

**METAONTOLOGY,
EMERGENCE AND THEORY
CHOICE:

IN DEFENCE OF
MEREOLGICAL NIHILISM.**

Richard Laurence John Caves

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

The University of Leeds

School of Philosophy, Religion and the History of Science

September 2014

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

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

© 2014 The University of Leeds and Richard Laurence John Caves

The right of Richard Laurence John Caves to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

Acknowledgements

This research was supported by the Arts and Humanities Research Council, whom I am grateful to for fully funding the first three years of my doctoral studies. (Initial preparation for my PhD was begun during my PhD-track M.A., which was fully funded by the University of Leeds Collection of Small Endowed Awards.) This research was also supported by many, many individual human beings, in one way or another, and I owe you all a huge debt of thanks:

Firstly, I must thank my (main) supervisors, Elizabeth Barnes and Ross Cameron, for your unwavering support and crucial guidance (including, but not limited to, insightful feedback, challenging supervisions, and occasional much-needed pep talks). You were my teachers long before you were my supervisors: my intellectual debt to both of you runs unfathomably deep, as does my gratitude for your generosity of time, patience, understanding, kindness and encouragement. You have been role models to me as philosophers, as teachers and as people.

Next I must also, of course, thank my *other supervisor* - John Divers - for keeping a weather eye on my progress: our annual meetings were always incredibly valuable and your support in the final months (with the absence of Ross and Elizabeth) indispensable.

I would like to thank all the other staff at Leeds, and many visiting academics, for helpful comments, support and advice, here and there, over the years— particularly regulars of the CMM seminar. If I attempt to name names exhaustively, I'll be setting myself up to fail: the number of people who made small but significant contributions is too great for my measly powers of recall to handle accurately. Nonetheless, Heather Logue must get a mention for supervising my M.A. dissertation and PhD proposal. Also, Scott Shalkowski and Jason Turner for helpful discussion and constructive criticism at my transfer meeting. Robbie Williams, too, for little things with a sum greater than the parts. My conversations with Ted Sider and Agustín Rayo (primarily during small graduate discussion groups) have proved especially helpful with respect to a number of issues relevant to the thesis. Chronologically last, but by no means least, Christopher Daly - and Scott Shalkowski again - should be thanked for incisive and thought-provoking discussion in my viva, and for help with editing.

Invaluable for their enlightened conversation and moral support: the past and present postgraduates of Leeds - philosophers, philosophers of science and historians of science

alike. Again, too many names. However, for helpful discussions or comments, at one point or another, directly relevant to the thesis, I should mention: Sarah Adams, Jonathan Banks, Michael Bench-Capon, Thomas Brouwer, Kerry McKenzie, Robert Pezet and David Race. And probably many others, including visiting postgraduate speakers, and regular audience members of the Leeds PG Seminar who are not on the above list - my apologies as well as my thanks.

I should add to the above Jade Fletcher, for keeping me on my toes with respect to matters metametaphysical - may you have an easier time grappling with these issues in your doctoral studies than I did! Meanwhile, extra special thanks is due David Cornell for extensive discussion of our respective partially overlapping thesis topics: when we agreed, it was reassuring; when we disagreed, I think we both profited from the exchange. Also, Danielle Adams for last minute proof reading and much needed encouragement.

Extra-extra special thanks goes to Ali(s) Kay for always challenging me to think truer and more interesting things. While last, but determinately not least, extra-extra-*extra* special thanks is to be heaped upon Carl Warom for frequent in-depth discussions on many pertinent subjects - often bordering on shameless procrastination - throughout the past four years. I'm infinitely indebted to him for reading every single word in this thesis (many tokens more than once), providing invaluable constructive criticism and going above and beyond in attempting to keep me sane through the final weeks and days of the writing up process. Very soon, I hope to be able to repay this infinite debt in kind.

Beyond the PG Room, I would also like to thank Andrew Sellers for his advice and instruction in matters of fundamental physics (and if, as likely, I've slipped up on any physics related issues in these pages it is certainly in spite of, rather than because of, his best efforts to educate me.)

Thanks also to my friends outside the philosophy ecosystem, especially:

Marcos - my old and dear friend. The things I'm indebted to you for are too many to count (yet somehow you are still managing to keep a tab). But I guess here is the place to thank you for being the sounding board for the crazy metaphysical views I formed long before I was trained in academic philosophy.

Shanti - for keeping me sane with regular provision of delicious snacks and awesome classic T.V. shows. I set out on this project with the aim of defending your favoured view that everything is One: I apologise that the final results were a little more ambiguous, if not downright contrary.

And, of course, my family:

My parents, Helen and Laurence: this thesis would never have been written without you, for more than just the obvious reasons. Thanks for putting up with a child who would talk your ear off about the paradoxes of the divine attributes and the possibility of mereological gunk from a ridiculously young age: you have always been loving, patient and encouraging.

My Uncle Gareth: Our long walks and philosophical debates were a defining feature of my childhood. Thanks for always keeping my interest in the big questions alive. (To my Other Uncle: I haven't seen you around much lately, but I'm thrilled that you're reading!)

Thank you to my new in-laws, Dave and Edie, for believing in me in every way.

And a huge thank you to the rest of my wonderful family - my grandparents and my siblings (Jeanette, Beryl, Robin, Christiana, Amy and Alex) - for your love and your support. Here I should give special mention to my late grandfather, John, who always pushed me to pursue my dreams, even in his last days.

Finally, yet most importantly: all my deepest love and thanks goes to Zoe, my wife and my soul mate, for sharing in this great journey, along with all its uplifting moments and heavy burdens— I look forward to continuing our journey together, long into the future, in expectation of all the joyful and exciting things to come. S.S. Warburg puts it best, I think you'll agree: *When I think something is important, you think it's important too. We have good ideas... sometimes we have good luck and sometimes we don't... I never met anyone sillier than me until I met you... I don't know why I like you. But I do.*

(To the people I've undoubtedly missed: again, I can only offer my humble apologies along with my sincere gratitude.)¹

¹ In postscript, let me note that I'm also immensely indebted to several non-human persons:

Dr Tiny - my constant companion, who completed seven doggie PhDs in the time it took me to write this thesis. Without you by my side throughout this process, I would have spent many hundreds of hours alone.

Master Maestro - you weren't even born yet when I embarked on this journey, but you have brought so much joy and excitement (and much needed additional exercise) to my life this past year.

Marley the Ferret (a.k.a. "Bendy-Dog") - you are the most intrinsically hilarious creature on the face of the Earth and probably beyond. Thanks for keeping me amused.

Lilly, Morris & Cleo - One can never have too many paws about the place. Your paws are all, without exception, exceptional.

Dylan - Thank you for the long walks and trying to teach Maestro how to behave. Getting to see you is always a treat for both of us (usually in the literal sense for you).

Willow - my fourth and furriest supervisor. That I couldn't possibly say you are also my *favourite* supervisor does not mean that it is not, in fact, the case.

Abstract

This thesis offers detailed remarks on debates in metametaphysics, the question of ontological emergence and the merits of extra-empirical theory choice criteria, such as ontological and ideological parsimony, in the course of defending Mereological Nihilism (nihilism).

Part I is primarily a discussion of the metaontology of the composition debate and how nihilism begins to look like a more attractive theory once a certain metaontological framework is adopted. §0. provides an overview of the background assumptions of my thesis and outlines the dialectical strategy I will be employing. §1. Explains why a strategy is needed to reconcile revisionary ontological claims, such as nihilism, with common sense: I argue for the Ontologese Strategy (OS) over the more established neo-Quinean paraphrase strategy. §2. Elucidates OS and defends it from objections, while §3. discusses the status of non-fundamental existential claims on OS and responds to further objections. §4. Employs OS to defend nihilism from epistemic dismissivism, which threatens to undercut the *prima facie* theoretical advantages of nihilism— demonstrating the utility of OS in advancing the first order debate.

In *Part II* I defend nihilism from what I take to be the major challenge to the theory once certain metaontological assumptions are in place: the epistemic possibility of ontological emergence. §5. sets out the main line of defence – a strategy employing plural instantiation. §6. explores alternative strategies such as entanglement relations, fundamental indeterminacy and extended simples. §7. Looks at arguments for monism vs. pluralism in light of putative emergence from quantum mechanics and argues there is conceptual room for a third way – Local Holism. §8. Tries to bolster the assumption of this thesis that parsimony considerations are relevant to the composition debate. In §9. I make some concluding remarks for the thesis as a whole.

Table of Contents

Acknowledgements	i
Abstract	iv
Table of Contents	v
List of Figures.....	viii
PART I: MEREOLGY AND METAONTOLOGY	1
0. Introduction.....	3
0.1. The Question of Ontology	3
0.2. Why Care about Ontology?	8
0.3. Beyond Ontology	12
0.4. Mereological Nihilism	14
0.5. Puzzle Solving	17
0.6. Solutions, Solutions... ..	19
0.7. Redundancy and Simplicity.....	21
1. Revisionism and Reality.....	23
1.1. Ontology and Everyday Objects: the need for reconciliation	23
1.2. Common Sense and the Moorean Objection	26
1.3. The van Ormwagean paraphrase strategy.....	30
1.4. Semanticism: merely verbal disputes	33
1.5. The Errors of Error Theory.....	38
1.6. Metaphysical Perspicuity and Ontologese.....	45
2. Speaking the Language of the Ontology Room.....	50
2.1. Ontologese	50
2.2. A Common Sense of Existence	58
2.3. Can Earthlings Conceive of Electric Cats?	62
2.4. Hofweber – more unintelligibility	65
2.5. External World Scepticism (more Moore)	67
3. Escaping the Fundament.....	73
3.1. Dependence, Fundamentality and Reality.....	73
3.2. Unreal and Imperfectly Natural Metaphysics (of A-Level Chemistry).....	79
3.3. Social Ontology.....	85

4.	Dismissing Dismissivism— the nihilist triumphant!	93
4.1.	Epistemic Dismissivism	93
4.2.	Competing Simplicity Considerations	96
4.3.	Solve a Problem, Make a Problem	98
4.4.	A Nihilist Response: Parsimony	100
4.5.	A Nihilist Response: Puzzle solving	103
4.6.	Causal Exclusion: A case study	105
4.7.	End of Part One	110
	PART II: NIHILISM AND EMERGENCE	113
5.	Emergence for Nihilists	115
5.1.	The Argument from Emergence	116
5.2.	Understanding Emergence	120
5.3.	The Plural Instantiation Strategy	123
5.4.	The Cost of Plural Instantiation	131
6.	Extended Simples, Entanglement Relations and Indeterminacy	138
6.1.	Emergent Simples	139
6.1.1.	Objection 1: Simples can't be extended	140
6.1.2.	Objection 2: Heterogeneous properties	141
6.1.3.	Objection 3: Why not just have composition?	143
6.2.	Entanglement Relations	144
6.2.1.	Objection 1: No smallest particles	146
6.2.2.	Objection 2: Losing the unity of properties	148
6.2.3.	Objection 3: Priority Monism & the Impossibility of Submergence	150
6.3.	Discussion of Indeterminacy	156
6.3.1.	Objection 1: Against metaphysical indeterminacy	157
6.3.2.	Objection 2: The right tool for the job?	158
6.3.3.	Objection 3: Deep indeterminacy	159
6.4.	A Quick Comparison of the Four Strategies	163
7.	Priority and Pluralism	166
7.1.	Priority, Parthood and our Intuitions	167
7.2.	Dependence and Fundamentality	170
7.3.	Monism vs. Pluralism	171
7.4.	Strange New State Spaces	178
7.5.	Screen Worlds and the Actual World	183
7.6.	Merging (an intuitive gloss)	191

7.7. Local Holism and Ideological Economy.....	198
7.8. Sider-Lewis Objection and Merging.....	200
8. Nihilism and Parsimony.....	205
8.1. Kinds, Primitives, Theories and Cosmologies.....	206
8.2. Ideological Parsimony.....	208
8.3. Quantitative Parsimony	217
9. Some Simple Restatements Arranged Conclusion-wise.....	223
10. Appendix to Chapter 7: Collapse and Decoherence.....	226
Bibliography.....	234
Abbreviations	243

List of Figures

FIGURE I. Ferrets arranging some extended simples Cthulhu-wise*	2
FIGURE II. Some simples (in a chair configuration).	41
FIGURE III. Overdetermination	106
FIGURE IV. Many baseballs	107
FIGURE V. Micro-overdetermination	108
FIGURE VI. Simple breaking	109
FIGURE VII. Cthulhu emerges: Father of Ferrets; prior to all*	114
FIGURE VIII. A circle divided in half by a dotted line.....	167
FIGURE IX. Star Fighter sprite – 148 x155 pixels*	168
FIGURE X. Three example Screen World states.	173
FIGURE XI. The four chessboard states.	178
FIGURE XII. Example switch world states.	180
FIGURE XIII. Paired pixel states	180
FIGURE XIV. Variety states (X = disallowed)	182
FIGURE XV. Extended world and many simples	188
FIGURE XVI. Chessboard and heterogeneity fact.	189
FIGURE XVII. Building a house.....	193
FIGURE XVIII. Making a cake	194

(* = Original artwork by Marcos Avlonitis, reproduced with permission.)

*PART I: MEREOLGY AND
METAONTOLOGY*



Figure i. Ferrets arranging some extended simples Cthulhu-wise*

0. Introduction

Unsuspecting readers should be warned that what follows is a rather dense overview of the assumptions and dialectical strategy of this thesis, placed within the relevant literature as far as the author has deemed practicable. Those in search of a more discursive and accessible introduction should skip ahead to the next chapter, returning here when – and only when – they have worked up an insatiable hankering for recondite background materials. For those sticking with me, here is what you have in store for the next few pages:

§0.1 discusses what is to ask an ontological question, while §0.2 asks why we should care about answering them. Ultimately I argue that what we take ontological questions to be, at the meta-level, will have a direct impact on the plausibility of certain first order ontological positions. §0.3 goes beyond questions of ontology to explore how we should think about the *ideology* of our best theories. §0.4 introduces the subject of our investigation: mereological nihilism and gives a brief overview of how others have argued for the view previously and the strategy I will be taking to defend it. §0.5 introduces a number of puzzles which have plagued the literature on composite objects, while §0.6 argues that nihilism is attractive because it offers a straightforward, systematic dissolution of these puzzles. §0.7 concludes this introduction by suggesting that nihilism is also attractive because it offers us a metaphysics of concrete objects with a simpler ontology and ideology.

0.1. *The Question of Ontology*

We are embarking on an ontological investigation: Naively, this might be characterised as an investigation into what exists. As we'll see, this naive characterisation carries over – qualified but more or less intact – to several nuanced and considered theories of what ontological inquiry consists in. While not without prominent rivals,¹ these existence-based

¹ For example, Kit Fine (2001) and (2009), while arguing for a view that is hard to put definitively in one camp or the other, points along the way to both an explicitly existential (though non-quantificational) account and an explicitly non-existential account in the vicinity of the proposed view. Schaffer (2010) and (2009) offers us an explicitly non-existential account of ontological inquiry, on which the ontologist asks what things depend on what other things (see Sider 2011, §8). I shan't

or existential accounts dominate both ontological practice (how metaphysicians go about trying to find out answers to ontological questions) and the metaontology debate (our attempt to come up with a satisfactory reflective theory of what it is exactly we're doing when we 'do' ontology). Most influential of the existential accounts, in recent times, is the quantificational account, inspired by Quine: on this view, the question of what exists is, at least for most relevant purposes, to be identified (conflated?) with the question of *what there is*. To ask 'Do aliens exist?' is just to ask 'Is there life on other planets?', while wondering if quarks exist is to wonder if there are quarks. Ordinary language sentences of a theory containing these quantificational determiner phrases, such as *there is* and *there are*, are then regimented into existentially quantified sentences of first-order logic: What a theory says exists - what it is *ontologically committed* to - are the values of the variables bound by its quantifiers. To ask what exists, in the most general and unrestricted sense, is to ask what your best theory of the world says falls in the domain of the unrestricted existential quantifier.

Here I will be defending a particular ontological view concerning material objects: Mereological Nihilism (nihilism). To put nihilism in naive existential terms, it is the view that composite material objects do not exist, or in ordinary language quantificational terms, that there are no composite material objects: according to the nihilist, then, everything there is - everything that exists - is mereologically simple, not composed of any smaller things. But metaontology matters in the formulation and evaluation of this claim: If the naive existential view is wrong, or can't be sufficiently refined, the above formulation of the view won't be tenable. Nonetheless, the existential view is useful for getting a conceptual grip on what the nihilist's position is supposed to be, even if it isn't the last word on the matter: as such, I shall often unreflectively express nihilism using just such language. However, it should be made clear from the outset that it is the *ontological* thesis of nihilism that I am wedded to, not any particular form of words used to cash out its meaning or implications: If it turned out that the most perspicuous way to ask - and answer - ontological questions is not in quantificational or existential terms, then given my intention to defend nihilism *qua* ontological thesis I should rephrase my characterisation of nihilism accordingly, retracting the irrelevant and misleading claims I made about the *existence* of composite objects. Yes, given certain background assumptions about ontology, I'm arguing that there are no composite objects— furthermore, regardless of how the cards fall in the metaontology debate, I think this characterisation of my view

attempt an exhaustive survey, as in the end I think more hangs on whether or not the proffered metaontologies cash out ontological questions in quantificational terms, rather than whether they accommodate the intuition that ontological claims are about what exists.

provides a good way to get an intuitive grip on what the distinctively ontological claims of nihilism are supposed to be. This is precisely because – for better or worse – we naively think about ontological debates in existential terms and the tendency to slip between existential and quantificational talk is a fairly natural one. Nonetheless, if when push comes to shove, it transpires that the assertion that ‘there are no composite objects’ does not adequately capture the ontological commitments of the nihilist, it should be unceremoniously ditched, not clung onto until the bitter end. Better, in that case, to put all my energies into the potentially troublesome task of figuring out how to say what I *really* want to say.

On the other hand, the plausibility of nihilism is not at all independent of how the metaontological issues pan out. In fact, I’ll be arguing pretty much from the get go that a lot hangs on what we take to be the correct account of ontological questions to be. As things stand, we can’t simply keep our heads down and proceed with first order theorising, trusting that whatever explication of ontological practice at the meta-level we settle on will allow us to proceed without reevaluating our assumptions, methods or conclusions. In the final analysis, we must be judged on the specific claims that we make in the course of first-order ontological inquiry, rather than on our good intentions of making claims that have a certain ontological *je ne sais quoi*. We must, therefore, hold out hope that progress in metaontology will bring better understanding of how to formulate and evaluate these first-order claims. Indeed, I will be arguing that, if we accept certain plausible modifications to the neo-Quinean view of ontology, radically ‘revisionary’² ontologies such as nihilism – those theories which make existential claims which significantly conflict with what we ordinarily say and believe exists – end up looking a lot more attractive than they otherwise would. That metaontology can shape first-order debates would seem to lend support to the slogan that metaontology is just more metaphysics.³ At the same time, also consistent with the slogan, we must be lead into thinking that metaontological questions can or must be independently settled before first-order theorising can get back on track— assuming it was ever on a track heading roughly in the right direction to begin with. However, ontological and metaontological questions cannot be neatly separated from one another and any attempt to do metaontology in isolation from first order theorising risks wandering into debates that are sterile and intractable. For this reason, I think it fruitful to consider the metaontological framework I am attracted to in the course of arguing for nihilism, a first-order theory which I also think has a lot going for it: My hope is that the

² See Williams (2012, §1.3).

