite* entities exist in virtue of their parts in the way relied upon in the previous argument. But we can make a similar point by claiming only that *properties of composite entities supervene (with metaphysical necessity) on properties of their parts*. This, I take it, is even less controversial than the claim that composite* entities exist in virtue of their parts (for example, even priority monists, who won't accept the in-virtue-of claim, can accept the supervenience claim).

So consider again the molecule. Its properties supervene on those of the atoms that compose* it; the properties of the atoms in turn supervene on the properties of the simples* that compose* them. Supervenience is clearly transitive (if there can be no A-difference without B-difference, and no B-

¹⁸⁵ Note that this claim is weaker than it might appear. In particular, the claim doesn't require that, if x exists in virtue of y , then facts about x are *derivable* (even in principle) from facts about y . The claim is simply that, by fixing the facts about y we thereby fix the facts about x , even if it's in principle impossible to know what those facts about x are given only facts about y .

¹⁸⁶ It's possible to hold that universalists should go in for a global rather than local in-virtue-of story: rather than holding, for example, that local facts about particular atoms fix the facts about the particular molecules they putatively compose*, perhaps we should take the view that, globally, quark-level facts fix the atomic-level facts, and that atomic-level facts fix the molecular-level facts (c.f. Potochnik [2010: 62-3]). But if such a view is correct my argument still demonstrates that one of the quark-level facts must be that the quarks instantiate *being arranged ((atom-wise) molecule-wise)*, which gets me my desired conclusion.

¹⁸⁷ The universalist might object that a fact about the simples* *other* than a fact about what (superplural) arrangement-property they instantiate can fix the fact about the molecule's complex structure, thus allowing her to avoid positing superplural arrangement-properties. But if such a fact could be found then the eliminative nihilist could make use of it too, thus also allowing *her* to avoid having to posit superplural arrangement-properties to capture the molecule's complex structure.

difference without C-difference, then there can be no A-difference without C-difference – see also e.g. McLaughlin and Bennett 2018: §3.2). Thus the properties of the molecule supervene on properties of the simples* that compose* it.¹⁸⁸ In other words: properties of the simples* *suffice* for the properties of the molecule.¹⁸⁹ One of the properties of the molecule is the one that characterises its complex structure, with its overall molecular structure and its atomic substructures. Thus the universalist needs to say that among the properties of the simples* is a property (or a set of properties) that suffices for the property of the molecule that characterises its complex structure. What property of the simples* could this be? Again, this is essentially the same question that the eliminative nihilist was faced with in §8.2.1: how to account for the putative molecule and its complex structure with only simples* and properties of the simples*. And again, the result is that only by saying that the simples* instantiate the superplural arrangement-property of *being arranged ((atom-wise) molecule-wise)* can we maintain that the properties of the simples* suffice for the properties of the molecule, and thus that the properties of the molecule supervene on the properties of the simples* (and if some other property of the simples*, P, would do the trick instead, then the eliminative nihilist could claim that the fact that the simples* instantiate P allows her to capture the complex structure of the putative molecule, thus escaping commitment to superplural arrangement-properties herself). So, again, universalists must accept that the simples* instantiate *being arranged ((atom-wise) molecule-wise)*, and must therefore posit this superplural property.

We thus have three independent arguments that demonstrate that the universalist must join the eliminative nihilist in positing superplural properties. Since the universalist *also* posits an extra kind of entity, i.e. composite* entities, her theory is less qualitatively parsimonious than the eliminative nihilist's.

Let me now consider an objection to my response to Bennett's redirected argument here.

I've suggested that composite entities, simples*, superplural properties, and non-superplural properties constitute distinct kinds from one another, and have used this to argue that eliminative nihilism is more qualitatively parsimonious than universalism. But one might wonder whether we could divide the world up into kinds in a different way, thereby getting different results about which theory is the more qualitatively parsimonious. For example, perhaps the relevant kinds for the purposes of qualitative parsimony comparisons are *entities* and *properties*; since eliminative nihilism and universalism both posit these, this would make the two theories come out as equally qualitatively parsimonious, even if I'm right about the universalist having to join the eliminative nihilist in positing superplural properties.

¹⁸⁸ Again (see fn 186), I talk here as though supervenience is local, such that if some simples* compose* a composite* *F* then the properties of that *F* supervene on the properties of those simples*. One might instead prefer to tell a global supervenience story, and to say that the set of all properties instantiated by composites* supervene on the set of all properties instantiated by simples*. But again, my argument here still demonstrates that among the properties of the simples* must be superplural arrangement-properties such as *being arranged ((atom-wise) molecule-wise)*. For brevity, then, I'll continue to use a notion of local supervenience.

¹⁸⁹ For suppose the properties of the simples* *didn't* suffice for the properties of the molecule. Then there is a possible world that is exactly the same as the actual world with regard to the properties of the simples*, but that which differs from the actual world with regard to the properties of the molecule. That would be an example of a difference with regard to the properties of the molecule without a difference with regard to the properties of the quarks, contradicting the supervenience claim.

I know of no way to conclusively establish that my way of carving the world up into kinds is the right way. But, given the dialectic, I don't really need to. No matter how we decide to count kinds, my argument in §8.2.2 leaves the eliminative nihilist in a strong dialectical position.

This is because, if my argument above is sound, the universalist must join the eliminative nihilist in positing superplural properties, which means that the universalist posits everything the eliminative nihilist does. So the universalist posits *at least* as many kinds of entity as the eliminative nihilist. But the universalist also posits composite* entities. Now, on many natural and plausible ways of counting kinds, this extra commitment to composite* entities brings with it a commitment to at least one extra kind of entity. If so, then my argument that eliminative nihilism is more qualitatively parsimonious than universalism goes through straightforwardly. But even if the universalist's commitment to composite* entities *doesn't* commit her to a new kind of entity (because composite* entities belong to a kind that she and the eliminative nihilist already posit), her theory is still only equally as qualitatively parsimonious as eliminative nihilism. And though this means that the eliminative nihilist has lost (or, really, *tied*) the qualitative parsimony battle, she has still won the ontological parsimony war. For I've already demonstrated that eliminative nihilism is more *quantitatively* parsimonious than universalism, and even those who doubt the importance of quantitative parsimony in theory choice ought, as Divers (1994: 388) convincingly argues, to at least regard it as a tie-breaker when trying to assess the ontological costs of two equally qualitatively parsimonious theories.