³ Due to Sider (2011).

case for each will be bolstered by the discussion of the other.

Before moving on to consider *why* we should care about ontological questions, I should say a bit about what sort(s) of... er... *thing* things are: So, at the most abstract and general level of consideration, are things *just things* (and, still, what's one of those?), or are there lots of irreducibly different kinds of thing? Well, when we pre-theoretically talk of 'things', we usually take our cue in the first instance from ordinary experience: as I've said, our intuitive world of things is one of medium-sized everyday objects such as teapots and lampposts and hamsters (other concreta are available).⁴ These things are concrete individuals – non-abstract, particular entities. These individuals of our experience have various qualities: being red or heavy or squishy or quiet. They also show up at particular places and times and have a habit of persisting and changing their qualities. We notice certain relations between them, such as some lampposts being taller than other ones, or some hamsters being the mothers of other hamsters, while some teapots *contain* hamsters (but most don't).

However, we may soon find ourselves quantifying over the properties and relations that we ascribe to the individuals in our ontology: We might say that this shade of red is deeper than another, or that all meters are shorter than a mile. We might start talking of the *state of affairs* of our hamster being hyperactive, or our teapot being only marginally above ambient temperature (so much for the worse for the tea). What about the *event* of the tea pot smashing? We might notice that the *number* of lampposts on the street now is one less than the *number* of lampposts on the street before the incident with the steam roller and the chimpanzee yesterday— do we have to believe there are numbers now too?

There are broadly three reactions you can have, once you've reached self-awareness about your propensity to talk about such entities: acceptance, paraphrase and weaseling. The Platonist about properties, for instance, accepts that such entities exist in addition to concrete individuals and treats them as abstract universals.⁵ The theorist who accepts a number of different ultimate kinds of entities beyond concrete individuals might take the primary goal of metaphysics to be to enumerate these basic *sui generis* categories— while those who are not so accepting might push for an alternative conception.⁶ The paraphraser, meanwhile, eschews (some of) these additional *sui generis* kinds and seeks to

⁴ As Quine (1960*a*) tells us, we all have to start somewhere.

⁵ Or multiply located concreta if you're Armstrong (1997).

⁶ Though L.A. Paul, e.g. (2012*a*), (2012*b*), is a prominent defender of both a one category ontology and a conception of metaphysics as centrally concerned with enumerating categories.

amend her way of talking such that she no longer has to mention them at all (this isn't usually an easy task⁷). Finally, the weasel refuses to believe in such entities, but goes on talking just as she always did— usually, she will provide a least a minimal gloss on why this OK.⁸ The nominalist about universals, for instance, might either analyse away such universals in terms of sets of individuals, or primitive similarity relations between individuals or particular concrete qualities, or she might go on talking as if individuals instantiate universals, but ultimately attribute patterns of property instantiation to the nature of individuals, without *analysing* the former in terms of the latter.

In general in this thesis, when discussing anything apart from concrete individuals (the objects of our investigation), I take myself to be saying something compatible with any of the above positions: For instance, I take everything I say about properties or property instantiation to be compatible with both Platonism and nominalism about universals. I will often speak *as* a Platonist, not because I am one, but because such talk will be acceptable to both the Platonist (at face value) and the weasel (in her weasely way), while it will be constitutive of having a paraphrase theory that the paraphrase theorist will have developed the tools to translate what I say into terms that are acceptable to her. Now it *may* be that, particularly in the case of trying to carry over something said in Platonist language to a given paraphrase, that not everything I say will straightforwardly work for everyone whatever their total ontology: In which case, there will be interesting further work to be done to see whether what I say for, e.g., the Platonist will work for a particular breed of nominalist. But if I'm sure of anything (more so even of my having hands which I am using to type this sentence) it's that you can't please everyone all of the time: even if - in the worst case scenario - these results are only useable for Platonists, that doesn't prevent them being of interest to non-Platonists. But I doubt the situation is this extreme: I try not to rest anything crucial on points where I perceive there to be irreconcilable differences between Platonist and nominalist frameworks. The same goes for any talk of facts, sets, numbers and so forth. In my heart of hearts I countenance very few of these things, but here my primary concern is making the argument for simplifying one's ontology of concrete individuals— insisting on deeper austerity from the outset would be dialectically damaging.

Note, however, that once one has all these additional categories of entities, one can pull up the ladder and rid oneself of concrete individuals altogether, by analysing them in

⁷ See, e.g., Melia (2000).

⁸ Weaseling is defended at some length in Melia (2000).

terms of entities from other categories. For instance, the bundle theorist tells us that individuals are just bundles of universals or particular properties that are compresent or held together by mereology or some *sui generis* relation.⁹ Again, I think everything I say can quite straightforwardly be understood by the bundle theorist: Nihilism can simply be stated as the view that no compresent compose further compresent bundles, or have other compresent bundles as parts. To the extent that not treating individuals as basic has the potential to complicate the dialectic down the line – for instance, if on a particular bundle theory the bundling relation is merereological, this doesn't play nicely with my claim that the ideology of mereology is dispensable – I simply set such issues aside as further work for another time. While I intend my arguments to be as neutral as possible concerning orthogonal metaphysical issues, space limitations prevent me from tracking down every loose end.

0.2. Why Care about Ontology?

But why care about ontological questions? Why does it matter what the ontological commitments of our theories are? Whether such commitments include composite objects or not? These aren't easy questions. The straightforward answer is that if we want to know what the world is like,¹⁰ and part of that requires what things are in it. Once we have a good idea about what sorts of concrete individuals exist – how many things of such and such qualities exist – and what the spatiotemporal relations and other relationships between them are, we have a pretty good idea of what the world of our experience is like. The ontological realist believes that what objects there are and what they are like is an objective, mind-independent matter – that our talk, thought and beliefs about objects are reflected in external reality. Indeed, that these objects *are* external reality.

So, is this picture sustainable? Leaving aside the thoroughgoing anti-realist who is sceptical about the idea of a real, mind-independent external reality in general, it is open for the opponent of ontological realism to question whether this is the only way – the best way, the *right* way – to describe said mind-independent reality. Do we really need to believe in a world of things? Well, it's unlikely that we're going to gain much traction on

⁹ See Paul (2012a).

¹⁰ I do keep talk about 'the world' don't I? I definitely don't mean to presuppose that there is a thing, the world, that is a concrete individual or any other kind of *thing*. (Though in *Part III* do discuss views on which the world is the *only* concrete individual.) It's just a handy way of speaking about reality or whatever you want to call it – whatever's going on beyond my own inner Cartesian theatre (no, there probably aren't any Cartesian theatres either). Call me a *world weasel* if you will, but I think I could get on fine without the world talk, I'm just rather attached to the various turns of phrase it facilitates.

such a high-level debate – certainly now is not the place to embark on a full-scale defence of ontological realism. However, I’m going to argue that we have *pro tanto* reason to be attracted to the ontological realist position. The gist of the argument is:

- 1) If we’re going to think or speak about reality then we need to use language as a vehicle for this thought and talk.
- 2) That language should allow us to make powerful, *systematic* statements about the world, since in theorising about anything we should strive for at least a basic degree of systematicity.
- 3) It’s very hard (impossible?) to speak a language that both eschews ontological commitment entirely (at least on a broadly quantificational model of ontology) and retains a basic degree of systematicity.
- 4) So, if we want to theorise about the nature of reality at all, we will end up incurring at least some degree of ontological commitment.¹¹

Our talk has a certain structure:¹² We combine determiner phrases – naming phrases such as ‘the cat’ or quantifiers phrases such as ‘some felines’ – with predicate phrases, which describe how the things picked out by the determiner phrases are. This is what Turner¹³ describes as a pegboard and rubber-band structure – we can imagine hanging rubber bands on pegs (representing the instantiation of monadic properties) or stretching them between pegs (relations), with the pegs themselves representing particular objects and the overall pattern created by the hanging of rubber bands, on and across various pegs, as representing the ontological structure of the world that we try to uncover through investigation.¹⁴ For those tempted to make a decisive break from an ontological picture of reality, a pressing question is whether we could – in principle – speak a more perspicuous language that did away with any linguistic machinery that could reasonably be

¹¹ At least by broadly Quinean standards – there’s room to quibble about how this fits with radically different conceptions of ontology.

¹² Importantly, I here mean structure in a fairly literal sense – i.e. pertaining to configuration, organisation, the relationship between components – and this is meant to carry over to my talking of structure in reality or ontology. I do not take myself to be committed – certainly at this point in the dialectic – to the fully fledged metaontological view of structure argued for in Sider (2011).

¹³ Turner (MS.), (2010), (2011).

¹⁴ I take it that all of this can be cast in a nominalist spirit, as per my remarks in the previous subsection. In the final analysis your average nominalists will go to great lengths to give the subject-predicate structure of our talk an ontological basis: most likely an object-property account in their preferred terms.

taken to correspond with an object/property structure in the world (i.e. a much more radical proposal than simply analysing individuals in terms of bundles, or analysing properties in terms of objective resemblance between individuals – for these picture still closely reflect the subject/predicate structure of our language).

Turner¹⁵ shows us that while we might want to be deny the existence of certain things – e.g. abstracta, holes and, yes, composite objects – stopping talking about things altogether is (un)surprisingly hard if you still want to be able to describe facts about the world in a systematic way. Take two of the languages considered: The first, a feature placing language, makes only statements of the form ‘it’s raining’ which indicate the satisfaction of a 0-placed predicate without picking out any particular object. So, instead of saying ‘here is a ferret’ we might say ‘it’s ferreting’. For any state of affairs, no matter how complex, we can describe it using a sentence of this sort. But say we have a sentence which stands in for a complex sentence such as ‘there are two ferrets on the beret’, we can’t move from this sentence to our sentence that stands in for there being *some* ferrets on the beret— for these sentences are not truth-functional compounds predicating things about ferrets but correspond to simple propositions. So logic can’t get us from it two-ferret-on-the-beret-ing to the fact that it is ferreting or beret-ing or some-ferret-on-the-beret-ing.

There is a worrying lack of inferential systematicity here. The ontological nihilist might try to rectify this deficit by coming up with a fancier language: Turner considers Quine’s functorese.¹⁶ But he argues that however this language appears to avoid quantification over objects and is as inferentially systematic as thing-talk, it in fact contains an expression which is exactly equivalent to the existential quantifier— so it fails to avoid ontologically committing talk after all. That is, it’s not just that *sentences* of functorese can be paraphrased into *sentences* of quantificational language: But on a subsentential level functorese contains a bit of vocabulary playing a role exactly equivalent to that of a quantifier. Speaking in a non-ontologically committing way, then, is an honour that looks like it perhaps only belongs to languages which are patently inadequate for systematically describing the world. So if one wants to make any attempt at speaking in a way that captures the structure of the world, you are going to incur ontological commitments in one form or another. Yes, there’s plenty of ways one could go about trying to resist this conclusion: denying the need for inferential systematicity; holding out hope of finding an

¹⁵ Turner (2011). See also Sider (2001, introduction); Sider (2011).

¹⁶ See Quine (1960*b*); (1971).

adequately systematic language that does not incur ontological commitment; and so on. But we do have a firm suggestion for why interest in ontological commitment is a healthy starting point for those with ambitions to discover and describe the structure of the world.

There are various ways you might go about modifying the above picture of ontological structure which are still in the same spirit: for instance, one version of ontological pluralism¹⁷ allows for there being multiple, mutually incompatible ways to exist (concretely or abstractly, say) and that existing in one way precludes you from having properties had by things that exist in another way.¹⁸ So: separate peg board, without rubber bands being stretched between pegs of different boards – pegs on different boards representing ways for objects to exist, in contrast to the single pegboard model that assumes all pegs are the same (exist in the same way) only differentiated by the rubber bands that hang on them (the different properties they have). Provided that the questions such as how many pegboards there are, how many pegs each has and the arrangement of bands on each board remain non-trivial questions then it would seem that what we're presented with here is just a more elaborate version of the original pegboard and rubber-band model: the basic idea of uncovering ontological structure remains intact.

What about non-quantificational views of ontology? On most of these views, it tends to be assumed that the facts about what objects there are and what properties they have is somehow already a given prior to our embarking on ontological inquiry.¹⁹ In what way are they given to us? Well, perhaps they are given to us by science, or by common sense, or the senses – perhaps, on the one hand, answers to such questions turn out to be trivial or obvious or, on the other hand, perhaps the questions are hard and we have better epistemic routes to such information than those available to us when merely sitting in the metaphysics room. Either way, such views are interested in what things ontologically ground what— that is (to employ some further terms of art which form a tight conceptual circle), what things are ontologically prior to other things or metaphysically explain them.²⁰

¹⁷ Discussed in Turner (2010).

¹⁸ This last bit isn't an essential part of the ontological pluralist's position – there are interestingly different ways of formulating the view.

¹⁹ This view is stated explicitly in Schaffer (2010).

²⁰ For Jonathan Schaffer the relevant 'things' here are straightforwardly concrete individuals— Kit Fine's view is more complicated as it involves facts (which may concern concrete individuals) grounding other facts. What follows, for sake of simplicity, is a characterisation closer to Schaffer's view.

To continue the analogy, we may still take there to be important questions to be answered about the pegboard-rubber-band structure that must be investigated in the ontology room, even with the pegboard structure laid out before us. Perhaps, again, there are multiple pegboards, but we know that they do not all belong on the floor— we must stack them one on top the other. Then questions arise about which level certain pegs and their bands belong on, and which pegboards stand under and support the others and which pegboard belongs on the ground, supporting the entire structure. Perhaps things get more complicated than this: Perhaps the support pillars don't go from board to board, but from one local region of one peg board to a local region of another, or even from peg to peg— in this case, we have to answer more complicated questions not just about what level boards belong on in the structure, but which bits of the board beneath directly support which bits of the board above.²¹ When all is said and done, though, these popular conceptions of non-quantificational ontology can be seen to have their roots very firmly in a picture of reality which contains worldly structure corresponding to the subject/predicate structure of our language, with rival definitions of ontological commitment being parasitic on this— e.g. by saying that only the things at the top/bottom of the grounding hierarchy are ontological commitments.

Making quantificational-type claims would seem to be at the heart of what gives our theories the desired systematicity: as such, when examining the doctrine of a theory²² it makes sense to single out claims it makes of this form – using the label *ontological commitment* – and to pay special attention to them in evaluating the merits of the theory. In so far as we take our theory to be holding a mirror to reality, we should take each ontological commitment of our theory to be representing a genuine and important feature of the world. In the final analysis, even if it's hard to see how we could get by without making *any* ontological commitments, there are substantive questions about *which* objects we need to be ontologically committed to: this is what we investigate when we do ontology.

0.3. Beyond Ontology

Before the real fun begins – i.e. mereological nihilism, the subject of this investigation, gets a look in – we should address one last issue. Does the extent to which the world is

²¹ For a further look at non-quantificational metaontology see §1.

²² The doctrine of a theory is what a theory says. The ideology of a theory is the theoretical machinery used to say it. Cf. Sider (2011).

structured, in this sense introduced in the previous subsection, go *beyond* the subject-predicate structure?²³ On the standard Quinean view our theories should be constructed in sentences of the canonical language— usually first order logic. However, this choice of canonical language prejudices a lot of issues that would appear to be *worldly* questions rather than just pragmatic concerns about what suits us in writing down and communicating our theories.²⁴ The language in which we formulate our theories determines what it is possible for us to say about the nature of reality and the ontology that is required to say it.

One significant issue is what we say about natural properties. For any true predication one can make of an object, there is in some sense a corresponding property. If I can truly say of a ball that it is red, then there is a property of being red which the ball has. But many a metaphysician has wanted to say that all these true predications are really only true in virtue of a small handful of properties. Perhaps just the properties like charge and spin, say, mentioned in our theories of fundamental physics. Perhaps other properties besides, but not a property for every predication. How to distinguish metaphysically between these abundant properties, on the one hand, which are two are penny, and the sparse properties which play a special role in our theorising, having predictive power and so forth? We might take our cue from the ontological language I used above - *there are* properties - and follow David Lewis' influential suggestion²⁵ of positing entities, universals, to mark the difference between the sparse and abundant properties: there aren't any *things* out there in the world corresponding to the abundant properties, just true predications in virtue of how the sparse properties are. The sparse properties, meanwhile, correspond to universals— a predication concerning a sparse property is true when particulars *instantiate* the corresponding universal.

But we should question why such a distinction should have to be made ontologically. If you're a dyed in the wool nominalist who doesn't think there are any such things as properties, full stop - universal or particular; concrete or abstract - then you will want a way to single out the predicates that correspond to natural properties without positing corresponding universals. This can be achieved if certain predicates can be given the status of naturalness by appearing in the theories' primitive vocabulary, while other predicates are left out of the theories' primitive vocabulary - perhaps with a recipe for

²³ For detailed discussion on this topic see Sider (2011, §6).

²⁴ See Sider (2011).

²⁵ Lewis (1983*b*).

being defined from it. This strategy only works without additional ontology to back it up if we are realist about the ideology of the theory – if we take its primitives to be decided by the world rather than the pragmatic concerns of theory makers.

But the implications of ideological realism reach beyond the worldly status of predicates. Quine’s use of ‘posit’ to talk of our ontological commitments smacks a little of word play – suggestive both of the claims of our theory and the things out there in the world of which our theories are claiming.²⁶ Our ideological commitments should be thought of as spanning the divide in much the same way: If we wholly endorse a theory, the language needed to describe the world on that theory should not just be considered a feature of that theory but a feature of the world itself. We can all agree, for instance, that history is in the past, but the language our best theories use to describe there being past events gives us a clue as to whether irreducible pastness – tense – is a genuine feature of reality. It is telling, for instance, whether a theory has – or lacks – primitive tense operators. By treating the primitive vocabulary of our best theories as worldly, analogous to the worldly status we give the values of variables in our theories, we furnish ourselves with the resources to make genuine metaphysical claims about the world that would be hard or ontologically costly to make otherwise.

The primary focus of this investigation is ontology. As such, much of what I have to say in this thesis relies only on realism about ontology rather than realism about ideology. However, where ideological issues must be confronted head on in the course of ontological investigation, I adopt an avowedly realist position.²⁷ This is for the reasons above, and because I believe this approach will give most traction in the debates that follow.

0.4. Mereological Nihilism

Mereological Nihilism (nihilism) is one answer to van Inwagen’s Special Composition Question (SCQ).²⁸ Van Inwagen asks us: When we have some things, under what

²⁶ See Azzoumi (2004: 125-6)

²⁷ Again, perhaps we can stop short of a fully fledged Siderian position here, which includes even logical vocabulary in the category of ideological notions which are ‘baked into the world’. See Brouwer (MS.) for a full appraisal.

²⁸ The view has been defended by Peter Unger (1979), (1980), Ross Cameron (2010a), Cian Dorr (2005), Ted Sider (2013). Van Inwagen (1990) and Trenton Merricks (2001) are nihilists about inanimate composite objects. Horgan & Potrc (2000) and Cornell (MS.) believe there exists only one concrete object and so are trivially nihilists.

circumstance would they compose some further thing?²⁹ On the face of it, this is quite a straightforward, commonsensical question: When I have some bricks, what do I have to do to get them to compose some further thing, a house? If I have two slices of buttered toast and I press the buttered sides together, have they now composed an additional object— a toasty sandwich, perhaps? What does it take for some oxygen, some carbon, some hydrogen and some nitrogen, along with a few trace elements, to compose some further thing, such as a dog or a human being? But the straightforward, commonsensical answers to this straightforward, commonsensical question is shown by van Inwagen to be lacking: One representative answer, *contact*, is unworkable since it is implausible – or at least, certainly not commonsensical – that when two people shake hands they compose some further object.³⁰ What about if two objects are fastened together sufficiently tightly— well, if two people had an accident with some superglue would they form a new object?³¹ You get how this is going to go...