In a sense, then, little is at stake here: if I'm wrong about how to count kinds, then eliminative nihilism still ends up being at least equally as qualitatively parsimonious as universalism, which is enough to guarantee eliminative nihilism a win in the overall ontological parsimony stakes. But my way of counting kinds is natural, and consonant with the intuitive thought that a commitment to composite* entities brings with it a commitment to at least one new kind of entity. So I see no reason to depart from my assumption that composite* entities, simples*, superplural properties, and non-superplural properties constitute distinct kinds from one another. I therefore stand by my claim that eliminative nihilism is more qualitatively parsimonious than universalism.

8.2.4 – Summary

Thus I conclude that the presumptive case in favour of eliminative nihilism's greater ontological parsimony in comparison to universalism survives Bennett's challenge. Moreover, my arguments here have not been sensitive to any features of universalism not shared by rival answers to the SCQ that posit at least one composite* entity, I draw the wider conclusion that the presumptive case in favour of eliminative nihilism's greater ontological parsimony than all such rival answers to the SCQ (other than those that accept CAI) is vindicated.

8.3 – Other Nihilist Views

Eliminative nihilism is not just in competition in the ontological parsimony stakes with non-nihilist views, but also with other forms of nihilism: emergence nihilism and existence monism.

8.3.1 – Emergence Nihilism

We need not spend too long considering emergence nihilism, because it's obvious that eliminative nihilism is more ontologically parsimonious than this form of nihilism. As we've seen (§2.3), emergence nihilism posits not only sub-atomic particles such as quarks and leptons, but also larger entities such as atoms and molecules, and even macroscopic entities like tables, chairs, planets, penguins, etc. It distinguishes itself from non-nihilist answers to the SCQ by denying that the non-

sub-atomic entities it posits have parts*. But whilst this allows emergence nihilism to qualify as a form of nihilism, it doesn't prevent the view from being just as ontologically profligate as the non-nihilist answers to the SCQ that we concluded in the previous section to be less ontologically parsimonious than eliminative nihilism. So emergence nihilism is less ontologically parsimonious than eliminative nihilism.

8.3.2 – *Existence Monism*

Existence monism presents a more difficult case, however. Recall (from §2.3) that existence monism is the view that the only material entity that exists is the universe. The universe is thus simple*: if it had parts*, then it wouldn't be the only material entity. The existence monist accounts for the *apparent* existence of parts* of the universe by appealing to vastly complex distributional properties that characterise the universe's internal heterogeneity. For example, though it might appear that there is a *table* in a certain region of space, existence monists claim that what is really going on is that the universe instantiates a complex distributional property that (among many other things) characterises the way in which the universe is 'aspected table-ishly' at that region of space (see Schaffer 2007: 179).

It may seem as though existence monism offers greater ontological parsimony than eliminative nihilism. Schaffer certainly seems to think so (2007: §VII).¹⁹⁰ In his words: 'the [eliminative] nihilist demands the simplest sufficient ontology, and the monist delivers it' (2007: 175). Indeed, although Schaffer doesn't make this explicit, existence monism is ostensibly both more quantitatively and qualitatively parsimonious than eliminative nihilism: more quantitatively parsimonious because it posits just one material entity (the universe) and one property (the distributional property instantiated by the universe) in place of a great many entities and properties (all of the sub-atomic particles plus all of their individual properties and the plural properties they instantiate); more qualitatively parsimonious because it posits just one kind of entity (the universe) in place of the various kinds of sub-atomic particle (e.g. the kinds *quark* and *lepton*).¹⁹¹

Nevertheless, I think that we should reject the claim that existence monism is more ontologically parsimonious than eliminative nihilism. My argument for this is a flat-footed one: it simply cannot be correct to say that consideration of ontological parsimony in general should push us towards favouring theories that posit one big thing in place of several smaller things. Indeed, consideration of the fact that theories can reduce the number of entities that they posit simply by claiming that what appear to be several small, simple* entities are really one, big, complex one tells us, it seems to me, that we should augment our way of measuring ontological parsimony in such a way as to be sensitive not only to the *number* of (kinds of) entity a theory posits, but also the *size* and *complexity* of those entities.¹⁹² Once we've done this, we'll see that the extreme reduction in the *number* of entities that existence monism posits is counterbalanced in the extreme increase in the *size* and *complexity* of the one entity it does posit.

¹⁹⁰ See also Tallant (2014: 1521-3), who agrees with Schaffer's judgement here.

¹⁹¹ Further, since existence monism posits just one property, it posits just one *kind* of property. That means that eliminative nihilism at best posits the same amount of kinds of property as the existence monist, and plausibly (depending on what the best way to carve up properties into *kinds* of properties is) posits more.

¹⁹² I note that by augmenting our way of measuring ontological parsimony in the way I've suggested would preserve the results established so far, because the question of *size* or *complexity* of the entities posited by eliminative nihilism and its non-nihilist rivals (and indeed of those posited by emergence nihilism) is not relevant to my argument for eliminative nihilism's victory over those theories in the ontological parsimony stakes.