It turns out to be incredibly hard to find simple, systematic answers to the question that respect our ordinary beliefs. Indeed, the simple, systematic answers that have found currency among metaphysicians would appear to flatly contradict our ordinary intuitions about when things compose: The believer in unrestricted composition says that when you have some things, they always compose some further thing— that the world more or less obeys the axioms of classical mereology. Van Inwagen proposes a restricted theory of composition on which the only composite objects are living things.³² While, as we've said, the nihilist says there are no composite objects at all, not even plants and animals... not even people. Yet, on the whole, great effort has been put in, by those concerned, into trying to assure readers that such positions are not genuinely incompatible with commonsense, or at least the statements made by ordinary folk can be in some way reconciled with what the correct theory of composition is: the universalist might appeal to tacit quantifier restriction³³ while the nihilist and the restricted composition theorist might appeal to paraphrase strategies, which we will soon be examining more closely.³⁴

²⁹ See van Inwagen (1990, pp.20-31) for the canonical formulation of SCQ and the intuitive 'practical' version.

³⁰ Van Inwagen (1990, p.35).

³¹ Van Inwagen (1990, pp.57-8).

³² Or, for Merricks (2001), persons; conscious beings.

³³ See Lewis (1991: 80).

³⁴ See §1, §2, §4.

Those partial to ‘extreme’ answers to the special composition question— i.e. the non-restricted answers of universalism and nihilism, where things either always compose or never compose, respectively – often argue for their position by elimination of restricted positions. Once restricted composition is eliminated, the debate can be narrowed to assessing the relative merits of the two non-restricted positions.³⁵ Although – and plausibly because – there are a potentially infinite number of candidate restricted theories of composition, the restricted theories are usually rejected *en masse* via a general argument against the possibility of any restriction upon composition. The most well-known of these is the Sider-Lewis argument,³⁶ which argues that for any plausible restriction on composition it would sometimes be vague whether a given collection of things composed: i.e. if things composed only when they are ‘strongly fastened’ together, given that ‘strongly fastened’ is a vague term, it would be vague when condition was met. If there’s vagueness in whether things compose, then there’s vagueness in what exists, but Sider tells us that there cannot be vagueness in what exists— not unless there is vagueness in the world itself.

As it happens, however, I’m doubtful that the Sider-Lewis argument is successful in knocking restricted theories of composition off the table. For starters, I’m not sure the argument is sound, even granting its background assumptions.³⁷ But there’s one big assumption I think is probably wrong: for it is entirely plausible that there is vagueness in the world itself.³⁸ Even if admitting worldly indeterminacy into your theory of composition is an ideological cost, as I elaborate on in the second part of this thesis,³⁹ it is of such great theoretical utility in many other areas of metaphysics I don’t think it should be counted as a local cost against any particular theory that makes use of it: at worst it’s a global cost for one’s overall metaphysical system to absorb.

I don’t propose to go into any further detail in this section about either the Sider-Lewis argument or worldly indeterminacy, since nothing I say hangs on it. If you’re convinced by the Sider-Lewis argument then everything I say here works out the same or better for me, since you will be convinced of the falsity of swathes of rival theories before I’ve even tried to persuade you of the virtues of my own. I will be setting up my argument against

³⁵ Or the universalist might dismiss the nihilist’s position more or less out of hand (cf. Sider (2001a)).

³⁶ See Sider (2001a).

³⁷ See Nolan (2006).

³⁸ See Barnes & Williams (2012)

³⁹ Particularly in §6.3 but also in a few place elsewhere.

the believer in composition generally (the believer) and arguing on grounds of simplicity and its ability to solve puzzles concerning composite objects that there is good defeasible reason not to clutter our ontology with composite objects at all.⁴⁰

0.5. Puzzle Solving

Puzzles concerning number and identity conditions of composite objects underpin one powerful reason for adopting nihilism, since it offers a systematic way of dissolving such puzzles. There are many such puzzles. Some representative examples include:

Statue and the Clay -

In the morning I buy a lump of clay (Clegg) and make it into a statue (Sturgeon). By noon I'm bored of Sturgeon and unceremoniously squash the clay, destroying Sturgeon. Clegg presumably remains throughout the whole affair, but when I create Sturgeon I don't bring some further object into existence, occupying the same space and having the same categorical properties. There's just one object on the table: the same old Clegg, now identical with Sturgeon. And, yet, if Clegg and Sturgeon are identical in the morning, what are we to say at noon? If Sturgeon and Clegg are numerically identical then they should have the same persistence conditions, but somehow by squashing the clay I managed to destroy Sturgeon while Clegg presumably survives, as clayey and lumpy as ever.⁴¹

How to reconcile the competing data? Perhaps the most common way out⁴² is to accept the strange proposal that there are indeed two objects occupying the very same space: these are constitution views. Common sense tells us it's absurd for two material objects to be in the same place at the same time: I don't fall with gravity to the centre of the Earth because I can't occupy the same space as the solid rock beneath my feet— pick any two material objects and we would ordinarily expect the principle to hold. The constitution theorist will object that this is no common-or-garden case of co-location - or, rather, it *is* a common-or-garden case of co-location, not a bizarre situation like two co-located cats - because the two objects, Clegg and Sturgeon, are made from the same material and composed of the same parts. But here the mystery deepens: If the two objects have all the same parts and share all their qualitative properties, how on Earth did

⁴⁰ For the well-known argument to universalism via the Sider-Lewis argument see Sider's (2001*a*). The most systematic and persuasive elimination argument for nihilism I've encountered is given detailed presentation in Cornell (MS).

⁴¹ See, e.g., Thomson (1983); (1998).

⁴² Wasserman (2013).

they come to have different persistence conditions, modal properties and so on? If they're made out of exactly the same physical material, what are the *grounds* for that difference? Here the defender of the constitution view may try to distinguish the co-incident objects via relational facts, or through formal or abstract parts that they are taken not to share.⁴³ Here things get tricky: we won't explore these options any further.⁴⁴

Alternatively, perhaps Clegg really can be destroyed by being crafted into a statue and a new lump is created when the statue is destroyed. But it's hard to believe a lump of clay could be destroyed simply by forming it onto a statue. On the dominant sortal view, one thing can be both a statue and a lump of clay, but the persistence conditions of an object are determined by its dominant sortal, such that a mere lump of clay is destroyed by being shaped into a statue, while a lump that is also a statue is created and is subsequently destroyed when it ceases to be a statue. On this account, we must be prepared to accept the seemingly anthropocentric upshot that material objects can literally pop into or out of existence when they become or cease to be the sort of object that we are wont to pay special attention to— such as a piece of art.⁴⁵ Other potential ways out, which we won't go into here, include resort to relative or temporary identity.⁴⁶ Another way out which I'll make some brief remarks about in the next subsection is Stage Theory.⁴⁷

Ship of Theseus -

On the weekends Sarah builds a replica of the Ship of Theseus. On its completion, she decides to sale it to the Mediterranean. As the voyage continues, parts of the ship begin to wear out and so she replaces them at ports *en route*. Eventually, having replaced every plank in her ship, she begins to wonder whether it is even the same ship anymore. To make matters worse, unbeknownst to her, Sarah's arch rival in the Ancient Ship's Restoration society, Danielle, has been collecting up the planks at each port and has by now constructed her own Ship of Theseus replica out of all the cast of planks of Sarah's.

⁴³ E.g. Koslicki (2007, §7).

⁴⁴ Though see, e.g. Sider (2001*a*, §5.3).

⁴⁵ See Burke (1994*a*); (1994*b*) for a defence of the view. For discussion and objections see Sider (2001*a*, §5.4).

⁴⁶ See Sider (2001*a*, §5.5) and Wasserman (2013) for discussion. For the views see, e.g., Geach (1980) and Gallois (1998), respectively.

⁴⁷ Hawley (2001); Sider (2001*a*).

So which ship is the *original* Ship of Theseus Replica— Danielle’s or Sarah’s?⁴⁸

Here, arguably, is an even more worrying puzzle, of a slightly different sort:

Problem of the Many -

Observe a cloud from afar by looking up at the sky— it looks as if it has clearly defined boundaries. Now look at the same cloud from up close (get in an aeroplane and... take a magnifying glass?). You will see that it is composed of a huge number of tiny water droplets. As you zoom in on the ostensible boundary of the cloud, you see that instead of a sharp delineation between cloud and non-cloud, there is just increasingly diffuse collections of water droplets, with some outlying water droplets far from the cloud’s centre of mass.⁴⁹ Which drops are part of the cloud and which are not? Wherever we draw the line, there is likely some outlying water drop that we could have allowed was part of the cloud, or some water drop we could have omitted and still correctly called the remaining aggregate of water droplets a cloud. This isn’t just the familiar problem of having an object with a vague boundary: for surely the collection of water droplets with a water drop added, or the collection with one omitted, are also clouds? In fact, when we think about it carefully we realise that there are many different, overlapping, collections of water droplets in the vicinity that, each taken on its own, we would say composes a cloud. So, in fact, there are hundreds (thousand? millions?) of clouds present where ordinarily we would say there is just one! This might not be such a worrying problem if it only applied to clouds, but all macroscopic ordinary objects such as cats and people are made up of swarms of subatomic particles, and thus succumb to exactly the same difficulty.⁵⁰

0.6. Solutions, Solutions...

Due to considerations of space, and because I suspect I would be contributing little of any novelty to the debate, I do not propose any detailed examination of each puzzle. Nor, as a consequence, will I attempt to argue that nihilism presents the best solution to each puzzle, taken on its own. Plenty of ink has been spilt in an attempt to find satisfactory

⁴⁸ For some background on this very old puzzle see Wasserman (2013). See Parsons (2000) for a statement of the puzzle and potential solution employing indeterminacy of identity.

⁴⁹ See Unger (1980) for the original formulation of the puzzle and compare Geach (1980). See also Lewis (1993) for commentary. For an overview of the history of the debate see Weatherston (2009). For an interesting recent attempt to make progress on the puzzle without resorting to eliminativism see Jones (2012).

⁵⁰ See Parsons (2000), Lewis (1993); Geach (1980).

solutions. Suffice, for our purposes, to note that there is little consensus over which if any of the solutions is satisfactory and that each comes with its *prima facie* costs. For instance, as we've seen, the permitting of co-incident physical objects or the admission that where we would usually say there is one cat there are in fact thousands. The fact of the matter is, whether they can ultimately be resolved to the satisfaction of believers or not, belief in composite objects opens up a viper's nest of additional theoretical questions and puzzles, exposing one's theory to potential theoretical costs. What nihilism offers is a sweeping recipe for systematically dispelling all such puzzles, by eliminating composition from our theorising altogether. This is a strong *prima facie* reason to be attracted to nihilism; good reason to be interested in whether it is a viable theory, which might itself be adopted with little theoretical cost in place of a bewildering and problematic ontology of composites.

So, I won't be offering a complete argument for nihilism on the basis of these puzzles: the lacuna is that I won't be seeking to show that less extreme solutions than nihilism, piecemeal solutions to the individual puzzles, aren't viable— it might very well be that some and enough of them are. But in the following chapters what I will be seeking to show is that nihilism is a plausible theory that comes without any significant costs of its own. Further, contra what others have recently suggested in the literature – particularly Bennett (2011) whom I respond to at length in Chapter 4 – that it is fit to play the role of systematically resolving composite objects puzzles without creating new, equally perplexing puzzles. In that case, nihilism will have the virtue of offering a cheap, systematic solution to the sorts of puzzles described above. Though, as I've said, whether it is the *best* solution to each of the puzzles is left as an exercise for the reader.

One notable theory that also offers a systematic solution to many of the puzzles of composite objects is stage theory.⁵¹ In Sider's argument for stage theory it's telling that nihilism – eliminativism – is a view that constantly has to be set aside or dismissed, not being vulnerable to the more persuasive and systematic arguments on display against other rival positions.⁵² To avail ourselves of the *nihilism as systematic puzzle solver* motivation for nihilism, it does look as if I have to set stage theory to one side in a similar way. But rather than dismissing it as implausible, as is Sider's (2001) general approach to nihilism,⁵³ I take it that the two views can happily co-exist in dialectical space playing the same role: Nihilism is attractive if it can rid us of composite object puzzles independent of

⁵¹ Advocated by Sider (2001) and Hawley (2001).

⁵² And, indeed, that Sider went on to endorse nihilism/eliminativism anyway in his subsequent work.

⁵³ For stage theory is in fact a rather elegant and attractive view in many ways, by this author's lights.

any highly committal views about persistence and time possessed by stage theory,⁵⁴ and even if we do in fact hold such committal views on persistence and time, the overdetermination means that the advantages accrued by dissolving such puzzles are less hostage to fortune— if stage theory need be abandoned, there is an alternative way out in nihilism.

0.7. Redundancy and Simplicity

In addition to its puzzle solving ability, nihilism is attractive because it rids us of entities which are redundant for most – I suspect *all* – purposes. One major way in which composite objects appear to be redundant would be with respect to causal powers. Trenton Merricks⁵⁵ offers an argument for nihilism on the basis that if composite objects have any causal effects they must overdetermine the effects of their simple parts: so either composites are causally redundant by not possessing causal powers, or because they never cause anything that does not already have a sufficient cause in its microphysical base. This argument is discussed in more detail in Chapter 4, where I defend it from an objection by Karen Bennett.⁵⁶ As we'll also see in the coming chapters, simples give us a complete explanation of why our ordinary beliefs and observations concerning the macroscopic world are correct – so there's no additional work for composites to do here.

It's widely agreed that it's a bad making feature of a theory for it to contain *strictly redundant* entities.⁵⁷ However, even if composite object don't turn out to be strictly redundant, they may turn out to complicate our theory in other objectionable ways. Controversially, decreasing the different types⁵⁸ or (more controversially) total number⁵⁹ of entities posited by a theory compared to its rivals is a virtue of that theory— qualitative and quantitative ontological parsimony respectively. Recently, Sider has argued for nihilism on the grounds of another sort of parsimony: ideological parsimony.⁶⁰ Since the nihilist has no need to mention composite objects in her fundamental theory, she does

⁵⁴ So much the worse for stage theory if it's not quite so non-committal, to say the least, about the mereology of concrete objects!

⁵⁵ Merricks (2001).

⁵⁶ Bennett (2011).

⁵⁷ See §8.3.

⁵⁸ Lewis (1973, p.87)

⁵⁹ Nolan (1997)

⁶⁰ Sider (2013).

not need the primitive of parthood in her ideology. By reducing the number of ideological primitives in our theory, the thought goes, our theories become simpler and thus more explanatory.

I follow Sider in taking there to be an ideological parsimony advantage to nihilism. Which, if any, parsimony claims we should take seriously (and, if so, why) is a long standing area of contention. I don't claim to have any groundbreaking contribution to make to this debate, though I make some modest efforts to say something original about why parsimony considerations should apply to nihilism in the second half of this thesis.⁶¹ But the fact that many scientists and metaphysicians have thought that parsimony considerations are relevant to the debate again provides us with defeasible motivation to be nihilists in so far as you think parsimony is a virtue of theories and that nihilism proffers potential parsimony advantages.

Without further ado, then, I will proceed to outline a metaontological framework in which I think austere ontological positions such as mereological nihilism can flourish. In the next chapter we will see why nihilism does not conflict with our ordinary beliefs about the macroscopic world.

⁶¹ See esp. §8.

1. Revisionism and Reality

In the chapter I ask whether and how revisionary ontological claims – such as nihilism – can be taken seriously in the face of common sense and the evidence of our senses. In §1.1 I explain the need for a strategy to reconcile common sense with revisionary metaphysical claims. §1.2 examines arguments from Moore that such reconciliation is impossible and therefore revisionary metaphysics is impossible, while §1.3 looks at the now standard neo-Quinean paraphrase strategy – an error-theory due van Inwagen (1990) – for reconciliation. §1.4 raises a problem for the error-theory due to Hirsch and sets out the standard Siderean strategy for saving paraphrase from this unfriendly conception of revisionary metaphysics. In §1.5 I say what I find lacking in paraphrase strategies, while §1.6 puts forward some preferred suggestions for reconciliation.

1.1. *Ontology and Everyday Objects: the need for reconciliation*

Are there mongooses? Yes. Cabooses? Definitely. Are there weasels? Without doubt. Easels? Those too! I believe in the existence of all these things and more besides: statues, bumblebees, pulsars, deuterium molecules, atomic nuclei, human beings... you get the idea. If I went around in my day-to-day life speaking and acting as if none of these things existed I would be viewed by others as weirdly eccentric at best, crazy at worst. It's hard to imagine we could ever be given plausible reason for believing that none of the above things – or any great number of them – do not, in fact, exist. We might be tempted to say that it's just 'common sense' that such things exist, but this doesn't quite do our certainty justice if the body of common sense claims is taken to include dubious folk wisdom like going out in cold weather will cause you to catch a cold: Rather, in practice, this is somewhere up there with my certainty that I am in pain now, having just had my kneecaps smashed in with a baseball bat, and that going around randomly hitting people in the knees with baseball bats is morally wrong. If someone yells at me that there is a car coming, I'm not going to decide not to look around or dodge out of the way because of any of the entailments – such as that there are no composite objects such as car – of a metaphysical theory I happen to endorse.

Yet, like many before me, I also think that having a minimal ontology is a virtue and am, largely as a consequence, attracted to a particular metaphysical thesis concerning ontological status of concrete objects— *mereological nihilism* (henceforth just 'nihilism'). As I suggested in the Introduction and will be defending further in the coming pages,

mereological nihilism is a parsimonious theory that offers an easy unified solution to the many puzzles concerning the ontology of composite objects. Nihilism, according to the standard and widely accepted definition, is the view that there are no composite objects: that when you have some things they never compose some further thing; that all concrete particulars are mereologically simple. If the nihilist is right then mongooses, cabooses and everything else on the above list do not exist. At least given the overwhelmingly plausible assumption that weasels and the like are mereologically complex: For everyone who believes in the existence of weasels will say they have legs, ears, whiskers, tails and other proper parts (the believer in atomic nuclei that they are made up of protons and neutrons; while clay statues are composed of hydrous aluminium phyllosilicates, and so on). That everything on the preceding list of proper parts are all themselves prime candidates for being mereologically complex (being composed of cells, molecules, quarks...) is not really any comfort – it just goes to show how few of the things we ordinarily believe in the nihilist takes to exist.¹ What does the nihilist believe in, then? Just mereological simples, *whatever these are*. (Usually, simples are taken to be microscopic, perhaps point-sized entities, arranged in a variety of interesting ways across space(time)— or sometimes they are identified with points of space(time), their differing intrinsic properties producing a sort of mosaic.)