I've reported my intuition on the matter; here's a more fleshed out argument. It concerns the fact that different forms of existence monism can surely differ with regard to ontological parsimony, despite all positing the same number of entities. Consider two theories, T1 and T2, that agree on everything except for the fact that T2 additionally posits an empirically undetectable teapot in front of you. T2's empirically undetectable teapot is completely explanatorily idle. This seems to be a paradigm case in which consideration of ontological parsimony should tell us not to posit the teapot, and to endorse T1 over T2. But now consider T3 and T4, which can be loosely described as the existence monist versions of T1 and T2 respectively. That is, T3 and T4 agree on everything – they both posit one material entity, i.e. the universe, and one vastly complex distributional property characterising how the universe is here 'aspected tablishly' and there 'aspected cup-of-tea-ishly', etc – except for the fact that T4 claims that the vastly complex distributional property instantiated by the universe is slightly more complex than the one posited by T3, in that it *additionally* characterises the universe as being 'aspected empirically-undetectable-teapot-ishly' right in front of you (or rather, right in front of the you-ishy-aspected bit of the universe). Now, anyone who thinks that T1 is more ontologically parsimonious than T2 should also find it intuitive to say that T3 is more ontologically parsimonious than T4. Otherwise how are we to explain why T3 is preferable to T4 (as it surely is)? It doesn't seem as though T3 is more *ideologically* parsimonious than T4, for example, for neither theory seems to posit more primitive concepts than the other.¹⁹³ So, despite positing the same number of (kinds of) things (one universe and one complex distributional property), T3 seems ontologically more parsimonious than T4. Thus ontological parsimony should be sensitive at least to the *complexity* of the entities posited. Further, suppose T5 is identical to T3 except that T5 claims that the universe is twice as big as the universe that T3 posits. Loosely speaking, T5 posits *two* of the universes that T3 posits, with the second one being completely explanatorily idle – except of course T5, as an existence monist view, describes the situation as one in which the One is 'aspected two-universes-ishly'. Again, T3 seems to be clearly more ontologically parsimonious, despite the fact that it posits the same number of entities as T5, because the one material entity that it posits is only half as big as the one that T5 posits. Thus ontological parsimony should be sensitive to the *size* of the entities posited.

What's more, within the literature there is some evidence in support of my claim that ontological parsimony should be sensitive to the complexity and size of entity posited, in addition to just the number of entities posited. For example, Bennett, when discussing ontological parsimony, writes that 'the very simplest theory of the world is that there is nothing at all. The next simplest theory is that there is one *qualitatively homogenous*, mereologically simple thing' (2017: 218, emphasis mine). Thus Bennett apparently takes it to be obvious that qualitative homogeneity is ontologically simpler than qualitative heterogeneity, which is tantamount to accepting that the complexity of the entities in one's ontology is relevant to assessing how ontologically parsimonious that ontology is. Even Horgan and Potrč, who defend existence monism as a part of their defence of blobjectivism, concede that it might be that pointillism – i.e. roughly speaking, a version of eliminative nihilism where the simples* are taken to be spatiotemporal points – 'is not significantly worse off than blobjectivism on the score of parsimony in positing concrete particulars, since pointillism repudiates the blobject itself [i.e. the universe] while also repudiating any composite objects' (2008: 183). This suggests that they too are at

¹⁹³ Perhaps you think that T4 posits the extra primitive concept of *being aspected-empirically-undetectable-teapot-ishly* in comparison to T3. I'm dubious that this is an extra primitive concept, but even if it is, we can sidestep the issue by tweaking the example so that T3 claims that the universe is aspected-empirically-undetectable-teapot-ishly *once* whereas T4 claims that it is so aspected *thousands of times*. Then both theories need the concept of *being aspected-empirically-undetectable-teapot-ishly*, but T3 still seems preferable.

least open to the thought that positing a great many small, simple* entities might be roughly equally as parsimonious as positing one big complex one.

Now, of course, I concede my comments here don't amount to a full-blown defence of the thought that any adequate measure of ontological parsimony must be sensitive to the size and complexity of the entities posited, rather than just the number of entities are posited. But I think that I have done enough to make this thought plausible. What's more, it seems plausible to me that once we understand ontological parsimony in this way, existence monism and eliminative nihilism will turn out to be equally ontologically parsimonious. For, intuitively, the one entity that existence monism posits seems to be exactly as complex as the entities that the eliminative nihilist posits are numerous. And on the strength of these plausible thoughts, I conclude that existence monism and eliminative nihilism are equally ontologically parsimonious.¹⁹⁴

8.4 – Conclusion

I therefore conclude that eliminative nihilism is the joint most ontologically parsimonious answer to the SCQ. It is more ontologically parsimonious than all non-nihilist answers to the SCQ, because it posits only simples* and properties of the simples* where non-nihilist answers additionally posit at least some composite* entities; it is more ontologically parsimonious than emergence nihilism because it posits far fewer simples*; it is just as ontologically parsimonious as views that accept CAI and existence monism. Since eliminative nihilism and reductive nihilism share an ontology, that means that reductive nihilism is also a joint winner in the ontological parsimony stakes. Since ontological parsimony is very plausibly a theoretical virtue, that means that we have defeasible reason to accept the disjunction of reductive nihilism, eliminative nihilism, existence monism, and CAI, over rival answers to the SCQ.

¹⁹⁴ There may be those who would claim that I've misdescribed the nature of my argument here: really I'm not arguing that existence monism is just as ontologically parsimonious as eliminative nihilism, because ontological parsimony concerns the *number* of (or number of kinds of) entities posited; really what I'm doing is conceding that existence monism is more ontologically parsimonious than eliminative nihilism but claiming that there's a broader theoretical virtue – call it 'ontological simplicity' – that encompasses both number and complexity of entities posited and arguing that existence monism and eliminative nihilism tie with regard to this broader theoretical virtue. But a dispute on this point would be a mere verbal one, and I hereby stipulate that I'm here taking 'ontological parsimony' to refer to this wider notion of ontological simplicity.

Chapter 9

Conclusion

Over the course of this thesis, I have argued for two central claims.