You could be forgiven for thinking (and some commentators give this general impression) that it would be better all-round if, in the service of wholesome and

¹ It would only be a problem for the nihilist if the world turned out to be made of atomless ‘gunk’ (Lewis 1991 : 20) – if every object were composed of smaller parts, such that there were no simples. For an argument from the possibility of gunk to the falsity of nihilism see, e.g., Sider (1993). My main reason for thinking that the nihilist need not worry about arguments from the possibility of gunk is detailed in Caves (*draft*). In a nutshell: the nihilist can explain away all *putative* cases of gunk, described in non-mereological terms, by appeal to expressive resources endorsed by both sides of the debate, while direct appeals to the possibility of gunk in mereological terms are dialectically inappropriate against the nihilist. A similar line of argument is found in Williams (2006). For an alternative defence, to which I’m also somewhat sympathetic – which argues from a deflationary view of modality – see Sider (2013). Thirdly, if you think answers to SCQ are contingent, as Cameron (2007) argues (again, I’m sympathetic), there will be no moving from the metaphysical possibility of gunk to the falsity of nihilism (see also Miller (2009) for remarks on contingentism in metaphysics more generally). For an argument, via the pessimistic metainduction, that there is no ‘fundamental level’ of ontology in the actual world see Schaffer (2003). McKenzie (MS.) shows Schaffer’s argument to be flawed, but offers her own argument by appeal to interpretational issues in Quantum Field Theory (QFT). It is beyond the scope of this investigation to evaluate the latter argument: I offer brief comment on what the nihilist should say about putatively hostile interpretations of QFT in §6.2.1. The overarching discussion in §7 is also of oblique relevance. While arguments from *actual* gunk sidestep the responses from Cameron and Sider, I suspect the responses in Caves (*draft*) and Williams (2006) might be adapted to meet the challenge from the epistemic – rather than metaphysical – possibility of gunk.

interesting debate, ontologists who endorse positions such as nihilism should not just deny the existence of composite objects but reject wholesale, in all contexts, the distinctive claims we would usually take to be entailed by there being composite objects— i.e. the folk all agree there's a wall here, and if there's a wall here I can't keep walking without hurting my face, but there are no such things as walls, so I can happily keep on walking without hurting my face. By trying to reconcile the non-existence of walls my inability to keep walking without hurting my face I'm at risk of being accused of making a subtle or esoteric claim (see §2.4)— or perhaps of making bold claims out of intellectual arrogance, or some such, then on realising its crazy entailments rapidly backpedalling to the point where it looks like I'm making no claim of any substance at all. This line leads to a sort of flatfooted dilemma for revisionary ontology: intelligible, intellectually honest and morally serious revisionary ontology is obviously false, while any attempt to reconcile its discourse with common sense is a delusional attempt to have one's cake and eat it (or even a dishonest attempt to eat one's cake and employ slight-of-hand to steal the cake of others). Contrary to this, however, I firmly believe that a reconciliation with ordinary existence claims is at the heart of understanding what we mean by ontological claims in the first place— it should not be seen as an attempt to weasel out of a stronger claim that the ontologist would really like to make but can't if her audience is to keep a straight face, but as part of explaining what claim she was trying to make all along. Even the idealist doesn't try to walk through walls or kick rocks with their bare feet, after all.

Indeed, we can all agree there's something very weird going in the interplay between ordinary belief and ontology that's not satisfactorily captured by the flatfooted rejection of ontology: yes, we're ordinarily confident that there's something very right about our ordinary discourse about composite objects, but the nihilist doesn't take herself to be saying something completely crazy when she earnestly puts forward her arguments for nihilism (and attributing genuine insanity or insincerity from the outset is hardly a charitable way to proceed). What would be nice here is some story that explains what's going on in both discourses - the ordinary person's everyday existential discourse and the revisionary ontologist's nihilist discourse - telling us to what extent they conflict and whether any reconciliation is possible.

This is a familiar story in metaphysics, not just in ontology: a radical claim is made which is, on the face of it, incompatible with our existing commitments, and where there are competing pressures to hold on to the common sense claim, on the one hand, and while accept the radical claim, on the other. The road travelled by such a metaphysician is a familiar one: Metaphysician becomes convinced of radical claim, is told it conflicts with exiting knowledge, then develops some sort of account to reconcile the new claim with the

putative existing knowledge. There are very many forms such a reconciliation strategy might take. In this chapter we're going to look at four such accounts: Mooreanism, van Ormwagean Error-Theory, Semanticism and, my preferred approach, the Ontologese Strategy (OS). Each of the four views differs with respect to at least one of the following questions which we should expect an account of the relationship between ordinary existential discourse and ontological discourse to answer:

- Q1. What is the alethic status of each discourse— is either of them *true*?
- Q2. Are the two discourses in *conflict*— does one being true ensure the other is (at least in some respect) false?
- Q3. Can we reconcile the two discourses in some way if there is conflict - show that there is at least *something* right or theoretically useful about the losing discourse that could be salvaged? If so, how?
- Q4. Is there any *substantive* issue at stake here - or once we understand the idiosyncratic statements of ontologists, do all the interesting issues dissolve?

First, then, we turn to Mooreanism.

1.2. Common Sense and the Moorean Objection

According to the Moorean:

- A1. Our ordinary discourse about the existence of objects is true, while revisionary ontological theses are just flat out false.
- A2. There *is* conflict: It's the truth of the ordinary claims that ensures the falsity of the ontological claims.
- A3. There can be no reconciliation since the ordinary speaker and the ontologist mean the very same thing by their conflicting statements, the ontologist is just very wrong about the facts.
- A4. There is something substantive at stake: If it turned out (contrary to fact) that there were no tables, that I have no hands, we would be in a radically sceptical scenario.

Mooreanism is therefore *hostile*, rather than *conciliatory* towards revisionary metaphysics, since it leaves open no possibility that the revisionary claims could be true (or even in some way correct or insightful, despite strictly speaking being false).

G. E. Moore defends the truth of a whole range of propositions on the grounds that

they are ‘Common Sense’:² propositions which he ‘know[s] with certainty to be true.’³ These include his own existence and the existence of the Earth and other material objects.⁴ The gist of the argument is that no one seriously doubts any of these propositions but that metaphysicians end up accidentally committing themselves to other propositions that are in fact in conflict with them: Thus, we should retract any subtle metaphysical claim once we realise it leads us astray in this matter, since we all agree that the propositions they conflict with are *obviously* true. As it’s often put, we should always consider it more likely that a given metaphysical argument is unsound, than that a common sense claim is mistaken.⁵

Whatever you make of the rest of the argument, it should be noted that the Moorean’s answer to Q2 – that ontological claims are straightforwardly incompatible with ordinary claims – goes more or less un-argued for. But this is a reasonable *pro tanto* assumption: Revisionary ontological claims sure *seem* incompatible with ordinary folk claims, which is what lead us down this road in the first place— they use, after all, the very expressions, the very same types of sentences, the folk would use to deny their ordinary existential claims if they were so inclined (or sufficiently intoxicated). Given a plausible story about how this is not in fact so, despite strong appearances, then we would be in a position to reevaluate to Moorean’s crucial assumption. I attempt to tell one such story later in the chapter, in the course of giving my preferred reconciliation strategy. However, the Moorean seems perfectly entitled to help herself to this assumption in the initial set-up of her view: Against conflicting views that share this incompatibility assumption (such as van Ormwagean’s error theory) the point will be moot; against views which take a compatibilist line, the debate will hang on the plausibility of the compatibilist’s alternative story about ontological claims.

Meanwhile, the assumption of substantivity, that something is genuinely at stake, is going to be something any *conciliatory* strategy is going to need too (the Moorean needs it for there to be any *prima facie* credibility to her claim that my not having hands would be surprising and worrying; while *conciliatory* strategies needs it if establishing the truth of revisionary ontological claims is going to be at all revisionary – that is, even if it does not actually lead us to the revision of our Moorean beliefs, it leads to the revision of our

² Moore (1925, p.42).

³ Moore (1925, p.32).

⁴ Moore (1925, p.33).

⁵ See, e.g., Daly & Liggins (*draft*).

overall theory in some interesting way.)

What's causing the problem for revisionary ontology, given incompatibility with our Moorean beliefs, is (of course) the assumed truth of our ordinary ontological claims. For this claim, at least, we have something of a more fleshed out argument from the Moorean. However, I want to offer some observations which cast doubt on the legitimacy of Moore's extreme confidence in the truth of (what he takes to be) common sense existential claims, especially relative to his credence in revisionary ontological claims. Firstly, as I've already hinted, I don't think it's a matter of us having some (infallible?) folk wisdom concerning what the true theory of composition is and nihilism coming into conflict with that. If we've done careful and methodical theorising why *shouldn't* we be happy to overturn our previous common sense convictions about what objects there are? Even if we have such a folk theory and nihilism is incompatible with it, there is just no *pro tanto* reason to think that our collective inherited prejudices carry significant evidential weight.⁶

Secondly, if my confidence in the chair I'm sitting in right now existing came primarily from some folk theory of composition, then on learning from van Inwagen⁷ that it's highly implausible that any theory of composition that respects our intuitions about composition could be true I should massively decrease my credence in the existence of said chair. But that is not how it works: Even if I came to think our folk theory of composition is false, I wouldn't stop believing I was sitting on a chair— or expect to suddenly fall on my behind like a cartoon character who was hanging in mid-air then plummeted to its death upon reading a book on Newtonian gravity. Thus, there is actually apparent *compatibility* between folk and ontological *beliefs*, even if linguistic considerations tell in favour of an incompatibilist reading of the competing claims. Even if we don't move straight from apparent compatibility to assuming genuine compatibility, just saving the phenomenon of apparent compatibility mitigates against Mooreanism: It strongly points either to falsity of the (*pro tanto* granted) assumption that the contents of the ordinary and ontological claims are incompatible (i.e. a story about genuine compatibility) or, assuming that believing a theory of composition entails believing its literal truth, to an error-theoretic conception of the veracity of folk beliefs (incompatibility, but with a strategy for conciliation.) Certainly, it puts paid to any notion that our belief in medium-sized objects is somehow driven by some 'folk theory' of composition that's

⁶ Cf. Van Inwagen (1990) and Sider (2013).

⁷ Van Inwagen (1990)

incompatible with revisionism – it suggests our belief in them is pre-theoretic in a strong sense (not just temporally prior to our first ever reading of *Material Beings*, or something like that).

Thirdly, everyday wisdom about the composition debate where it is evidenced seems as diverse in its opinions as it is among philosophers. Despite its apparent incompatibility with our ordinary, deeply-held beliefs about what concrete objects exist, nihilism is also *obviously* true: Or so say the great majority of my friends, family and acquaintances who are both reasonably scientifically literate and *not* specialists in any of the following fields: philosophy, physics or the philosophy of physics. Why do this group say that mereological nihilism is true?⁸ The justification they give for their conviction tends to be relatively straightforward: *Science (fundamental physics) tells us that all that exists are particles*. Yet, like me, this anecdotally defined group – call them the sciencey folk – don't go to the coffee shop overtly talking and acting as if tables, chairs and coffee cups don't exist. Moreover, in the next breath, they are happy to speak of science discovering that the objects they've just denied the existence of are made of molecules and atoms and that these molecules and atoms are made up of even smaller particles. And sometimes it even seems like the fact that science has discovered that these big things are made up of smaller things forms part of the *justification* for not believing in the bigger things. Is this all hopelessly confused? Whatever we make of the details, the spirit of the sciencey folks' position is closely aligned with that of the austere metaphysician who seeks to reconcile her attraction to a minimal ontology with our deeply held convictions about what concrete objects exist: Again, it points to a *conciliatory* conception of ontology, whether compatibilist or error-theoretic: It might even be that nihilism is the *intuitive view* among the folk!

While I concede I think Moorean convictions have a certain *prima facie* force to them, I would be far from the first (including defenders of common sense metaphysics like Eli Hirsch, whose rival *hostile* account of revisionary ontology – semanticism – we'll come to later on in the chapter) to conclude Moore's argument lacks the dialectical force to convince opponents.⁹ Luckily for his opponents, we can keep the intuitive thought – what

⁸ Well, as all its members are, as stated, my friends, family and acquaintances this should perhaps be explained away as them wanting to be nice about my work or my inadvertently contaminating this small and haphazard sample with my heterodox opinions – but, as suggested by my previous point, it's not as if we're trying to find some 'uncontaminated' sample of people with the *true folk intuition* about composition. All engagement with SCQ is all, in an important sense, *post-theoretic*.

⁹ Again, see Daly & Liggins (*draft*).

we can call the core ‘Moorean’ intuition – that folks talk and beliefs are either true or as-near-as-damn-it-true by adopting any of the other three options on the table. Meanwhile, we can jettison his suspect claim that reconciliation is impossible while taking our pick both with respect to the substantivity of ontological questions and the precise alethic status of folk belief (i.e. compatibilism vs. error-theory). Friends and foes of revisionary metaphysics – whether they share the core ‘Moorean’ intuition or not – should stand united against the Moorean.

1.3. *The van Ormwagean paraphrase strategy*

The neo-Quinean paraphrase strategy, as characterised in distinctive error-theoretic terms by van Inwagen (1990) – what we might call the van Ormwagean strategy – gives the following account of ontology and ordinary existential statements:

1. Our unrestricted quantification claims made in the service of serious ontological investigation can be true even if this means coming into apparent conflict with our everyday existential claims. Meanwhile, even the everyday existential claims we feel Moorean certainty about can sometimes be false.
2. Ordinary claims and ontology *conflict*: If the ontologist truly says there are no tables, it follows that our pre-theoretic belief in the existence of tables is false.
3. However, reconciliation between the folk and the ontologist is possible: If an everyday existence claim turns out false, we can none the less say that it is *correct* – which is enough to secure what we’re ordinarily sure about when we’re sure about an everyday existence claim – provided we can give a systematic paraphrase of the ordinary existence claims within our preferred theory (e.g. despite the truth of nihilism, I’m sure I won’t go hungry for breakfast because there are, at least, *simples arranged toaster-wise*.) Since we’re sure of at least the *correctness* of Moorean-style existence claims (such as my having hands), it thereby acts as a *constraint* on true ontological theories that they deliver us at least a systematic paraphrases of Moorean claims.
4. Ontological debates are taken to be substantive: it matters what the ontological commitments of our best theory are (and it is a mind/language-independent matter).

What is a metaphysician aiming to do by making an ontological claim? According to

the neo-Quinean, she is trying to capture our *ontological commitments*, where these are whatever our overall best theory has to quantify over.¹⁰ And what are we doing in everyday speech by making existential claims? Just the same thing as the metaphysician. Or, perhaps the best way of spinning it: It turns out that what we've just established is that what the metaphysician is doing is no different from what ordinary folk do with their existential claims. This is arguably, in essence, the view taken by Quine and has been highly influential: a broadly neo-Quinean approach has dominated metaphysics ever since, at least until the recent resurgence of interest in non-Quinean metaphysics.¹¹

So, Quine tells us the question of ontological commitment is to be 'merely asking what, according to that theory, there is'¹² and according to Quine we should assess the ontological commitments of a theory simply by looking at the things it quantifies over.¹³ Once we've regimented ordinary folks' existential claims into existentially quantified sentences of first order predicate logic, we're in a position to judge the ontological commitments of the things that ordinary people say about the world. The ontologist might think much of what ordinary people say is rubbish: if so, she's free to discard these committal sentences and come up with a different theory. However, in the final analysis, there's no fundamental difference between what the ordinary folk do when they make existential claims and what a metaphysician does when she makes ontological claims.

It's important to this way of cashing out the reconciliation story that the folk frequently go wrong, or at least in some way do worse in some systematic way, than the metaphysician according to this shared assessment criterion. Else there would be nothing for the metaphysician to do, not in this case because metaphysics is impossible, but just because it is so very easy. The most straightforward strategy here is to give some sort of error theory: our ordinary existential claims are flat out false, yet for the better ones we can see how and why they went wrong. By understanding this, we can get a better idea of why appropriate claims concerning composites, such as "the cellar door is shut!", impart useful information, such as the physical inaccessibility of the cellar's region of space, even though there are no such things as doors or any other composite objects.

¹⁰ Recall §0.1, 0.2.

¹¹ See Chalmers, et al. (2009).

¹² (Quine 1966, pp.203)

¹³ See §0.1.

The Neo-Quinean tradition, though arguably not Quine himself,¹⁴ provides us with such an error theory. On this strategy ordinary speakers' claims about composite objects are taken to be strictly false, but can be, nevertheless, 'correct under right circumstances.'¹⁵ For instance, if someone says 'there is a table' and there are simples 'arranged tablewise'¹⁶, the statement is correct, though strictly false; whereas if the same speaker were to say 'there are horned horses' her statement would be both false and incorrect, as there are no simples arranged in the requisite fashion. There may not be any parsnips, but the nihilist will still sit down for dinner, since she believes that simples arranged parsnip-wise are delicious. This belief makes the same qualitative demand on the world as her non-nihilist friend's belief that parsnips are delicious: it's just that they disagree about the number of things on their plate. While ordinary speakers are flat out wrong about what there is, we have a systematic way of connecting up their existential claims about composite objects with true existential statements about simples arranged in certain ways via the paraphrase scheme. This special class of paraphraseable statements gets some sort of 'alethic commendation'.¹⁷

We can follow Quine and say that this alethic commendation is in fact truth - but as Williams has pointed out this relies on his own idiosyncratic views on the philosophy of language.¹⁸ Quine's own view is the limiting case of a Quinean-inspired error theory: it can't properly be called an error theory as it recovers the truth of the folk claims, but in an important way which we'll see it is still very much in the spirit of an error theory since it claims that the purpose of existential claims is to deliver ontological commitments, but we learn that many of the folk sentences can't do this without regimentation or paraphrase—some true sentences are more equal than others. Similarly, Van Inwagen allows for some such sentences to be true, but not all, reserving correctness for the remainder.¹⁹ But the question of truth is less important (to the metaphysician *qua* metaphysician), I think, than the need to differentiate between three classes of existential claims: the utterly misguided, the ordinarily appropriate but in need of correction or paraphrase by the metaphysician and, finally, the sort which is in a good enough state for the metaphysician to read his ontological commitments off. By saying that paraphrasable claims are in fact true - a

¹⁴ See Williams (2012).

¹⁵ Or one might prefer some other 'term of alethic commendation' (van Inwagen 1990, p.102)

¹⁶ Van Inwagen (1990, p.111). See also Sider (draft).

¹⁷ Van Inwagen (1990).

¹⁸ Williams (2012).

¹⁹ Van Inwagen (1990).

position I will later endorse, but not for Quine's reasons - we're still left needing to distinguish those sentences which make ontological commitment perspicuous and those that do not.

In the next subsection, we'll look at a hostile challenge to this broad neo-Quinean picture from the semanticist. The core of the semanticist argument is that the neo-Quinean is wrong to think that what our ontological commitments are is a substantive matter: The charge is Carnapean in spirit - yes, our theory has certain language-relative ontological commitments, but we could have spoken a different language and ended up with a completely different set of commitments, and what language we speak is merely a conventional matter. It turns out, therefore, that the ontologist and the nihilist are having a *merely verbal* dispute.

I'll suggest - taking more or less Sider's line in the Sider-Hirsch dispute - that the neo-Quinean conception of ontology can be rescued from the jaws of neo-Carnapian semanticism. However, I'll subsequently go on to argue that we have independent reason to be unhappy with the neo-Quinean paraphrase strategy: The proposed replacement will take on board aspects of both the neo-Quinean and neo-Carnapean strategies.

1.4. Semanticism: merely verbal disputes

The semanticist gives us the following *hostile* account of revisionary ontology's relationship to our ordinary existential claims:

1. Our ordinary existential claims are by and large true— all the Moorean truths, such as my having hands or there being tables are certainly true. The same goes, however, for the revisionary ontologist's claims.
2. Since both sets of claims get to be true there is no conflict - this is a *compatibilist* account.
3. How on earth can both, e.g., 'there are chairs'? be true and a theory that explicitly states that 'there are no chairs'? The revisionary ontologist and the ordinary folk simply mean different things by their sortal terms and/or quantification phrases— so the proposition expressed by the first sentence is not the negation of the proposition expressed by the second (and *vice versa*, of course).
4. Does this save the day for revisionary ontology? Alas not, since according to the semanticist, revisionary ontological disputes are not *substantive*: which quantifier to use is an arbitrary/pragmatic choice— so nothing hangs on whether

we use a ‘nihilist quantifier’ or an ‘English quantifier’ to play the quantification role in our language / idiolect. As such, the revisionary ontologist’s claims are trivial and give us no reason to change our linguistic practices or theories or conceptual scheme.