The first of these claims (detailed in chapter 2 and argued for across chapters 3-5) is that, if microphysical nihilism is true, then the reductive nihilist's non-standard semantics follow. That is, if microphysical nihilism is true, then so are:

- (RN1) 'Existence' and its cognates express the plural existential quantifier, $\exists xx$.
- (RN2) 'Object' is not synonymous with 'entity', but is rather a plural predicate that can be satisfied by proper pluralities of entities. Roughly, xx satisfy 'object' iff it is intuitive that the region of space exactly occupied by xx is exactly occupied by an object.
- (RN3) Mereological terms such as 'parthood' and 'composition' are defined by PARTHOOD, COMPOSITION, etc.
- (RN4) Sortal predicates that purport to apply to single composite F s express properties that are instantiated by xx iff xx are both arranged F -wise and are an object.
- (RN5) Predicates that putatively ascribe properties to single things express plural properties that are satisfiable by proper pluralities (thus, for example, 'weighs 10kg' express *jointly weighing 10kg*).
- (RN6) Referring expressions of the form 'N' that purport to refer to a single composite F refer plurally to xx iff xx are arranged N -wise and xx are an object.

Given (RN1)-(RN6), microphysical nihilists can accept both that composite objects exist, and that reality genuinely has a mereological structure (in that some objects are parts of other objects, and some pluralities of objects compose different objects, etc.).

This has consequences for nihilists and non-nihilists alike. For microphysical nihilists, the consequence is that they should be reductive nihilists. For the semantic theses that reductive nihilism adds to microphysical nihilism are well-motivated by (at least) two different influential theories of meaning and reference-fixing, and are additionally robust in the face of objections, so anyone who accepts microphysical nihilism should accept the reductive nihilist's extra semantic theses too.

For non-nihilists (and even for existence monists and emergence nihilists), the consequence is that when they criticise microphysical nihilism and recommend their own theories of composition in its place, they should do so on the understanding that microphysical nihilism, properly understood, is reductive nihilism. Thus, for example, they shouldn't criticise microphysical nihilism on the grounds that it rejects the existence of tables, because microphysical nihilism, properly understood, is reductive nihilism, and reductive nihilism accepts the existence of tables (of course, if the microphysical nihilist with whom they are arguing explicitly accepts eliminative nihilism, then all bets are off).

The second conclusion that I have argued for in this thesis is simply that reductive nihilism is true. That is, it's not just that reductive nihilism is the best way to understand microphysical nihilism, it's that reductive nihilism is the best answer to the SCQ.

Though I endorse this second conclusion, I should admit to doing so with slightly less confidence than that with which I endorsed the first. That's because, although I have argued that reductive nihilism is both more intuitive and more ontologically parsimonious than (or at least, just as ontologically parsimonious as) all of its rivals, it remains open that some consideration *besides* intuitiveness and ontological parsimony will ultimately cause us to reject reductive nihilism. One such consideration might be that of (atomless) *gunk* (Lewis 1991: 20). An object is an example of gunk iff it is an instance of infinitely descending parts*. Sider (1993) has argued that gunk is conceivable, and therefore possible, and that this possibility implies the (actual) falsity of nihilism. Although I do not consider gunk to pose an insurmountable problem to reductive nihilism (for responses to the gunk objection to nihilism, see e.g. Cameron [2007: 101-2], Williams [2006], Sider [2013a: §§9-10]), it's probably fair to say that more work would need to be done to conclusively establish this (and certainly, I haven't done that work here).

That said, though, it's probably true of any theory about the world, no matter how well-confirmed and attractive, that there is some room for doubting whether it can answer any and all actual and possible objections to it. That seems like a poor reason to withhold belief from any and all theories. The fact that reductive nihilism satisfies our intuitions regarding the SCQ better than any other theory, and does so whilst remaining the (joint) most ontologically parsimonious answer to the SCQ, gives us strong reason to believe it. So we all have strong reason for accepting reductive nihilism as the right answer to the SCQ.

References

- Armstrong, D. (1989) *Universals: An Opinionated Introduction* Boulder, CO: Westview Press
- Audi, P. (2012) 'Grounding: Toward a theory of the in-virtue-of relation' in *Journal of Philosophy* 109: 685-711.
- Baron, S. and Tallant, J. (2018) 'Do Not Revise Ockham's Razor Without Necessity' in *Philosophy and Phenomenological Research* 96(3): 596-619
- Baxter, D. (1988) 'Identity in the Loose and Popular Sense' in *Mind* 97(388): 575-582
- Baxter, D. (2014) 'Identity, Discernibility, and Composition' in Cotnoir, A. and Baxter, D. (eds) *Composition as Identity* Oxford: OUP pp. 244-254
- Bennett, K. (2009) 'Composition, Colocation, and Metaontology' in Chalmers, D., Manley, M., and Wasserman, R. (eds) *Metametaphysics* Oxford: OUP pp. 38-76.
- Bennett, K. (2017) *Making Things Up* Oxford: OUP
- Benovsky, J. (2006) *Persistence Through Time, and Across Possible Worlds* Ontos Verlag
- Boolos, G. (1984) 'To Be Is To Be A Value Of A Variable (Or To Be Some Values Of Some Variables)' in *The Journal of Philosophy* 81(8) 430-449
- Braddon-Mitchell, D. & Nola, R. (eds) (2009) *Conceptual Analysis and Philosophical Naturalism* MIT Press
- Braun, D. (1993) 'Empty Names' in *Noûs* 27(4): 449-469
- Brenner, A. (2015) 'Mereological Nihilism and the Special Arrangement Question' in *Synthese* 192: 1295-1314.
- Brenner, A. (2017) 'Simplicity as a Criterion of Theory Choice in Metaphysics' in *Philosophical Studies* 174: 2687-2707
- Brenner, A. (2018) 'Science and the Special Composition Question' *Synthese* 35(4): 447-463
- Bricker, P. (2016) 'Composition as a Kind of Identity' in *Inquiry* 59(3): 264-294
- Calosi, C. (2016) 'Composition is Identity and Mereological Nihilism' in *The Philosophical Quarterly* 66: 219-235
- Cameron, R. (2007) 'The Contingency of Composition' in *Philosophical Studies* 136(1): 99-121
- Cameron, R. (2008) 'Truthmakers and Ontological Commitment: Or How to Deal with Complex Objects and Mathematical Ontology without Getting into Trouble' in *Philosophical Studies* 140(1): 1-18
- Cameron, R. (2010) 'How to Have a Radically Minimal Ontology' *Philosophical Studies* 151(2): 249-264
- Cameron, R. (2012) 'Composition as Identity Doesn't Settle the Special Composition Question' in *Philosophy and Phenomenological Research* 84(3): 532-554

- Cameron, R. (2014) 'Parts Generate the Whole But They Are Not Identical to It' in Cotnoir, A. and Baxter, D. (eds) *Composition as Identity* Oxford: OUP pp. 91-110
- Carey, S. (2009) *The Origin of Concepts*. Oxford: Oxford University Press
- Carmichael, M. (2015) 'Towards a Common-Sense Answer to the Special Composition Question' in *Australasian Journal of Philosophy* 93(3): 475-490
- Carnap, R. (1958) 'Beobachtungssprache und Theoretische Sprache' in *Dialectica* 12(3-4): 236-248
- Contessa, G. (2014) 'One's a Crowd: Mereological Nihilism Without Ordinary-Object Eliminativism' in *Analytic Philosophy* 55(2): 199-221
- Cornell, D. (2017) 'Mereological nihilism and the Problem of Emergence' in *American Philosophical Quarterly* 54(1): 77-87.
- Correia, F. and Schneider, B. (2012) 'Grounding: An Opinionated Introduction' in Correia, F. and Schneider, B. (eds) *Metaphysical Grounding: Understanding the Structure of Reality* Cambridge University Press pp. 1-36
- Cotnoir, A. (2014) 'Framing the Debate' in Cotnoir, A. and Baxter, D. (eds) *Composition as Identity* Oxford: OUP 3-23
- Cowling, S. (2013) 'Ideological Parsimony' in *Synthese* 190: 3889–3908
- Da Vee, D. (2020) 'Why Ockham's Razor should be preferred to the Laser' in *Philosophical Studies* 177: 3679-3694
- Davidson, D. (1973) 'Radical Interpretation' in *Dialectica* 27(3-4): 313-328
- Devitt, M. (1981) *Designation* Columbia University Press
- Devitt, M. and Sterelny, K. (1999) *Language and Reality* 2nd Edition Oxford: Blackwell
- Divers, J. (1994) 'On the Prohibitive Cost of Indiscernible Concrete Possible Worlds' in *Australasian Journal of Philosophy* 72: 384-9
- Donellan, K. (1974) 'Speaking of Nothing' in *The Philosophical Review* 83(1): 3-31
- Dorr, C. (2005) 'What We Disagree About When We Disagree About Ontology' in Kalderon, M. (ed.) *Fictionalism in Metaphysics* Oxford: OUP pp. 234-286
- Eddington, A. (1929) *The Nature of the Physical World* Cambridge: Cambridge University Press
- Fahrbach, L. (2005) 'Understanding Brute Facts' in *Synthese* 145(3): 449-466
- Fiddaman, M. and Rodriguez-Pereyra, G. (2018) 'The Razor and the Laser' in *Analytic Philosophy* 59(3): 341-358
- Frege, G. (1953) *The Foundations of Arithmetic* trans by J. L. Austin, 2nd Edition, New York: Harper and Brothers
- Grandy, R. (1973) 'Reference, Meaning, and Belief' in *Journal of Philosophy* 71: 439–452

- Grandy, R. (2016) 'Sortals' in *The Stanford Encyclopedia of Philosophy (Winter 2016 Edition)*, Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/sortals/>> Accessed 10.10.2020
- Grupp, J. (2006) 'Mereological Nihilism: Quantum Atomism and the Impossibility of Material Constitution' in *Axiomathes* 16: 245-386
- Hawley, K. (2001) *How Things Persist* Oxford: OUP
- Hawley, K. (2014) 'Ontological Innocence' in Cotnoir, A. and Baxter, D. (eds) *Composition as Identity* Oxford: OUP pp. 70-89.
- Hawley, K. (2020) 'Temporal Parts' in *The Stanford Encyclopedia of Philosophy (Summer 2020 Edition)*, Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2020/entries/temporal-parts/>> Accessed 3.3.2021
- Hawthorne, J. (1994) 'A Corrective to the Ramsey-Lewis Account of Theoretical Terms' in *Analysis* 54(2): 105-110
- Hawthorne, J. and Michael, M. (1996) 'Compatibilist Semantics in Metaphysics: A Case Study' in *Australasian Journal of Philosophy* 74(1): 117-134
- Hazen, A. (1997) 'Relations in Lewis's Framework Without Atoms' in *Analysis* 57(4): 243-248
- Hewitt, S. (2012) 'Modalising Plurals' in *Journal of Philosophical Logic* 41: 853-875
- Hirsch, E. (2002) 'Quantifier Variance and Realism' *Philosophical Issues* 12: 51-73
- Hirsch, E. (2009) 'Ontology and Alternative Languages' in Chalmers, D., Manley, M., and Wasserman, R. (eds) *Metametaphysics* Oxford: OUP pp. 231-258
- Hirsch, E. (2011) *Quantifier Variance and Realism: Essays in Metaontology* Oxford: OUP
- Hoffman, J. and Rosenkrantz, G. (1997) *Substance: Its Nature and Existence* Routledge
- Horgan, T. (1999) 'On What There Isn't' in *Philosophy and Phenomenological Research* 53(3): 693-700
- Horgan, T. and Potrč, M. (2000) 'Bobjectivism and Indirect Correspondence' in *Facta Philosophica* 2: 249-270
- Horgan, T. and Potrč, M. (2008) *Austere Realism: Contextual Semantics Meets Minimal Ontology*, Cambridge, MA: MIT Press
- Hossack, K. (2000) 'Plurals and Complexes' in *British Journal for the Philosophy of Science* 51(3): 411-443
- Jackson, F. (2000) *From Metaphysics to Ethics* Oxford: OUP
- Jackson, F. & Pettit, P. (1995) 'Moral Functionalism and Moral Motivation' in *The Philosophical Quarterly* 45(178): 20-40
- Johnston, M. (1992) 'How to Speak of the Colors' in *Philosophical Studies* 68(3): 221-263