There are many things that in everyday discourse we are happy to admit exist which some metaphysician or other has denied the existence of by way of an elaborate philosophical argument. Examples include, but are by no means limited to, properties (e.g. being red), numbers (e.g. π), and composite objects (e.g. tables, molecules, planets and human beings). Semanticists such as Hirsch²⁰ think that such arguments – especially those directed at the everyday objects of direct experience – can be dismissed easily, since the existence of all these entities is trivially settled by the rules of the English language, provided we assume that ordinary speakers don’t go around making implausible, systematic mistakes about what’s right in front to their eyes. Following a roughly Carnapian line of thought,²¹ it’s assumed that where the actual rules of English don’t determinately settle the matter, we’re left with an arbitrary choice of which precisification of the English language we want to speak – one where the existence of the entities in question is true, or one in which it is false.

Hirsch holds that (a good many) ontological disputes are ‘merely verbal’ (p.232). In the paradigm case this is due to the fact that those on one side of the dispute are speaking²² ‘an alternative language’ (p.235) to those on the other side of the dispute.²³ When the nihilist disagrees with the universalist over whether composite objects exist, each side (unwittingly) employs a different *existence-like concept* from the other when they use the word “exist”. Here we’ll refer to the universalist concept of existence as U-existence, and the nihilist concept of existences as N-existence. Hence, the universalist means something like “there *U-exists* composite objects”, while the nihilist means something like “there doesn’t *N-exist* composite objects”. On the assumption that *U-existence* isn’t *N-existence*, there’s no straightforward sense in which they can be said to be engaging in a substantive disagreement. Meanwhile, ordinary users of English can rest assured in the truth of their ordinary existential assertions – vindicating the spirit, if not the letter, of the Moorean

²⁰ Hirsch (2002), (2005), (2008), (2009). Compare Thomasson, e.g., (2009), (2007).

²¹ See, e.g., Carnap (1965).

²² Or, at least, attempting to speak as if they were users of such a language – perhaps they cannot avoid speaking ordinary English.

²³ Eklund (2009) raises difficulties for the possibility of there being alternate languages in this sense. See Sider (2009) for one way of responding to the objection.

complaint against revisionary ontology.

Hirsch argues for this conclusion via interpretive charity, which is a generally accepted assumption on both sides of the metaontological debate.²⁴ On the assumption of charity, we should interpret a community of speakers²⁵ such that the majority of their utterances come out true, and we don't suspect them of making systemic mistakes about very easy or obvious matters – such as mistakes about what's 'right in front of their eyes'.²⁶ Another way of putting it is that we should attribute the most plausible meaning to a term used by a language community, given their patterns of utterance of that term in particular contexts. If a community of speakers use the term “car” frequently when gesturing towards automobiles, and only rarely in the presence of giraffes, we should take it to be more likely that they mean their term “car” to refer to automobiles, rather than that they mean to use it to refer to giraffes, but are very often mistaken about whether there is or isn't a giraffe present.

Importantly, interpretive charity is not supposed to be 'a matter of generosity' but rather a principle that is 'partially constitutive of linguistic meaning'.²⁷ In the case of the nihilist and the universalist, although they see themselves as engaged a substantive non-verbal dispute – and mean to be interpreted as such – Hirsch thinks that charity demands we interpret both sides as being right in their own language rather than as making systematic mistakes about what exists. For Hirsch, the only good question in the vicinity of their dispute is, “Is 'exists' in English referring *U-existence* or *N-existence*?” In this case he would say that the answer is clearly that 'exists' in English means neither *U-existence* nor *N-existence*, since ordinary speakers countenance the existence of some of the universalist's mereological fusions (e.g. chairs), but not others (e.g. the mereological sum of the chair I'm sitting on and all the hydrogen atoms in the Horsehead nebula), and therefore evidently use 'exists' neither in the way the universalist uses it²⁸, nor in the way that the nihilist uses it.

The challenge to the van Ormwagean, and the neo-Quinean more generally, is clear: If ontological commitment is a mind/language-dependent phenomenon – and thus ontological disputes are non-substantive – then ‘revisionary’ ontology simply can't be

24 See, esp., Sider (2001*a*, Introduction); (2001*b*).

25 See Hirsch (2002), (2005).

26 Hirsch (2005, p.75)

27 Hirsch (2009).

28 Arguments to the effect that the ordinary folk tacitly restrict their quantifier notwithstanding.

meaningfully revisionary: We'll be just as well sticking with our ordinary quantificational claims as endorsing the nihilist's, since the choice of one ontology or another – when we consider very general questions such as whether there are tables in addition to simples arranged table-wise – is really just a choice of language, and there's no (not merely pragmatic or incidental) reason to prefer one language over another. However, it's by challenging this last assumption that proponents of revisionary ontology might hope to rescue themselves from falling into semanticist indifference about ontological commitment.

Sider argues, following Lewis, that interpretive charity – 'use' – is not the only the factor that fixes the meaning of our words.²⁹ We must also consider 'eligibility', for there are some things that are intrinsically more eligible to be meant than others. Sider takes eligibility to coincide with naturalness: one meaning is more eligible to be meant than another if it is more natural than its competitor; if it *carves the world closer to its joints*. Sider takes an *electron* (from physics) and *existence* (or, at any rate, the unrestricted existential quantifier, \exists , from logic) to be instances of perfectly natural features that are irresistibly eligible to be meant, if usage facts locate a term roughly in their vicinity.³⁰

The Siderean believes that whatever language we are speaking (including ordinary English) and in whatever context (whether we're debating statue-lump puzzles in the ontology room or counting cucumbers in the kitchen) eligibility has a very strong role to play in fixing the meanings of our terms. This puts the Siderean view very much at odds with the Hirschian position we have set out, which recognises no such role for eligibility. But it's worth pointing out that merely conceding *some* role for eligibility in fixing the meanings of our terms isn't enough to ensure the substantivity of ontological debate: we can imagine a concessive Hirschean (or timid Siderean) allowing some role for eligibility but saying that when it comes to interpreting the different sides of the composition debate charity considerations win the day overall and we must still interpret the two sides as talking past each other. Hence, the Siderean solution proceeds in four steps, each of which is open to challenge:

1. Appeal to Lewisian naturalness.³¹
2. Extending Lewisian naturalness to the logical vocabulary of our theories

²⁹ Sider (2001a, introduction).

³⁰ Sider (2001*b*)

³¹ Recall §0.3.

and to quantification in particular.³²

3. Defending the metasemantic role of Lewisian (or, now, Lewis-Siderean) naturalness— i.e. natural vocabulary as reference magnets.
4. Arguing (or assuming) that the most natural meaning for the quantifier is sufficiently reference magnetic to get to be the meaning of our term '∃'.

On the Siderean view, then, it's overwhelming likely that the nihilist and the universalist are not engaged in a merely verbal dispute, since there's only one plausible thing they – and, for that matter, ordinary English speakers – could mean by 'exists' and that's *exists*, in the perfectly natural sense. This yields a unique, substantive answer to SCQ. The Hirschian, on the other hand, will maintain that the role of eligibility, if a factor at all, can't plausibly be strong enough to ever bring it about that the vast majority of speakers' existential statements turn out mistaken— so even if there is a most natural meaning for our quantificational vocabulary out there, we could never get to mean it (which might itself constitute a *pro tanto* reason for not believing in it in the first place).³³ The upshot remains, for the Hirschian, that interlocutors in the composition debate are inevitably speaking past each other and do not make claims that conflict with our Moorean intuitions.

Here we appear to reach something of a dialectical stand-off concerning what the *right* account of revisionary ontology versus Moorean-style claims is: we can either go semanticist, or endorse a Siderean version of the error-theory,³⁴ depending on how strong we think the pull of reference magnetism is (if we buy Sider's naturalness story at all). For the revisionary ontologist obviously only the second of these two options is going to seem attractive. However, dialectically, it would be nice for revisionary ontologist's

³² For a thoroughgoing explication of this see Sider (2011): though the complaint from some quarters what was an intuitive notion when applied to predicates is abstruse when applied to e.g. logical vocabulary is likely to be recalcitrant.

³³ Dorr (2005) responds to Hirsch by arguing that if each language community (i.e. a nihilist language community, a restrictivist community, and universalist community), on encountering the others, is happy to translate their talk without considering the other language to be expressively impoverished, then this is evidence for the truth of nihilism: If believers in composition think that nothing of importance is missing from the nihilist's language, this suggests the missing terms like 'parthood' and 'compose' are semantically defective, in the same way as defunct theoretical terms such as 'phlogiston'. If this is right, it suggests that those sympathetic to Hirsch's way of framing the debate should strongly consider adopting nihilism.

³⁴ Ted Sider has discussed some details of how this might end up looking in conversation. I take Sider's actual view on this to have more in common with the non-error theoretic view to be proposed.

(collectively) to have a possible way of telling the story that doesn't rest quite so explicitly on the controversial reference-magnetism metasemantic machinery. Additionally, in the next subsection I'm going to give some independent reasons for disliking an error-theoretic reconciliation strategy. Therefore, in the final subsection I'm going to put forward a positive suggestion for a non-error-theoretic, substantive account of revisionary ontology that sets out to be more neutral about reference magnetism.

1.5. *The Errors of Error Theory*

What I have against the error theory³⁵ (even in its less strictly error-theoretic versions) is that it largely misses the theoretical role that the two levels of 'alethic commendation' are playing - or at least should be playing - in the debate. We need a term reserved for those sentences that get *perspicuously* at our ontological commitments, as those rescued by paraphrase fail to do, and a term for those sentences which may not succeed at the former but, as Hirsch would put it, 'get the facts right'. What I don't think is helpful is to make a division between which sentences perspicuously represent our commitments and those which deserve commendation for partly doing so or being connected to a statement that does so via a process of regimentation or paraphrase.

This sets up is a kind of spectrum from exact correspondence with reality to non-correspondence with reality: At some point along this spectrum we, as a language community,³⁶ happen to draw the line between truth and falsity. But where this line is drawn seems beside the point for the purposes of metaphysical debate: we have some sentences which meet the ideal without needing any degree of tolerance (wherever we draw the line) and those which don't; of those which don't some still get *something right* - deserving alethic commendation - whether true or not. Thinking of these as two 'levels' of alethic commendation with the second improving on the first is not especially helpful, however. Rather, it would be better to consider each as tracking along two different axes: Let's call these '*perspicuity*' and '*factuality*'.

Perspicuous Unfactual Statements:

³⁵ For a robust defence of error theory in ontology and elsewhere see Daly & Liggins (2010). While they refute many general arguments against error theory, my argument is based on focused considerations for the present debate and not motivated by a general suspicion of error theory.

³⁶ Or, on Horgan & Potrc's (2008), this is set by contextual features of the discourse. I somewhat more sympathetic to this view, as it at least has the resources to guarantee that when we are doing ontology truth/correctness is perfect correspondence with the world.

We can say something is metaphysically perspicuous but be dead wrong about the facts. Such as, supposing nihilism is true, I say that there are no chairs in the room (even though I am sitting on some simples arranged chairwise). While this is right with respect to ontology – there aren't really any such entities in the room or anywhere else – this statement conveys, in English, that the room (reality) is a certain way – i.e. *non-furnished* – which obtains when there are no simples arranged chairwise: But, *of course*, there *are* some simples arranged chairwise.

Perspicuous Factual Statements:

If I had just come out and said there are some simples arranged chairwise, in the above case, I could (*ex hypothesi*) have said something that was both perspicuous and got the facts right.

Unperspicuous Unfactual Statements:

Or, in the case above, if I was feeling contrary (or had just been convinced by an argument for gunk), I could have said there are no simples arranged chairwise and conveyed misleading information both about ontology and one's prospects of sitting somewhere with adequate back support.

Unperspicuous Factual Statements:

We can also, of course, be right about “the facts”, but wrong about the underlying metaphysical structure of those facts, and that's just to be right about one thing and wrong about another, not to be “sort of right” about the latter thing. Asserting that there are chairs in the room would be the obvious example of such a statement in the above case: It conveys that wordly conditions are optimal for back-supported sitting, but by quantifying over entities that do not really exist.

To pursue the idea of unperspicuous factuality to a greater extreme: Consider LC! – by our lights a somewhat bizarre language community, but who in many matters are quite intelligible to your average monolingual English speaker. Members of LC! expresses the presence of a chair by stating the proposition *that the angels are dancing counter-clockwise on red and purple dragons flying northward.*³⁷ As such, members of LC! are right to say ‘the angels are dancing counter-clockwise on red and purple dragons flying

³⁷ Or something reasonably close to this proposition (see next footnote).

northward when confronted with a chair'.³⁸ How well – or, rather, how badly! – they do with respect to perspicuously representing the ontological commitments of talking about chairs, or angels dancing counter-clockwise on red and purple dragons flying northward, is besides the point. For LC! it is something in the vicinity of a Moorean fact that there are angels and dragons doing their thing if they have a place to sit with back support. Perspicuity can go badly wrong without adversely affecting factuality.

Supposing nihilism is true, both LC!-ers and ordinary English speakers have an unobscured but adequate way of expressing the fact that there are some simples arranged chairwise. The cosmic furnishings conspiracy imbedded in their mythology also leads LC!-ers to express the existence of tables as 'There are some abstract forms of ferret-like creatures stacked on a slowly waking Cthulhu.' This might be a very good pragmatic reason to prefer speaking English, but it's not straightforwardly apparent whether English speakers do much better than LC! in perspicuously and systematically representing our ontological commitments. Such talk goes right for sure, when it does go right at all, in exactly the same way that dancing angel and Cthulhu talk goes right: It gets the facts right!

To be clear about the nature of factuality as I'm conceiving of it: On my view, it's not just that "there is a table" gets it right about the practical consequences of there being simples arranged chairwise, or how the world appears qualitatively. That would be enough, in a way. But, as realists, we can help ourselves to a little more. We conjecture

³⁸ The case I have in mind is one which, by hypothesis, the community in question use the terms 'dragon' and 'angel' and 'red' and 'dancing', etc., in various compound expressions in such a way that an ideal radical-interpreter would assign to these terms meaning recognisably close to the meanings of those terms in English. However, by hypothesis, certain of their express views on theology, chemistry, metaphysics, cosmology, etc., combine to make it plausible that the community are not speaking metaphorically when they use the phrase 'the angles are dancing [...]' – they are asserting that it is strictly and literally true that these specifically-shaded airborne individuals are participating in the stated recreational activity, so as to bring about (by some complex chain of occult events) all the local phenomenal / empirical / pragmatic consequences we would ordinarily associate with the presence of a chair. I happily concede that the community's linguistic behaviour would have to deviate from ours in a many practical situations in lots of other fabulous (though here unspecified) ways – ways which, taken as a whole, would have to have an impressive internal logic and invite outside onlookers to establish certain loose and informal 'bringing principles' between the community's intangible mythology and events in the natural world. But I see no reason to suppose that such a case is an in principle impossibility. Again: all that's needed for the example to be cogent, besides that 'the angels dancing[...]' is taken literally, is that the meaning assigned to terms such as 'angel' and 'dancing' bear some minimum of resemblance to the meanings of those terms in English, such that we can at least make an approximate stab at guessing what their function would be in most (though, as noted, by no means all) other compound expressions uttered by this community.

that these practical / empirical / phenomenological states are caused in us by *something* out there in the world, and that whenever the world is in that state, it will provide us with the corresponding phenomenology and empirical evidence, etc. So, what our existential sentences comes to express is that the world is in *such and such* a state, where this is a very precise sort of state, but where we as the speakers of the language don't have any direct access to the nature of *such and such* state beyond it being a state which causes events like E. To use a phrase from Williams (2012), "there is a chair" has exactly the same *reality requirements* as "there are simples arranged chairwise" if the former is a statement of everyday English while the second it (let's just suppose) an ontologically perspicuous sentence aiming to express the same fact, i.e., both say that this is what reality is like:



Figure ii. Some simples (in a chair configuration).

That is: (*ex hypothesi* as we said that "there are simples arranged chairwise" is a metaphysically perspicuous way of expressing this fact) that there are simples arranged chairwise. You can't get more *right* than that about the state of the world, at least qua whether there's furniture in the room, so it's odd to single one out as true on these grounds. It's just that for one of these sentences its surface grammar corresponds perfectly to the underlying structure of the fact in question. As for how perspicuous the everyday English sentence is, we remain agnostic at this point since it is a further question, perhaps of interest (we'll explore this in §3.2), but not relevant to assessing the revisionary ontological claim nor to assuaging the Moorean. Factuality certainly suffices for truth in so far as the Moorean cares about truth, so factuality is all that's need for the latter task, regardless of whether we decide to follow the error-theorist by building perspicuity into

truth or not. We get to be very certain about our Moorean claims and understand this comes from more than just epistemic conservatism or, worse, a reliance on prejudice or dodgy folk wisdom. Disparage “common sense” as much as you like, I’ll still know I have hands – and, as we’ll see, this goes for sceptical brain-in-vat hypothesis as much as radical, but not intentionally sceptical, ontological theories.³⁹

Furthermore, metaphysics gets to proceed without the constraint of providing a paraphrase scheme: supposing a statement about chairs can be easily paraphrased into a perfectly perspicuous statement about simples, if anything this points to it having fairly high perspicuity on some relative measure – but it tells us nothing about its factuality. It shouldn’t be a constraint on perfectly perspicuous theories that they can show statements of our ordinary discourse to be perspicuous to any degree – since that’s just not something we can take as data.

The flipside of these advantages –the putative disadvantages – are: Firstly, that we give up any kind of default assumption of partial or indirect correspondence between the structure of our ordinary speech and the structure of reality. While I expect that there is a fair amount of correspondence, I think it was wrong, as I’ve just said, to have this as a default assumption in the first place – we’re just not in position *prior* to metaphysical inquiry to have any confidence in where such partial or indirect correspondence might lie. That ordinary speech *gets the facts right* is enough to recover all Moorean intuitions without claiming some special metaphysical insights for ordinary language.

Second, we lose either “data” for our metaphysical theorising or an elaborate system of checks and balances for ensuring a radical metaphysician is worth her salt.⁴⁰ But the job of metaphysicians should be to uncover any partial or indirect correspondence wherever it might lie – in the construction of a metaphysical semantics⁴¹ – and where there is none to be found, to simply note the absence. Mapping out what, if any, systematic correspondences there are between the structure of ordinary language and the structure of the language of the world (and thus the world itself) will be project that that might proceed once we have a theory of reality, including our ontological commitments, in place. It should not be a prerequisite that any ontological theory worth taking seriously should assume the existence of a strong correspondence from the outset and be immediately at

³⁹ See §2.5.

⁴⁰ Van Inwagen (1990).

⁴¹ Sider (2011, §7.2.2.)

pains to elaborate before proceeding further.⁴²

Yes, it would be nice to have a simple, systematic metaphysical semantics that links ordinary talk with reality. Williams notes⁴³ that simpler, finitary theories of this sort will be easier (that is, possible) to evaluate compared to unsystematic non-finite theories and so should be preferred on these grounds, all else equal. I agree with this, but we should be very cautious about changing our theory about what reality is like arrived at on independent grounds to make our metaphysical semantics look pretty: if we have a theory of the world already, by all means pick the simplest metaphysical semantics available for that theory— we should prefer to say sentences of that theory correspond to sentences of English in the most systematic and straightforward way possible. But I see no particular motivation to revise our ontology to service the construction of a simpler metaphysical semantics— yes, we're always free to revise any part of our theory if it enables us to do better overall, but simply pointing to this Quinean truism ignores that in general we'll have strong independent reasons *not* to revise our ontology - i.e. the reasons we had for endorsing the ontology in the first place - and it's not clear how one should weigh up the value of a simpler metaphysical semantics against these more parochial considerations.

My strong suspicion is that the considerations from metaphysical semantics don't count for very much since there a plenty of good independent reasons to suppose, prior to any metaphysical investigation taking place, that the relationship between ordinary language and ontologese will be unsystematic and messy: All is not equal. We have poor epistemic access to the ultimate structure of reality in the first place; that access was probably worse in the past when the majority of our language conventions were formed than it is now with the advent of mathematically informed physical sciences; and members of the language community patently care more about practical and efficient communication of facts than they do about ontological perspicuity— much less about maintaining some indirect but systematic link to reality. It's little more than wishful thinking to assume that reality *forces* on us such a link by making the most practical languages for a linguistic community necessarily the (indirectly) ontologically perspicuous ones, given the diverse range of interests and needs a single speaker could have, let alone different speakers within a language community, or different language communities as a whole.

So, insisting on indirect correspondence as a prerequisite for revisionary metaphysics is

⁴² More on this in §3.

⁴³ Williams (2012).

understandable but ultimately misguided.⁴⁴ And it's misguided to insist on a systematic paraphrase scheme connecting sentences of ordinary language with sentences of your ontological theory if your penchant for this result leads you into otherwise unmotivated revision or censoring of your ontology.

To elaborate on this point, consider this instance of paraphrase (paraphrased statement → resulting sentence):

'the cat is on the mat' → 'some of the simples arranged cat-wise are almost in contact with the simples arranged mat-wise and all the simples arranged cat-wise are directly above the simples arranged mat-wise.'

A *systematic* paraphrase scheme for dealing with such expressions might give you some rules for taking sentences of the form 'the *a* sat on the *b*' and producing sentences of the latter form concerning arrangements of simples. Without some systemacity, you would just end up with an infinite list matching up each possible unperspicuous sentence with some perspicuous sentence. The van Ormwagean error-theorist needs to take as a constraint on revisionary ontological theories that a *systematic* paraphrase scheme be available for taking you from 'correct' (i.e. factual unperspicuous) sentences to 'true' (i.e. factual perspicuous) sentences: without this constraint, there's simply no way of evaluating the adequacy of the paraphrase scheme by seeing that it holds in certain toy cases and generalising from there – as such, no way of evaluating whether the theory successfully preserves the factuality of Moorean claims. Similarly, if a finite unperspicuous sentence can be mapped only to an infinitely long (perhaps highly disjunctive) perspicuous statement, how are we to understand the new sentence and assess it for factuality? If the error theorist can't meet these constraints, it's not clear that when they say 'there are no motor vehicles' that they are not in fact asserting something with the implication that if you stand in the middle of the motorway, you won't be run over.

But all that's needed is that one's total theory account adequately for the data that feed into our judgements of factuality: i.e. that it can account for my having sensations as of motor vehicles and my sometimes finding myself hurtling along near the ground at 70mph, etc. There need be no presumption whatever that there's any systematic relationship between (say) the nihilist's way of accounting for the data and the sentences that English speakers use to express particular factual situations. Similarly, that a range of subtly different factual situations that an English speaker sums up in one simple phrase

⁴⁴ See Horgan & Potrc (2000) for a metaontological view that explicitly advocates a reconciliation strategy based on indirect correspondence.

doesn't correspond to an infinite disjunction of incomprehensibly long perfectly perspicuous descriptions: e.g. for 'a ball' a hugely long description of the spatial relations between billions of simples such that an English speaker would recognise them as arranged in the shape of a ball disjoined to an almost identical description of a very similar 'ball' with one simple out of place and so on.

As we'll see in §4, if one uses systematic paraphrase as a guide to metaphysics, one risks sacrificing a simple and elegant theory of the underlying metaphysics to the goal of ensuring that our theories have sufficient resources for quickly and conveniently expressing certain parochial, anthropocentric facts. But our ontological theories need not play the dual role of perspicuously theorising about underlying reality and of being of a convenient tool for human beings to convey factual information in their day-to-day lives: Humans convey factual information in their everyday lives just fine without any direct reliance on metaphysical or folk-metaphysical theory.

1.6. Metaphysical Perspicuity and Ontologese

In this subsection, I set out my preferred reconciliation strategy for revisionary ontology and common sense existence claims. It contains the following basic commitments:

1. Revisionary ontological talk can be factual, as can Moorean existence claims.
2. There is no incompatibility between revisionary metaphysics and Moorean claims: 'There are no chairs!' uttered by the nihilist in the course of theorising may carry exactly the same factual content as 'But of course there are chairs!' uttered by a perplexed bystander.
3. Since there's no incompatibility with respect to factuality, there's no need for an error-theory or paraphrase scheme to reconcile the two discourses.
4. However, contra semanticism, ontological debates are substantive: The nihilist and the universalist disagree over whether saying 'there are tables' perspicuously represents reality rather over its factual content.

The two ills of paraphrase and overreaching Mooreanism are closely entwined, acting as unwarranted conservative constraints on revisionary metaphysics: I have no issue with a general weak principle of conservatism which counsels not adopting outlandish, wildly speculative new ontologies without principled and compelling reason to adopt such an ontology over our existing scheme. But I fear the more specific van Ormwegian and Moorean constraints cause us to systematically misjudge the merits of radical ontologies

by imposing obstacles that are at once too difficult to navigate and irrelevant to what we should be trying to assess. The overreaching Moorean assumes “common sense” has some default insight into metaphysical perspicuity and therefore that metaphysical claims that come into apparent conflict with common sense claims instantly have a mark against them. The van-Ormwagean error-theorist similarly assumes that the structure of our ordinary statements worthy of alethic commendation are those not so far gone in the perspicuity stakes that they can at least be made maximally perspicuous by systematic regimentation or paraphrase. This is weaker than the Moorean’s restriction but more insidious, since it has the air of plausibility to those who are attracted to revisionary metaphysics but who are understandably cautious of it getting out of control and leaving common sense behind altogether.

Whatever metasemantic story we want to tell about how ordinary language terms get their meaning, it’s clear that ‘I have hands’ in English conveys certain factual content – namely, that factual content the Moorean is fist-bangingly certain of when she is moved to protest that if she’s sure of anything, she’s sure she has hands. We could piggyback on Hirsch’s story, identifying factuality with truth and say that interpretive charity is the prime mover when it comes to determining the meaning of ordinary language sentences: But I’m more certain that ‘I have hands’ conveys the Moorean proposition that *I have hands* than I am about the correctness of any metasemantic just-so story or the truth of any particular theory of truth. So, by hook or by crook, we have an assignment of factuality (= truth?) values to English sentences:

‘There are chairs’ = Factual

‘There are no tables’ = Not Factual

‘There are unicorns’ = Not Factual

‘There is a ferret in my house’ = Factual

‘There are exactly 1003 grains of rice in my house’ = [I don’t know, but I can go count.]

We then have our ontological theories: formal artefacts aimed at systemising data from experiences, scientific observations, etc. These should be taken, in the first instance, as analogous to mathematical formalisms in science rather than as straightforward statements of natural language. As ontological realist⁴⁵ we believe that the best – most systematic – of

⁴⁵ Recall §0.2.

these theories will likely include some bit of ideological machinery with the inferential role of the existential quantifier. Once we've identified that '∃' in the theory plays this quantificational role, we can just stipulate, along more or less Quinean lines, that the objects in the domain of that quantifier are that theory's ontological commitments. We thus have a theory relative criterion for ontological commitment. Again, whatever story underlies how the meanings of the formal components of our theories get fixed, it's surely bizarre to think that charity considerations determining the meanings of our natural language existential claims could force '∃' in the formalism to have a domain other than the one specified by the theory. As such, we should treat the theory as a 'different language' from ordinary English in the semanticist's sense: We can't just 'read off' existentially quantified statements of the theory as ordinary language existential claims.

Rather, when an ordinary English speaker says 'there is a chair in the building' we should identify the factual situation that is being conveyed and try to describe that factual situation using the theory. The theory will spit out an existentially quantified statement carrying certain theory-relative ontological commitments. Whatever these ontological commitments are, we should say these are the ontological commitments of 'there is a chair in the building' in English.

Now we could, perhaps, have some pragmatic recipe for regimenting the English claim into a sentence in predicate logic, which would allow us to see whether the regimented English claim carries ontological commitments (relative to regimented English) which mirror the commitments (relative to our Best Theory) of the sentence produced by our Best Theory. If this is the case we can say that the English claim is ontologically perspicuous: it may thus license the claim that 'there are chairs' is really true in English, or that chairs are real. If there is not exact similarity between the commitments of the theory and the commitments of the regimented English claim, then we should say the English claim is not perspicuous - there aren't really any chairs. But this will in no way undermine the factuality of 'there are chairs'. Whether these moves are licensed, however, depends on how felicitous our regimentation of the English claims is: It must be one that respects the intuitive correspondence between ordinary existential statements and existential quantification (with exceptions, perhaps, for obviously non-standard cases such as overt metaphor).

Our ontological commitments, though, are - in true Quinean spirit - whatever our Best Theory quantifies over. The (optional) realist - or perhaps I should say hardcore realist - move is to say that our Best Theory can only be true if it's structurally isomorphic to the world in some (metaphorically) analogous way to how we demand regimented English be structurally analogous to our Best Theory, if its quantificational claims are to

carry genuine ontological commitment. Only if this criterion of world-theory perspicuity is met will its theory-relative commitments be genuine hardcore realist ontological commitments. However, we could opt for lighter form of ontological realism that only recognised the ontological perspicuity of claims relative to our best theory: theory-theory perspicuity. On this light view, the only constraint on Best Theory is that it gives an adequate account of the facts that is most systematic, unified, simple, etc.: Even if we only seek more virtuous theories for pragmatic reasons, we can still makes sense of a substantive debate about whether ‘there a chairs’ is true according to the best systematic representation of the facts according to mutually agreed criteria. Importantly, our approach to doing ontology stays fixed whether we’re hardcore realist or realists light: There’s no sense in which our ordinary language claims guide or constrain ontological theorising, and no sense in which ontological theorising threatens the factuality of any ‘Moorean truths’.

We can roughly fit this conclusion into the existing dialectical positions in the literature by taking the statements of our Best Theory to correspond to statements of ‘Ontologese’. The Ontologese picture is that Ontologists supposedly speak a different language from the ordinary folk when they enter the ‘Ontology Room’ and start making ontological statements: This is a privileged language, the language we should speak if we want to do ontology, and so gets around Hirschian semanticism. Dorr (2005) cashes out Ontologese as a language that’s not impoverished in its theoretical vocabulary relative to any other language.⁴⁶

Meanwhile, according to Cameron we should think of our ontology as consisting of truthmakers.⁴⁷ If I say “there are mereological simples” what makes this true are mereological simples (let’s reasonably suppose), so mereological simples have real existence. However, if I say “there are cats” the nihilist will say this too is made true by mereological simples and not by any cat – so the truth of this sentence once again entails the real existence of simples, while it does not support the real existence of cats. As existential claims about cats are made true by simples, we might say that simples ground the existence of cats, while also grounding their own existence by making true sentences about simples. On this scheme, we can only read off our ontological sentences in the Quinean way, on the assumption that everything is its own grounds: in this case, sentences

⁴⁶ I think this general view could be given readings favourable to either hardcore realism or to realism light, though I take Dorr’s is concerned in the paper to be with illuminating the hardcore position.

⁴⁷ Cameron (2010a). Compare Melia (2005).

about cats would be made true by cats (at least partially— simples and whatever else could be involved too, so there might still be hidden ontological commitments in a sentence). We could imagine speaking a language – again, we can call it Ontologese – in which it was a requirement for the truth of any existential claim that the object being quantified over was among the truthmakers for that sentence. Conducting ontological debates in this language would make our ontological claims more perspicuous, in our *heavyweight* sense.

Sider (2011) also elaborates the strategy in a hardcore way: He thinks the Ontologist can simply stipulate that she is speaking a language that *carves at the joints* – where every term, including the quantifier, has its most *natural* meaning in the Siderean sense already elaborated. Sider employs Ontologese as a *backup* if it turns out reference magnetism is weak, a language where reference magnetism is stipulated to be super-strong, so that he can make claims in a joint carving language, and thus be able to do metaphysics, even if it turns out that English isn't such a language.

The above is a good further illustration of why not to get hung up on the issue of whether accurate statements of the facts cast in unperspicuous terms are deserving of being called true or not—⁴⁸ for the rest of the thesis I will proceed with this assumption, but only for the sake of simplicity. Everything will come out more or less the same if not. I shall also be assuming the *hardcore* version of the Ontologese Strategy, as I elaborated it, but I think so little will hang on this that it's scarcely more than a point of autobiography: I don't think it will serve our purposes to get into a close quarters fight over whether mine, or Sider's or Cameron's way of gesturing to the *hardcore* nature of ontological commitment is most illuminating. My take home message is rather that understanding revisionary ontological claims as statements of Ontologese in the way I elaborated allows us to retain the factuality of Moorean claims while allowing us to see the theoretical utility of making revisionary ontological claims: Since ontology is aiming at perspicuity, it can 'revise' our theory of the world by enhancing perspicuity while preserving factuality. Importantly, as we've seen, the enterprise can carry on constrained only by the need to fit the empirical data rather than the need to have some systematic relationship with ordinary English claims we take to be factual.

⁴⁸ I say more about the status of existential statements cast in non- (or partially) joint carving terms in §3 where I discuss non-fundamental ontology.

2. Speaking the Language of the Ontology Room

The last chapter concluded by endorsing the construal of ontological disputes as conducted in a maximally metaphysically perspicuous language— Ontologese. Some have raised objections to strategies of this sort, claiming that introducing a 'language of the ontology room' only confuses the issue since a) we have no idea how to interpret – let alone speak – this language, and b) all the interesting ontological questions are questions best phrased in ordinary language, anyway. Here I hope to show that speaking the language of the ontology room is in fact very simple, and to persuade you that all the really interesting ontological questions are best asked in this language. I will finish by setting out some advantages of employing the Ontologese Strategy (OS).

Section §1.1 briefly recaps the Ontologese Strategy and discusses the difficulties for such an approach posed by Korman (*forthcoming*); in §2.2 I respond to those difficulties, putting forward a positive thesis concerning how it is we can understand and speak 'the language of the ontology room', while §2.3 elaborates on a possible way of fleshing out the schematic response and §2.4 shows how the same response can deal with a related worry from Hofweber (2009). §2.5 discusses how employing OS can help us respond to Moorean-style objection to revisionary ontology.

2.1. *Ontologese*

In the previous chapter I advocated my own variant of a broad strategy for reconciling ordinary existential claims with revisionary ontological claims: According to the advocate of the Ontologese Strategy (OS), when the nihilist says 'there are no composite objects' this is to be understood not as a statement in English, but as a statement in this other language – let's call it *ontologese*. In devising his error-theoretic conception of revisionary ontology van Inwagen (1990) takes himself to be taking to heart the lessons of Quine, who tells us we are stuck with the ontological commitments of existential claims if we cannot 'devise some way of so paraphrasing the statement' to show that the putative existential

claim was 'an avoidable manner of speaking'.¹ Van Inwagen and followers of his methodology take this mantra to apply to the existential claims we assert in everyday contexts.² However, as we've seen, we should construct our ontological theories without concern for the truth or falsity of ordinary language existential claims: These can effectively be taken as stated in a different language from the statements of our ontological theories, such that putatively incompatible existential claims are not in conflict as they don't mean the same thing by the quantifier. We can be assured that the truth (factuality) of the ordinary existential claims we are confident in won't be undercut by revisionary metaphysical theorising, since belief in the revisionary ontological theory by itself allows us to remain indifferent as to alethic status of our ordinary language claims and agnostic as to whether they are ontologically committing.

No member of any language community has direct and infallible access to the ultimate nature (or structure) of reality, and the majority of members are not concerned with trying to speak about the world around them in terms that are maximally perspicuously. Thus, practical considerations and accidental quirks of usage may – for all we know – have lead us to speak a language in which the objects its true to say exist have no direct one-to-one correspondence with the ultimate furniture³ of the world. If this is right, and there may be a gap between true existential claims in Ontologese and those in ordinary language, then ontological debate should not be sidetracked by debates about what it's seemingly true to say exists in ordinary language. This would allow the nihilist to hold a theoretically virtuous ontological position without coming into any obvious conflict with the claims of ordinary folk about what there is. Thus, when the nihilist says “there are tables”, she should be taken not to be speaking ordinary English, but an ontologically perspicuous 'language of the ontology room'.

In this chapter, I will be defending this general strategy of introducing Ontologese ontological debates against the criticism that answers to ontological questions given in ontologese are both unintelligible and irrelevant. In doing so, I hope to flesh out the above story about why our ontological commitments are not our ordinary existential commitments, and why as metaphysicians we should care about our ontological

¹ Quotations from Quine (1948).

² It's not at all clear this was Quine's own view: this is not the place for Quine exegesis, but he may arguably have believed something much closer to my proposed account on which we should only look to our final theory for our ontological commitments without primary concern for how arbitrary statements of ordinary language can be reconciled with that theory – thanks to Scott Shalkowski.

³ <groan>

commitments, into something a little more explanatory and convincing.

Korman (*forthcoming*) raises two objections to employing a 'language of the ontology room'. First, he wonders how it is that we could possibly come to understand this 'fundamental' language. Secondly, he takes it as obvious that all the interesting metaphysical questions – which have been around much longer than the recent overt attempts to employ a language of the of the ontology room – are questions asked using ordinary language, with expectation of receiving replies in ordinary language. Korman uses metaphysical debates over 'Statue-Lump' cases to illustrate his concerns. The set-up should be familiar: Suppose you have a lump of clay (Clegg) at t_0 that turned into a statue (Sturgeon) at t_1 . All of the following seem true:

E1) Sturgeon exists.

E2) If Sturgeon exists then it coincides exactly with Clegg.

E3) If Sturgeon exactly coincides with Clegg, Sturgeon is identical with Clegg.

E4) If Sturgeon is identical with Clegg, then Sturgeon existed at t_0 .

E5) Sturgeon did not exist at t_0 .

Yet the above claims are mutually incompatible. This is the basis for (a version of) the well known Statue-Lump puzzle – all of E1 to E5 are intuitively compelling, but one of them needs to be given up. There are lots of ways you might try and go about this: for instance a perdurantist would give up E2, since she would take the statue to be a proper temporal part of the lump. However, Korman illustrates how a nihilist who endorses the OS approach to ontology might go about resolving the puzzle.⁴

The standard nihilist solution to this Statue-Lump puzzle would be eliminativism: to deny E1, and say that it's not the case that the statue exists – since no composite objects exist and if there were any statues (at least of the sort indicated in this puzzle) they would be composite objects. With E1 denied, no inconsistency arises. Yet, of all the possible ways out, this is probably the solution most at odds with common sense: *of course* there are statues. Perhaps our commonplace intuitions are wrong about how objects persist over time, or about what it takes to create or destroy an object through re-arranging smaller objects, but it would surely be crazy to deny that there can be such things as clay statues?