- Kaplan, D. (1979) 'Transworld Heir Lines' in Loux, M. (ed.) *The Possible and the Actual: Readings in the Metaphysics of Modality* Cornell University Press pp. 88-109
- Korman, D. (2008a) 'Unrestricted Composition and Restricted Quantification' in *Philosophical Studies* 140(3): 319-334
- Korman, D. (2008b) Review of *Austere Realism: Contextual Semantics Meets Minimal Ontology* by Horgan and Potrč, in *Notre Dame Philosophical Reviews*
- Korman, D. (2009) 'Eliminativism and the Challenge from Folk Belief' in *Noûs* 43(2): 242-264
- Korman, D. (2015a) 'Fundamental Quantification and the Language of the Ontology Room' in *Noûs* 49(2): 298-321
- Korman, D. (2015b) *Objects: Nothing Out of the Ordinary* Oxford: OUP
- Korman, D. (2016) 'Ordinary Objects' in *The Stanford Encyclopedia of Philosophy (Spring 2016 edition)*, ed. E. Zalta. URL = <plato.stanford.edu/archives/spr2016/entries/ordinary-objects/>. Accessed 3.3.2021.
- Korman, D. (2019) 'Easy Ontology without Deflationary Metaontology' in *Philosophy and Phenomenological Research* 99(1): 236-243
- Koslicki, K. (2008) *The Structure of Objects*. Oxford: OUP
- Kriegel, U. (2008) 'Composition as a Secondary Quality' in *Pacific Philosophical Quarterly* 89: 359-383
- Kripke, S. (1976) 'Is There a Problem About Substitutional Quantification?' in Evans, G. and McDowell, J. (eds.) *Truth and Meaning* Oxford: Clarendon Press pp. 379-380
- Kripke, S. (1980) *Naming and Necessity* Oxford: Blackwell
- Lewis, D. (1966) 'An Argument for the Identity Theory' in *Journal of Philosophy* 63(1): 17-25
- Lewis, D. (1968) 'Counterpart Theory and Quantified Modal Logic' in *Journal of Philosophy* 65: 113-126
- Lewis, D. (1970) 'How to Define Theoretical Terms' in *Journal of Philosophy* 67(13): 427-446
- Lewis, D. (1971) 'Counterparts of Persons and Their Bodies' in *Journal of Philosophy*, 68(7): 203-211
- Lewis, D. (1972) 'Psychophysical and Theoretical Identifications' in *Australasian Journal of Philosophy* 50(3): 249-258
- Lewis, D. (1973) *Counterfactuals* Oxford: Blackwell
- Lewis, D. (1974) 'Radical Interpretation' in *Synthese* 23: 331-344.
- Lewis, D. (1983) *Philosophical Papers*, Volume I, Oxford: Oxford University Press.
- Lewis, D. (1986) *On the Plurality of Worlds* Oxford: Blackwell

- Lewis, D. (1991) *Parts of Classes* Oxford: Blackwell.
- Lewis, D. (1997) 'Naming the Colours' in *Australasian Journal of Philosophy* 75(3): 325-342
- Liggins, D. (2008) 'Nihilism Without Self-Contradiction' Le Poidevin, R. (ed.) *Being: Developments in Contemporary Metaphysics (Royal Institute of Philosophy Supplement 62)*. Cambridge: Cambridge University Press, 177–196.
- Liggins, D. (2020) 'Against Hermeneutic Fictionalism' in Armour-Garb, B. & Kroon, F. (eds.) *Fictionalism in Philosophy* New York: Oxford University Press pp. 81-102
- Linnebo, Ø. (2004) 'Plural Quantification Exposed' in *Noûs* 37: 71-92.
- Linnebo, Ø, (2017) 'Plural Quantification' in *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2017/entries/plural-quant/> Accessed 3.3.2021
- Linnebo, Ø. and Nicolas, D. (2008) 'Superplurals in English' in *Analysis* 68(3): 186-197
- Locke, J. ([1690] 1959) *Essay Concerning Human Understanding* A.C. Fraser (ed.) New York, Dover
- Long, J. (2019) 'One's an Illusion: Organisms, Reference, and Non-Eliminative Nihilism' in *Philosophia* 47: 459-475
- Loss, R. (2018) 'A Sudden Collapse to Nihilism' in *The Philosophical Quarterly* 68(271): 370-375)
- Lowe, J. (1995) 'Coinciding Objects: In Defence of the Standard Account' in *Analysis* 55(3): 171-178
- Lowe, J. (1997) 'Ontological Categories and Natural Kinds' in *Philosophical Papers* 26(1): 29-46
- Markosian, N. (1998) 'Brutal Composition' in *Philosophical Studies* 92(3): 211-249
- Markosian, N. (2008) 'Restricted Composition' in Hawthorne, J., Sider, T., and Zimmerman, D. (eds) *Contemporary Debates in Metaphysics* Oxford: Blackwell pp. 341-363
- McDaniel, K. (2007) 'Brutal Simples' in Zimmerman, D. (ed.) *Oxford Studies in Metaphysics* Volume 3 Oxford: OUP pp. 233-265
- McDaniel, K. (2010) 'Composition as Identity Does Not Entail Universalism' in *Erkenntnis* 73: 97-100
- McDaniel, K. (2013) 'Existence and Number' in *Analytic Philosophy* 54(2): 209-228
- McLaughlin, B. and Bennett, K. (2018) 'Supervenience' in *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2018/entries/supervenience/> Accessed 3.3.2021.
- McKay, T. (2006) *Plural Predication* Oxford: OUP
- Menzies, P. (1996) 'Probabilistic Causation and the Pre-emption Problem' *Mind* 105: 85-117
- Menzies, P. and Price, H. (2009) 'Is Semantics in the Plan' in Braddon-Mitchell, D. & Nola, R. (eds) *Conceptual Analysis and Philosophical Naturalism* MIT Press pp. 183-200