The OS ontologist can take something akin to the eliminativist way out. By employing

⁴ Recall §0.5, §0.6.

Ontologese she hopes to preserve the crucial insight of the eliminativist strategy while avoiding the difficulties it presents. Thus, the nihilist way forward on OS is to deny that the Ontologese sentence 'there are composite objects' is true. In English we might translate this sentence as 'there are *really* no composite objects.'⁵ Of course, this doesn't constitute a denial of E1, since E1 is simply the English claim that the statue exists, not the claim that it *really* exists. So the OS ontologist has not in fact made the apparently crazy move recommended by the standard nihilist approach. Though, given that each of E1 to E5 have yet been denied, you might be wondering at this point what the solution to the puzzle is supposed to be. Well, the important move is this:

The OS ontologist points out that since this is an *ontological* debate, we shouldn't really be conducting it in ordinary English. We should be conducting it in a language that is maximally ontologically perspicuous. Accordingly, the five claims should be recast in ontologese, employing the ontologese quantifier:

- O1) Sturgeon exists_o.
- O2) If Sturgeon exists_o then Sturgeon exactly coincides with Clegg.
- O3) If Sturgeon exactly coincides with Clegg, Sturgeon is identical with Clegg.
- O4) If Sturgeon is identical with the lump, then Sturgeon really existed at *t*_o.
- O5) Sturgeon did not exist_o at *t*_o.

Once the puzzle is recast in this way we can solve it by denying O1 (*English translation*: The statue doesn't *really* exist), without making the seemingly crazy claim that E1 is false—that statues don't exist. There's still the matter of what to say about the original English puzzle: we still have E1 to E5, a set of mutually incompatible propositions. But the OS ontologist can say something along these lines:

Firstly, since the problematic set of inconsistent claims is in English, we shouldn't be worried about this as metaphysicians.⁶ Secondly, why be so confident that rules of ordinary English must allow us to construct a consistent ontology, given that its users can go around happily communicating factual claims without a second thought to these puzzles? Thirdly, once you see that the real puzzle is best stated in Ontologese, it may be that the denial of some of the other claims – read in English – are not so implausible after

⁵Recall previous chapter. Cf. Cameron (2008a, 2010); Sider (Draft).

⁶ See §3, §4.

all. For instance, you might think that E3 is suspect. Many find it hard to give up E3 due to a suspicion of exactly coincident (numerically distinct) objects – especially the idea that there could be two distinct objects that not only occupy exactly the same space, but are supposedly constituted by exactly the same matter.⁷ Not everyone will share this suspicion, and some will be prepared to bite the bullet: but it is this worry that is one of the driving intuitions that gets the Statue-Lump puzzle off the ground. But if a suspicion of coincident objects is justified at all, surely this is best captured by asserting O3, not E3? By asserting O3 we ensure that there cannot be two distinct things in our fundamental ontology that are exactly coincident and made up of the same matter. But once we have O3, what does it matter if it's correct to *describe* certain situations as containing coincident objects according to the rules of the English language?⁹ We might choose to speak in such a way that we *say* that there are two objects present in Statue-Lump cases: but the most ontologically perspicuous way of describing the situation – the Ontologese description – makes clear that that's not what is really going on. Once we have Ontologese available as a tool in an ontological debates, we may see that some of our intuitions we were trying to express in English are best expressed in Ontologese.

Here, in the spirit of autobiography as illustration, is what I'm inclined to say about each of the English puzzle's statements as an OS ontologist:

1. Sturgeon exists = '*Moorean truth*'
2. If Sturgeon exists then Sturgeon exactly coincides with Clegg = *underdetermined in English (perdurantism vs. endurantism)*⁸
3. If Sturgeon coincides with the Clegg, Sturgeon is identical with Clegg = *false or underdetermined (e.g. in English statue sortals may well dominate⁹ vs. identity in English may well only be defined relative to sortals¹⁰)*
4. If Sturgeon is identical with Clegg, then Sturgeon existed at t_0 = *maybe false?*² (*just maybe identity in English temporally relative?*)¹¹
5. Sturgeon did not exist at t_0 = '*Moorean truth*'

⁷ Recall §0.5.

⁸ Cf. Hirsch (2009).

⁹ Cf. Sider (2001).

¹⁰ Cf. Thomasson (2009).

¹¹ Cf. Sider (2001).

Korman finds this general style of response to the Statue-Lump problem deeply unsatisfying. Since Ontologese is introduced as this brand-new language, especially for the ontology room, how are we to understand any of the words in the Ontologese sentence 'there exist no composite objects'? It seems the OS ontologist invites us to understand the phrase 'composite objects' more or less as we would understand it in English. However we are expressly cautioned that there 'there exist' should *not* be understood in the English sense. If that's so then, absent further guidance, how are we to avoid treating the phrase as anything except gibberish? We might as well have been told that there 'onxist' no composite objects.¹²

What's more, if we take the original Statue-Lump case above to have been set out entirely in English, it seems that the OS ontologist's answer makes no progress whatsoever. Since she accepts the truth of the English sentence 'there are composite objects' the inconsistent set of premises remains intact – there *are* statues, she admits. Claiming that 'there are really no statues' is neither here nor there, as it merely sidesteps the original problem. Korman argues that the attempt to reinterpret some of the premises as sentences of Ontologese is absurd: How can people have been puzzling for so long about a problem that's set up in a language they don't understand and have never dreamt exists? The premises all seem intuitive in ordinary English, but are mutually incompatible: thus do we not have a puzzle set up in entirely in ordinary English, deserving a response in that same language?

How can we speak the language of the ontology room? Well, to make the problem seem a little more manageable, let's clarify that we're only concerned with the meaning of the existential quantifier (\exists) and related vocabulary, such 'there are' and 'exists', when they appear in ontological statements. For the sorts of questions we are concerned with are whether there are any *xx* in the English sense of "there are" and whether there are any *xx* in the Ontologese sense of "there are"¹³ Thus Korman's challenge is: What does the Ontologese word existo mean? How did we come to know it had this meaning?

One could respond that we're using the quantifier in its most metaphysically perspicuous sense (quantifying in the way that carves nature at its joints, or whatever your story is about metaphysical perspicuity). Indeed, that's precisely what the OS ontologist

¹² (To borrow a memorable expression from an earlier draft of Korman's paper.)

¹³ See Sider (2013)

takes herself to be doing.¹⁴ We know that when we employ locutions such as “there *really* are” that we’re quantifying in the most perspicuous way because we’ve *stipulated* that these technical phrases have the required meaning: Or, as I prefer to think of it, evaluating a theory by perspicuity in addition to factuality is in some sense *constitutive* of what it is to do Ontology. This may be all well and good, as far as it goes, but I don’t think it really gets to the heart of the worry: If the meaning of the Ontologese existential quantifier is entirely unrelated to the meaning of the English existential quantifier, simply knowing that the Ontologese meaning is the most metaphysically perspicuous meaning isn’t going to get us very far in understanding the Ontologese sense of existence. Knowing that the Ontologese meaning is the most metaphysically perspicuous tells us something *about* its meaning but it doesn’t tell us *what it means*. Telling someone that you are using the word “indefatigable” to have its most common meaning is going to be no help unless they already understand what people usually mean when they use the word. So, how do we grasp the meaning?

Perhaps more can be said to bridge the gap between our understanding of the ordinary English sense of quantification and in the Ontologese sense. After all, whatever quantifier we’re using, our quantifying will still take the same logical form: it’s just the extent of their domains that differ. Similarly, whether we’re making a list of everything that exists or what “*onxists*” we’re still performing exactly the same sort of task. This goes some way to showing that the Ontologese sense of existence and the English sense are not entirely unrelated: they function in exactly the same way grammatically, at least. All one needs to grasp the Ontologese sense of existence – if one already has a grasp of the English sense – is an understanding of how the Ontologese criteria differ from English criteria when it comes to what makes it onto the list. Once we’ve gone this far, it might be hoped that knowing the Ontologese criteria are concerned with perspicuity will help us move from our ordinary understanding of existence to an understanding of the Ontologese concept. “Ordinary English gives you the concept of carving,” the caricature Siderean might say in full metaphorical swing, “in Ontologese we just carve closer to the joints!”

But is this really getting us very much closer to understanding the Ontologese sense of existence? Well, perhaps, but it also leaves us with some rather conspicuous holes in our understanding. For one thing, this explanation on its own isn’t much help in figuring out how we should think about the relative status of fundamental and non-fundamental objects. What would it tell us about the nature of ordinary objects if we were to learn that

¹⁴ At least if she’s a hardcore realist – recall §1.6.

the most natural way of carving reality (the Siderean again) doesn't (so to speak) carve around them? After all, we can and touch and taste ordinary objects. We know lots of facts concerning ordinary objects. What's the upshot of saying that they're not real? Certainly, it would mean that we wouldn't expect to see them popping up in your fundamental *physical* theories: but that's pretty uncontroversial.¹⁵ Few would hold that tables are bearers of perfectly natural properties of the sort that would get mentioned in our best physical theories anyway. If asked what sort of kind terms (if any) might appear in a fundamental physical theory, even the universalist, who believes in the common sense objects and many non-commonsensical ones besides, would plumb for terms such as “quark” or “lepton” or what have you, rather than “table” or “roast dinner.”¹⁶

Indeed, on this rather sketchy picture, it looks hard to block Korman's worries: What grounds do we have for thinking that only things that exist in the Ontologese sense are definitely ontological committing? How can we have any sort of intuition about whether objects that “onxist” can be co-located or extended or enduring, or anything else? We're told that concepts such as parsimony, parthood, spatial extension, persistence and the like – which we've formed in our everyday lives – can be seamlessly applied to the realm of Ontologese. In fact, we're told that any time we've idly puzzled about matters involving such notions in relation to objects, we were thinking about those objects that exist not in the ordinary sense, but in the Ontologese sense. Yet, as things stand, we have no guide to how to apply these concepts to the realm of the fundamental, nor any explanation of why we should expect them to be applicable. If we really don't have any grip on what it is for something to exist in the Ontologese sense – we just know that the Ontologese sense is stipulated to be the most perspicuous or natural sense – then it seems to me that Korman's criticisms go through. If we should be having ontological debates in a language we have no understanding of, then ontological debate is (for us) next to impossible.

And yet I'm not pessimistic about our ability to engage in ontological debate, and unlike Korman I don't see a need to retreat from OS in order to retain my optimism. In fact, I think that properly understood, such approaches help to make clear why we can have serious ontological debates about issues that would seem to be trivially settled by common sense. So, how is it that we are able to understand and speak the language of the ontology room?

¹⁵ This point is discussed in §3.

¹⁶ More on this in the next chapter.

2.2. A Common Sense of Existence

To meet Korman's challenge, let's recap the extensive similarities between the English 'exists' and the Ontologese 'exists': whichever quantifier we're using, it will still have the same inferential role; similarly, our sortal 'object' enable us to count the number of distinct things that fall under that sortal. The difference just seems to lie in determining what the domain of the quantifier is and how we should count using our most generic sortal:

This isn't like being told that the thing my friend Rob refers to when he says 'Tuesday' is eight-foot tall - in that case I'd be at a complete loss to understand him when he says 'See you next Tuesday' since my understanding of the term Tuesday plays a role in my mental and linguistic life that simply can only be filled by periods of time in the week, not anything that could be eight-foot tall. However, if I were told that what my friend Dave means by 'duck' is a type of duck, understanding would come much more quickly. I would get the gist of what he's saying 'there are three ducks on the pond.' For one thing, it implies in my idiolect the truth of 'There are at least three ducks on the pond'. What I'd want to know is what extra conditions my friend puts on the sortal 'duck' to be counted as a duck: when it turns out that the extra constraints exclude all ducks that don't have the physiology of mallards, I can translate my friend's (incorrect by the rules of English) use of 'duck' as 'mallard.' Everyone's happy.

While I'm not holding these up as a perfect analogies, the task of trying to understand what the ontologist means by 'exist' is much closer to the 'duck' case than the 'Tuesday' case— it's not that the ontologist is hijacking a common term to get across a completely different concept, but simply that what she means by 'exist', while playing broadly the same role in her cognitive life as what I mean by 'exist', is slightly more fleshed out or refined. When I see an existence question as weighty and important - 'Is there any such thing as the Higgs Boson?' - she does too, but sometimes when I (qua ordinary English speaker) see a question as having an easy or trivial answer - 'there are simples arranged table-wise in front of me— but is there a table?' - she claims there is no easy answer: metaphysics must be done. So, I want to know, what is it that this OS metaphysician packing into her idiosyncratic use of the quantifier or her generic sortal that will enable me to understand her better? I contend it is that sufficient to grasp the meaning of the Ontologese quantifier is a grasp of the ordinary English quantifier *plus* an understanding of the constitutive theoretical role of the Ontologese quantifier.

Consider: My other (*other*) friend, Ali - I have no shortage of weird friends - used to tell me often and insistently that my *red* jumper isn't red. At first I was at a complete loss:

I would point out that my jumper is approximately the same colour as this *red* apple, but Ali would respond by claiming the apple wasn't red either! We went on in this vein for quite some time, until one day we walked past a post box. I said "Look! My jumper is definitely the same colour as this *red* post box. She shakes her head. I expected her to say next, as now predictable, that the post box isn't red either: but, no, she says "yes, the post box is red, but your jumper isn't." To cut a long story short, it transpired that, having been working a long time and overtime at her job with Royal Mail, she had gradually slipped into thinking of red as the colour of post boxes, since this was the most *useful* meaning of 'red' for someone primarily interested in collecting letters - the particular shade of red that post boxes are painted just seemed like the best thing to mean by 'red'. Of course, she could have named the particular shade and continued to use 'red' the same as anyone else: But red things you can't collect mail from are just not *red* in any sense of interest to the postal worker— the subtly different shades of post boxes are all shades of red, but it confuses things to lump the non-post box shades in with these as all "shades of red" (surely post boxes should be uniquely picked out by one of the three primary colours?).

Ali has inadvertently slipped into speaking *Postologese* - a mail delivery perspicuous language. It might seem crazy to the bystander at first, but if you already understand what red is and you then come to appreciate the *theoretical role* that 'red' and red things play in the science of mail collection, you begin to see the internal logic and practical advantages of this way of proceeding for someone entirely absorbed in thinking and speaking about mail collection: and - most importantly - you could grasp the concept of Postologese 'red' from just a more general familiarity with red things and the English 'red', plus an understanding of red's constitutive theoretical role in mail delivery theory. This is the case *even if you had never actually seen a postbox* (sure, you can't imagine the precise shade, but you'd know enough to understand what Ali was on about when she used the term 'red'.)

The meaning of 'exist' in Ontologese, then, stands to the meaning of the ordinary language 'exists' as determinate stands to determinable: To grasp the Ontologese meaning is just to grasp the English meaning and some extra constraints on the application. As I said in the last sub-section, merely being told that the constraint is that all and only the real things get quantified over or counted isn't informative: that's like my 'duck'-friend telling me that he only counts the *genuine* ducks as ducks. But in both cases we can do better. Once I understand that he thinks all non-mallard ducks are in fact alien imposters in duck-suits, it's clear why he thinks that his non-mallard exclusionary use of 'duck' gets at something about the world that our ordinary use of 'duck' doesn't: Ours is a gerrymandered sortal that's all well and good for counting bird-shaped things on ponds,

but is unsuited for the additional theoretical work of combating the alien invasion (it does reasonably well at counting duck-shaped things on ponds, too, as a great many ducks are mallards – which is why his use of ‘duck’ doesn’t immediately strike me as unintelligible).

In the Ontologese case, the additional theoretical work is that of providing a complete and consistent set of answers to paradigm metaphysical puzzles: In the case of ordinary English it’s plausible that either nothing in our patterns of use of our sortal terms would provide an outside interpreter with enough evidence to say that we’ve determinately settled any of these puzzles one way or the other, in which case the constitutive rules for the application of quantifier and sortal concepts leave the answers to such puzzles necessarily undefined – ‘unanswerable’ in the semanticist’s strong sense, e.g. Thomasson (2009) – else English usage pushes us towards a position where, e.g., we accept all the premises in Statue-Lump argument and the English concepts turn out to deliver incoherent answers. By contrast, it’s stipulated of the Ontologese quantifier that it has this enhanced theoretical role and therefore must deliver *both* answers to standard quantificational and counting questions – ‘Are there Higgs Bosons?’, ‘How many cats on the table?’ – *and* paradigmatically ontological questions – ‘Is Clegg identical to Sturgeon, or are there two things here? Which, if either, is the original ship? Is there some further thing these simples compose?’ The former is a more anaemic theoretical role, but does enough to allow ordinary speaker to get by in their daily lives, while the latter is a more overarching theoretical role that allows one to build a richer, more systematic model of reality in quantificational terms (but is probably useless for ordering takeaway— even if the person taking your order is fluent in Ontologese).

It’s the very absence of these additional constraints on the use of the English ‘exists’ which explain why we needn’t worry too much about solving ontological puzzles when employing the English quantifier. As Thomasson (2009, p.451) writes, in full semanticist flow:

[...] if the application and coapplication conditions for our terms are just those embedded in the rules mastered by competent speakers, others may object that this simply isn’t enough to generate answers to all our metaphysical questions. This, I think, is true: the conditions for application and co-application – at least for most of our ordinary and common sense terms – are typically vague, and often highly incomplete.

I agree with Thomasson that there seems to be a strong possibility that neither English sortal terms, nor English quantificational terms, provide complete or determinate answers to many ontological questions, such as material object puzzles. On this assumption we needn’t worry about answering such questions using English terms as there are simply no good answers to be had. The concepts we employ are simply not geared towards

providing complete or coherent answers to such questions. Why? Because as a linguistic community we had no need to be able to deliver a satisfying set of answers to such questions, yet did have the pressing need to be able to talk about whether there is a tiger in front of us or how many apples are left on the tree. Any children or metaphysicians inclined to engage in such an activity are not abiding by the community standards.

On first encountering the puzzles, it occurs to us that we could provide answers, and reasoning from these answers to discover surprising things about the non-existence of tables, or the co-location of statues and lumps, or the persistence conditions of ships, seems like a perfectly straightforward thing to do. However, we soon realise that we're running far ahead of the usual work we put terms such as 'exists' or 'is identical with' to and coming into conflict with the intuitive, unspoken rules for applying such concepts whatever answer we give. Here we begin to apprehend the additional rigour and stipulations that would have to be built into our usual concepts in order to carry on the debate sensibly— that is, we come to have a grasp of Ontologese.

Now, the crucial point here is to distinguish our understanding of the Ontologese quantifier, as having this more comprehensive theoretical role, from our views on whether it has any theoretical utility - that is, whether the additional theoretical role is one that needs filling. The semanticist is realist enough to admit that whether there are subatomic particles or chairs according to the English sortal is a matter that is in some sense determined by the way the world is, rather than simply how we choose to speak. However, she thinks that the world doesn't fix the answer to paradigmatic ontological question - frame-level questions, in Thomasson's (2009) terms - at all. Rather, these are best a language relative matter - that is, we may speak the nihilist language or the universalist language, etc., as Hirsch would have it, but one is not picked out as privileged by the mind/language-independent world. Here the Ontologese speaker and the semanticist are employing the same refined sense of existence, privileged by its overarching theoretical role, and so understanding each other just fine: they disagree only over whether, once this revised concept is employed, the new questions that become answerable in this framework could be settled by the world in a way that more straightforwardly empirical existence claims are, or whether they could only be settled by our linguistic decisions. Hence, understanding that there are two sides to this debate is all that's needed to understand the ontologist when she says she wants to use a quantifier privileged *by the world*.