- Merricks, T. (1999) 'Composition as Identity, Mereological Essentialism, and Counterpart Theory' in *Australasian Journal of Philosophy* 77(2): 192-195
- Merricks, T. (2001) *Objects and Persons* Oxford: OUP
- Merricks, T. (2005) 'Composition and Vagueness' in *Mind* 114: 615-637.
- Merricks, T. (2015) *Propositions* Oxford: OUP
- Morton, A. (1975) 'Complex Individuals and Multigrade Relations' in *Noûs* 9(3): 309-318
- Nolan, D. (2014) 'Balls and All' in Kleinschmidt, S. (ed.) *Mereology and Location* Oxford: OUP
- Noonan, H. (1991) 'Indeterminate Identity, Contingent Identity, and Abelardian Predicates' in *The Philosophical Quarterly*, 41(163): 183-193
- Noonan, H. (2013) 'Moderate Monism, Sortal Concepts, and Relative Identity' in *The Monist* 96(1): 101-130
- Noonan, H., and Curtis, B. (2018) 'Identity' in *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2018/entries/identity/> Accessed 3.3.2021
- O'Leary-Hawthorne, J. and Cortens, A. (1995) 'Towards Ontological Nihilism' in *Philosophical Studies* 79(2): 143-165
- Oliver, A. and Smiley, T. (2013) *Plural Logic* Oxford: OUP
- Page, S. (2006) 'Mind-Independence Disambiguated: Separating the Meat from the Straw in the Realism/Anti-Realism Debate' in *Ratio* 19(3): 321-335
- Parsons, J. (2014) 'The Many Primitives of Mereology' in Kleinschmidt, S. (ed.) *Mereology and Location* Oxford: OUP pp. 3-12
- Papineau, D. (1979) *Theory and Meaning* Oxford: Clarendon Press
- Potochnik, A. (2010) 'Levels of Explanation Reconceived' in *Philosophy of Science* 77: 59-72
- Putnam, H. (1975) 'The Meaning of Meaning' in *Minnesota Studies in the Philosophy of Science* 7: 131-193
- Quine, W. V. O. (1953) *From a Logical Point of View* Cambridge, MA: Harvard University Press
- Quine, W. V. O. (1976) 'Worlds Away' in *Journal of Philosophy* 73(22): 859-863
- Ramsey, F. ([1929] 1978) 'Theories' in F. Ramsey & D.H. Mellor *Foundations: Essays in Philosophy, Logic, Mathematics and Economics* London: Routledge pp. 101-125
- Rayo, A. (2006) 'Beyond Plurals' in A. Rayo and G. Uzquiano (eds) *Absolute Generality* OUP pp. 220-254
- Rea, M. (1998) 'In Defence of Mereological Universalism' in *Philosophy and Phenomenological Research* 58(2): 347-360

- Richard, M. (2006) 'Context, Vagueness, and Ontology' in Greenough, P. and Lynch, M. (eds) *Truth and Realism* Oxford: OUP pp. 162-174
- Rose, D. & Schaffer, J. (2017) 'Folk Mereology is Teleological' in *Noûs* 51(2): 238-270
- Rosen, G. and Dorr, C. (2002) 'Composition as a Fiction' in R. Gale (ed.) *The Blackwell Guide to Metaphysics* Oxford: Blackwell pp. 151-174
- Rumfitt, I. (2005) 'Plural Terms: Another Variety of Reference?' in Bermúdez, J. (ed.) *Thought, Reference, and Experience: Themes from the Philosophy of Gareth Evans* Oxford: OUP pp. 84-123
- Salmon, N. (1998) 'Nonexistence' in *Noûs* 32(3): 277-319
- Schaffer, J. (2007) 'From Nihilism to Monism' in *Australian Journal of Philosophy* 85(2): 175-191
- Schaffer, J. (2009) 'On What Grounds What' in Chalmers, D., Manley, M., and Wasserman, R. (eds) *Metametaphysics* Oxford: OUP 347-383
- Schaffer, J. (2010a) 'Monism: The Priority of the Whole' in *The Philosophical Review* 119(1): 31-76
- Schaffer, J. (2010b) 'The Least Discerning and Most Promiscuous Truthmaker' in *The Philosophical Quarterly* 60(239): 307-324
- Schaffer, J. (2012) 'Why The World Has Parts: Reply to Horgan and Potrc' in Goff, P. (ed.) *Spinoza on Monism*, Palgrave, pp. 77-91
- Schaffer, J. (2015) 'What Not to Multiply Without Necessity' in *Australasian Journal of Philosophy* 94(4): 644-664
- Schaffer, J. (2016) 'Grounding in the Image of Causation' in *Philosophical Studies* 173: 49-100
- Schaffer, J. (2017) 'Laws for Metaphysical Explanation' in *Philosophical Issues* 27(1): 302-321
- Sider, T. (1993) 'Van Inwagen and the Possibility of Gunk' *Analysis* 53: 286-289.
- Sider, T. (1996) 'All the World's a Stage' in *Australasian Journal of Philosophy* 74: 433-453
- Sider, T. (2001) *Four-Dimensionalism* Oxford: OUP
- Sider, T. (2007) 'Parthood' in *The Philosophical Review* 116(1): 51-91
- Sider, T. (2009) 'Ontological Realism' in Chalmers, D., Manley, M., and Wasserman, R. (eds) *Metametaphysics* Oxford: OUP pp. 384-423
- Sider, T. (2011) *Writing the Book of the World* Oxford: OUP
- Sider, T. (2013a) 'Against Parthood' in D. Zimmerman and K. Bennett (eds) *Oxford Studies in Metaphysics vol. 8* Oxford: OUP pp. 237-293
- Sider, T. (2013b) 'Reply to Jonathan Schaffer' in *Analysis* 73(4): 760-770
- Sider, T. (2015) 'Nothing Over and Above' in *Grazer Philosophische Studien* 91: 191-216
- Simons, P. (1987) *Parts: A Study in Ontology* Oxford: OUP

- Stanford, K. and Kitcher, P. (2000) 'Refining the Causal Theory of Reference for Natural Kind Terms' in *Philosophical Studies* 97(1): 99-129
- Stanley, J. (2001) 'Hermeneutic Fictionalism' in *Midwest Studies in Philosophy* 25: 36-71
- Sterelny, K. (1983) 'Natural Kind Terms' in *Pacific Philosophical Quarterly* 64(2): 110-125
- Strawson, P. (1959) *Individuals* London: Methuen.
- Swinburne, R. (1995) 'Thisness' in *Australasian Journal of Philosophy* 73(3): 389-400
- Tallant, J. (2011) *Metaphysics: An Introduction* 1st edition Continuum
- Tallant, J. (2013) 'Qualitative Parsimony and the Metaphysics of Time: Motivating Presentism' in *Philosophy and Phenomenological Research* 87(3): 688-705
- Tallant, J. (2014) 'Against Mereological Nihilism' in *Synthese* 191: 1511-1527
- Thomasson, A. (2007) *Ordinary Objects* Oxford: OUP
- Thomasson, A. (2008) 'Existence Questions' in *Philosophical Studies* 141: 63-78
- Thomasson, A. (2009) 'Answerable and Unanswerable Questions' in D. Chalmers, D. Manley, & R. Wasserman (eds) *Metametaphysics* Oxford: OUP pp. 444-471
- Thomasson, A. (2015) *Ontology Made Easy* Oxford: OUP
- Thunder, S. (2017) 'Mereological Nihilism: Keeping it Simple' in *Thought* 6(4): 278-287
- Thunder, S. (forthcoming) 'There is No Reason to Replace the Razor with the Laser' in *Synthese*
- Tooley, M. (1987) *Causation: A Realist Approach* Oxford: OUP
- Unger, P. (1980) 'The Problem of the Many' in *Midwest Studies in Philosophy* 5: 411-67
- Uzquiano, G. (2011) 'Plural Quantification and Modality' in *Proceedings of the Aristotelian Society* 111: 219-250
- Van Cleve, J. (2008) 'The Moon and the Sixpence: a Defense of Mereological Universalism' in Sider, T., Hawthorne, J., and Zimmerman, D. (eds) *Contemporary Debates in Metaphysics* Oxford: Blackwell pp. 321-340
- Van Cleve, J. (1985) 'Three Versions of the Bundle Theory' in *Philosophical Studies* 47(1): 95-107
- Van Inwagen, P. (1981) 'The Doctrine of Arbitrary Undetached Parts' in *Pacific Philosophical Quarterly* 62(2): 123-137
- Van Inwagen, P. (1987) 'When Are Objects Parts?' in *Philosophical Perspectives* 1: 21-47
- Van Inwagen, P. (1990a) *Material Beings* Ithaca: Cornell University press
- Van Inwagen, P. (1990b) 'Four-Dimensional Objects' in *Noûs* 24: 245-255.
- Van Inwagen, P. (1994) 'Composition as Identity' in *Philosophical Perspectives* 8: 207-220.

- Van Inwagen, P. (1998) 'Meta-Ontology' in *Erkenntnis* 48: 233-250
- Varzi, A. (2001) 'Parts, Counterparts, and Modal Occurents' in *Travaux de logique* 14: 151–171.
- Varzi, A. (2019) 'Mereology' in *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2019/entries/mereology/>> Accessed 3.3.2021
- Wallace, M. (2011a) 'Composition as Identity: Part 1' in *Philosophy Compass* 6(11): 804-816
- Wallace, M. (2011b) 'Composition as Identity: Part 2' in *Philosophy Compass* 6(11): 817-827
- Wallace, M. (2014a) 'Composition as Identity, Mereological Essentialism, and Modal Parts' in A. Cotnoir and D. Baxter (eds) *Composition as Identity* Oxford: OUP 111-129
- Wallace, M. (2014b) 'The Argument from Vagueness for Modal Parts' in *Dialectica* 68(3): 355-373
- Wallace, M. (2019) 'The Lump Sum: A Theory of Modal Parts' in *Philosophical Papers* 48(3): 403-435
- Wasserman, R. (2018) 'Material Constitution' in *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2018/entries/material-constitution/>> Accessed 3.3.2021.
- Weatherson, B. [n.d] 'Stages, Worms, Slices and Lumps' URL = <<http://brian.weatherson.org/swsl.pdf>> Accessed 3.3.2021
- Westerhoff, J. (2005) *Ontological Categories: Their Nature and Significance* Oxford: OUP
- Wiggins, D. (2001) *Sameness and Substance Renewed* Cambridge: Cambridge University Press.
- Williams, R. (2006) 'Illusions of Gunk' in *Philosophical Perspectives* 20: 493-513
- Wilson, J. (2014) 'No Work for a Theory of Grounding' in *Inquiry* 57: 535-79.
- Yi, B-U. (1999) 'Is Mereology Ontologically Innocent?' in *Philosophical Studies* 93(2): 141-160
- Yi, B-U. (2005) 'The Logic and Meaning of Plurals Part I' in *Journal of Philosophical Logic* 34: 459-506
- Yi, B-U. (2006) 'The Logic and Meaning of Plurals Part II' in *Journal of Philosophical Logic* 35: 239-288
- Zimmerman, D. (1996) 'Could Extended Objects Be Made Out of Simple Parts' in *Philosophy and Phenomenological Research* 56(1): 1-29
- Zimmerman, D. (2008) 'The Privileged Present: Defending an A-Theory of Time' in Sider, T., Hawthorne, J., and Zimmerman, D. (eds) *Contemporary Debates in Metaphysics* Oxford: Blackwell pp. 211-225