Meanwhile, to reiterate the possibility of *realism light* as conceived in the previous chapter, even the semanticist could concede, as Thomasson (2009) comes close to doing, that there is still pragmatic value to be had in employing a concept of existence that allows

these questions to be systematically answered, even if there's no worldly reason to prefer one systematic answer over any other: So, even if you don't think such questions could be settled by the world, you can still see the case for a privileged quantifier. While you might think an advantage of a deflationary picture of the debate is that we no longer have to answer awkward questions, because they are revealed to be pseudo-questions, this is not obviously the correct answer: It may be preferable to introduce the linguistic apparatus to enable us to ask more fine-grained ontological questions purely for pragmatic gain (just as the physicist often gains certain pragmatic advantages from acquiring more fine-grained mathematical representations of reality).

Finally, of the ontologist who does see the answers to such questions as both worthy and worldly, yet who takes herself to be speaking ordinary English: she is in a very good position to understand the Ontologese quantifier, as she has made a subtle and forgivable mistake about the quantifier in her native language in assuming that it is suited to play the theoretical role she wants it to play, when in fact it has too coarse a theoretical role. Her intuitions about how to apply the relevant vocabulary in English may lead her astray in all sorts of ways in her theorising, since she is simply not using the right tool for the job.¹⁷ But she can easily come to understand the meaning of the Ontologese term by continuing to use existence just how she always has – which is close to the Ontologese sense – but coming to realise that the linguistic intuitions of ordinary speakers aren't relevant to settling the questions she is trying to get purchase on.

2.3. *Can Earthlings Conceive of Electric Cats?*

One way to flesh out this picture is to draw a distinction between our concepts and our conceptions: our concepts being 'our mental vocabulary, the basic units of cognition'¹⁸, while our conceptions attach additional information to our concepts, helping to flesh out a more determinate mental picture. Turner (2013), in drawing this distinction, notes that are conceptions and concepts can come apart in the following way:

as Putnam [...] suggests, if all catlike creatures turn out to be robots controlled by Martians, then the meaning of our concept cat will cover these robots, but our conception will be radically mistaken.

Following a certain way of drawing the *concept-conception* distinction, which will be

¹⁷ Recall the previous chapter.

¹⁸ (Turner 2013)

useful for our purposes,¹⁹ it can be employed to explain how our concepts can survive radical change in our beliefs concerning the nature of things that fall under that concept.²⁰ In the Putnam example, we have a *conception* of cats as flesh and blood animals, which is suddenly replaced with a *conception* of cats as Martian automata: this enables us to meaningfully and correctly say that we have come to learn – much to our surprise – that cats are, and have always been, robots and not flesh and blood animals. If we are forced to replace our entire concept of *cat*, however, we’d have to semantically ascend in order to describe our discovery, saying something like: the things we have always refer to with the term ‘cat’ turned out not to be *cats* at all, but *robots*— so ‘cats’ does not and never has referred to *cats* (there are no *cats*).

At least some of the time, the former (non-semantically ascending) thing seems to be the right thing to say: We understand what a cat is perfectly well, we’ve just learned surprising new facts about cats. We don’t want to say that we never properly grasped the concept corresponding to our term ‘cat’ by failing to realise that everything falling under that concept is a robot. Cats are the meowing things we keep as pets with a penchant for tuna. If you understand that then you understand my crazy conspiracy theory about the Martians and how to the untrained (in advanced Martian cybernetics – veterinary school isn’t sufficient) eye they appear to be flesh-and-blood but they’re not. On the other hand, if I started raving to you about my theory that all squares appear to have four sides, but in fact have three, or that Tuesday is in fact an African elephant, not a period of time in a week, you will rightly judge me to have fundamentally misunderstood the concept of *square* or *Tuesday*: Most charitably I could be seen as suggesting that the terms ‘square’ and ‘Tuesday’ don’t correspond to the concepts we ordinarily take them to correspond to.

Concepts are what we grasp when we come to understand a meaning of a term. By contrast, *conceptions* are more determinate – more fleshed out – mental representations that may guide us in thinking about a thing, but are inessential to our understanding of the terms in our language or having basic thoughts about the thing *qua* that thing: our conception of cats as flesh-and-blood things may play a big role in our thinking about cats, but the concept *cat* can be grasped without any understanding of biology— animated clay

¹⁹ The literature is in fact very divided on how to employ the distinction, sometimes using the terms interchangeably, but I don’t think too much hangs on this collective indecision.

²⁰ I’m going to italicise all uses of ‘*conception*’ so as to avoid confusion with ‘concept’ – apologies to readers whose aesthetic sensibilities are wounded by this stylistic choice, or who find it more distracting than useful.

cats, robot cats, God-like spirit cats are all perfectly coherent notions. We can have a concept of *flesh-and-blood cat* or *God-like cat* or *robocat* that does determine their internal constitution, but plausibly at least this isn't the concept picked out by the English word 'cat'. The cat concept is much less determinate, even though most individual speakers have a much more fleshed out *conception* of cats.²¹

My response to Korman, then, on the concept-*conception* picture, is that while the ordinary English speaker doesn't have a very determinate *conception* of existence, the Ontologese speaker necessarily does. Thus, the English speaker is analogous to the pet owner who doesn't give a damn whether her pet cat is fleshy or electric – she'll still love him and, more to the point, still consider him a cat. In the OS case, she doesn't have any pre-theoretic views on whether she speaks a pedurandist language or an endurantist language; a restrictivist language or universalist language; but just gets on with making and evaluating everyday existence claims. While the Ontologese speaker is like the bounty hunter charged with hunting down and "retiring" robocats when the Martian plot is first suspected but the full extent of it is unknown: for him true cats are necessarily flesh-and-blood. In the OS case, the Ontologese speaker has a conception of existence which requires that paradigm ontological questions are answerable language-independently.

The pet owner and the bounty hunter share the same *concept* of cat, here, but the bounty-hunter has a richer conception— since, not realising that *everything* we have ever called a cat is a robot, he had good reason due to his job and partial knowledge of the Martian plot to explicitly rule out that robots can really be cats. Hence, the pet owner and the bounty hunter understand each other fine when they use 'cat' since both grasp the same concept, despite differing conceptions. However, we can imagine the bounty-hunter stipulating that now on by 'cat' he means flesh-and-blood cat. Since he has merely introduced a more determinate concept of cat, to fit with his more fine-grained conception of cat, all he need do to get the pet owner to grasp the new concept is to fill her in on the plot and why the distinction is important, even if the pet owner is sceptical

²¹ N.B. My *conception* of a cat might well include that being a flesh-and-blood animal in an *essential* property of cats – that if something is a cat it is necessarily a flesh-and-blood animal. But this idiosyncratic metaphysical view of the modal profile of cats plays no indispensable role in my grasping what 'cat' means in English and therefore doesn't mean that I am working with a different concept of cat from someone who doesn't believe cats are essentially flesh-and-blood. Given that all actual cats are flesh and blood, it may be plausible to hold that all possible cats are flesh and blood – that a given cat could not have been anything other than a flesh-and blood being – but someone with this view can still make sense of robot cats being cats by recognising that their views on the modal features of cats may well rest on a false presupposition about their actual properties.

about the reality of the plot.

The ontologist and the ordinary speaker, meanwhile, start off with the shared English concept of existence, and so understand each other perfectly well when using the English term ‘exists’, even though the ontologist has in mind a more fine-grained conception of existence in terms of its theoretical role. The ontologist can then start doing ontology, in which case she will now use ‘exist’ in the Ontologese way that corresponds with her more fine grained conception: all that the ontologist need do in order to get an English speaker to understand the Ontologese term is to get them to grasp the more determinate *conception* of her existence concept, in order to thereby grasp the Ontologese concept of existence, which *builds in* this more fine grained conception. The fact that we start with a *shared concept* of existence – the English concept – and that the Ontologese concept is identical to a more determinate conception of that shared concept means that grasping the Ontologese concept is a relatively straightforward matter. Certainly more straightforward than trying to grasp an entirely new concept that is not merely a refinement of an existing concept— such as eight-foot tall Tuesdays (very plausibly not a candidate refinement of our existing concept of Tuesdays). There’s really is no more to grasping the Ontologese concept than grasping the additional theoretical role.

2.4. Hofweber – more unintelligibility

The defence of OS offered in the previous subsection can be deployed generally against those who doubt the intelligibility of Ontologese terms. In this section, we'll see how this helps us stave off a challenge to the assumption that ontological questions are non-trivial. Hofweber (2009) argues that the ambitions of ontology must be seriously reined in, to the extent that the task of the ontologist becomes – almost exclusively – to discover which existence questions belong properly to the 'domain of metaphysics'. Any question belonging to this domain will, according to Hofweber, be invariably answered in the negative. Thus, a metaphysician is never justified in making a positive ontological claims when wearing her “metaphysical hat” – if she wishes to posit the existence of an entity she do so on the grounds of common sense or scientific evidence.

I think this conclusion is too strong. I won't construct the entirety of Hofweber's argument here, as I think it can be shown that he makes a wrong step early on. Namely, Hofweber rejects introducing a primitive concept of *real* existence or Ontologese existence as a way to do 'ambitious' metaphysics, as he sees this as turning ontology into an 'esoteric' discipline. However, once it is understood that the ordinary person can easily come to a conceptual grasp of what it is to “*really* exist” or “exist in the fundamental sense” charges of esotericism against those views are hard to sustain.

Hofweber warns metaphysicians to avoid taking an 'immodest' attitude to the sciences. That is, being prepared to overrule scientific claims on the basis of metaphysical enquiry, or holding the successful of scientific theories as provisional on vindication by metaphysics. The best way to achieve modesty, in his opinion, is to follow the example of the higher-level science, such as the biological sciences and the social sciences, and carve out a distinctive domain of ontological questions that - although they may be indirectly impinged upon by the sciences - are properly and directly answered by metaphysics.

Perhaps the most popular way of achieving this in the contemporary debate, notes Hofweber, is to introduce a primitive concept of reality, such that questions of the form "what exists?" fall into the distinctively metaphysical domain if the sense of "exist" in such sentences is the *real* or OS sense.²² Thus, the question "Do numbers exist?" is properly in the domain of ontology if we mean "Do numbers *really* exist?". Hofweber rejects this method of demarcation on the grounds that as the concepts of reality are theoretical primitives introduced by philosophers, reliance on them turns ontology into an obscure and inaccessible practice. This goes against Hofweber's insistence that metaphysics should be an 'egalitarian' discipline, in which '[t]he questions are accessible to all, even though not everyone cares equally about finding an answer to them.' (p.266) Having dispensed with this way of answering the domain question, he proceeds to offer his own solution as the only viable alternative: This has the rather large downside - from the perspective of those of us who are optimistic about the epistemic value of ontological enquiry - of making all questions in the domain of ontology trivial.

Suppose we grant Hofweber's stipulation that metaphysics be 'egalitarian', in his sense. Why think that a grasp of the concept of real existence is not 'accessible to all'? The fact that real existence is a theoretical primitive cannot, in itself, be taken to entail an unacceptable degree of inaccessibility. All theories - folk, physical, and social scientific - have their primitives. To single out metaphysical enquiry and deny it the use of any primitive concepts would be to hold the discipline to an implausibly high, if not downright dubious, standard of accessibility. Perhaps the trouble is supposed to be that the primitive is a peculiarly "metaphysical" one, in some pejorative sense; that it is being manufactured artificially by philosophers, rather than latching on to a primitive found in our ordinary inventory concepts? This charge, which mirrors Korman's, cannot be sustained given the central claim of the last subsection:

²² Of course this is an infelicitous characterisation for the proponent of OS who sees her quantifier as unrestricted, but we shan't dwell.

All you really have to do to grasp real existence, or understand ‘exists’ Ontologese sense, is to *start caring* about answers to ontological questions, such that you grasp the more elaborate theoretical role – the fuller *conception* – of your ordinary concept of existence that the ontologist has in mind. So, the criteria of being non-esoteric is met: '[t]he questions are accessible to all, even though not everyone cares equally about finding an answer to them.' It's the widespread non-caring that's responsible for our anaemic concept of existence in the first place, but by simply trying to find answers to difficult ontological questions using this shared concept you can easily come to grasp the more heavyweight conception of existence that enables you to understand ‘exists’ in Ontologese— *seek, and ye shall find*. Charges of esotericism only have merit if this task proves inordinately difficult or achievable only by a select few. In the case of understanding questions about what the metaphysician means by '*really* exists', however, I think the concept should be readily understandable to anyone who cares enough to enquire. This should be enough to satisfy any plausible version of Hofweber's accessibility constraint.

So, if Hofweber's (2009) argument causes difficulties for anyone, it is for those like Korman who explicitly advocate doing ontology by asking existence questions without employing the concept of real existence. They must either endorse the trivialization of metaphysical questions, find a convincing way of avoiding Hofweber's background assumptions, or provide a novel answer to the domain question. Meanwhile, the fundamental language ontologist can safely get on with investigating what *really* exists.

2.5. External World Scepticism (more Moore)

I've defended OS by showing how it is that we're capable of understanding its existentially quantified claims, and explaining why ontological debates would be best conducted in that language. A little earlier we saw how adopting OS could impact favourably on the first-order debate: We resolved an apparently difficult puzzle about coincident statues and lumps by denying that statues really existed, and adopting principled agnosticism about the English puzzles. Korman challenges both of these moves: Why is it relevant whether or not there are really statues? And why do you think you think you are in any better a position than everyone else with no sure-fire answer to the English puzzle? With regard to the first challenge, we've seen that if we care at all about having a perspicuous representation of reality we should care about what really exists, since ordinary existence simply can't provide us with a complete or coherent picture. With respect to the second, we've seen that trying to insist on a determinate answer to ontological questions in ordinary language is potentially to misunderstand the theoretical

role of quantificational concepts in English.

Now, you might well have a different way of resolving statue-lump puzzles - with or without the help of this sort of metaontological strategy - which you take to be satisfactory. The point is that this approach puts a new option on the table which seems to preserve a lot of the intuitions other have got into trouble trying to hang onto, and it should be fairly easy to see how such an approach might be applied elsewhere.

In this subsection, I want to focus on how OS is additionally suited to making Moorean answers to sceptical puzzles more principled. Mooreanism, as we've seen, puts excessive weight on our common sense existence claims when it comes to answering ontological questions. Moorean positions can also be set up in opposition to radical external world scepticism. The sceptic says that we can't know anything about the external world, since we can never know that we aren't being systematically misled by our senses. For instance, even the common sense assertion that you have hands might be false, since you might be being deceived by an all powerful Cartesian demon or an evil scientist. The Moorean responds by insisting that we '*know*, with certainty' (Moore 1925, p.32) we have hands. It's just common sense. No sceptical hypotheses could ever be entertained that contradicted such an obvious truth as our having hands.

Previously,²³ we've been interested in those instances where the Moorean decides to get involved in revisionary ontological debates - such as the composition debate. The Moorean treats the Cartesian sceptic much as she treats the revisionary ontologist. When the nihilist says that there aren't any tables or chairs, say, just "simples arranged chairwise", the Moorean takes this to be just as obviously wrong as the sceptic's mooting of the possibility that we might not have hands and that there is just a systematically produced illusion. Of course we hands! Of course there are tables and chairs! I agree with the Moorean that it's obviously true (or at least factual, modulo the discussion in chapter one) that chairs and tables exist, but I don't think it follows from this that the nihilist is wrong to offer a revisionary ontology, as I've said. It's obvious that chairs exist, since it's obvious that the worldly conditions needed for English speakers to assent to the sentence "chairs exist" are met. However, it's not clear that chairs exist in the privileged Ontologese sense that provides perspicuous ontological commitments.

What we capture, with this way of explaining the debate, is the nihilist's insistence that ordinary folk are mistaken about the way things ultimately are: When folk look at a region

²³ §1.2

of the world that gives them chair-ish perceptions, they imagine that there is a single thing there – a chair – which exists in the way that fits with their concept of being. This is right. But the nihilist Ontologist, employing her more fine-grained concept, notes that there are lots of things in vicinity which do exist in the way conceived of by the folk, but none of them have all of the properties of the chair – instead, they act in concert to provide onlookers with chair-ish sensations. We also capture the Moorean's insistence that it's obviously true that there are chairs, hands, etc. For it's obvious that the folk speak in such a way that the truth (factuality) conditions of “there is a chair” require that reality have certain specific qualitative features, and that these conditions are evidently complied with. Thus, the nihilist who employs OS is set free from this suspicion that she is either making false claims about our language, or advocating a reform of our language. For it would be implausible to claim that we don't speak in such a way that it's obvious that there are chairs, and since ordinary language suits our purposes a categorical claim that we *should* speak as if there are no chairs looks suspect. But an OS ontologist makes no such claim – she is just making the conditional claim that *if* we want our ordinary existential commitments to match our ontological commitments, we should speak as if composite objects don't exist.

So, by casting nihilism as a thesis about what there is *according to Ontologese*, we capture the spirit of the naïve nihilist's claim that ordinary folk are wrong about what there is, whilst insulating nihilism from Moorean objections.²⁴ A radically strong form of Mooreanism could of course insist that we have certain knowledge of the fact that there are *really* tables and chairs, but in the absence of further argument, I can't find any reason to consider this plausible. We can see how the OS ontologist provides a moderating voice when it comes to another response to scepticism. Putnam (1981) offers an argument against a certain brand of sceptical scenarios often called 'Brain in a Vat' cases. Our Brain in the Vat sceptic thinks that maybe there are no chairs or tables because (for instance) we might be in some sort of simulated reality – disembodied brains in a laboratory that contains no tables or chairs, in a world that doesn't contain any tables or chairs either.

Putnam responds by claiming that such scepticism is easily refutable. For if there were a group of people (brains) raised in a vat, they'd learn to speak a different language from those outside the vat, since the meaning of their words would be fixed by their experiences of the simulation, rather than by the world outside the simulation. When

²⁴ The error-theorist still has the responses available set out in §1.2 to Mooreanism, of course. But the OS response is more comprehensive, satisfying and secure from doubt.

those outside the vat say 'look at that tree', they mean to point out a living leafy thing; whereas when the vat-dwellers say 'look at that tree!' they mean to point out a leafy-looking element of the simulation (for the simulation is all they've ever known). If they speak ordinary English outside the vat, we might say those inside the vat speak vat-English. Since 'tree' in vat-English means something entirely different from 'tree' in ordinary English, let's keep track of this by writing 'tree' in vat-English as 'tree*'.²⁵

Since the Great Furniture and Fixings Fire of 1863, those outside the vat have spoken truly when they've said that there aren't any tables any more. In ordinary English it's true to say that there aren't any tables. Does this mean that those inside the vat are in a bad epistemic situation? After all, they look at the chair-ish elements of the simulation and say "Lo! There are so many chairs around these days." Taking on board Putnam's response to the sceptic, we should say that they are not mistaken at all. For the sentence they have uttered is one in vat-English: 'Lo! There are so many chairs* around these days.' Since the vat-English word 'chair*' refers to chair-ish elements of the simulation, the vat-dwellers are quite correct to say there are lots of chairs* around these days, as the simulation is indeed littered by chair-ish elements. Since the vat-dwellers couldn't possibly have the vocabulary to make claims about 'chairs' proper, none of their existentially committing statements will ever go wrong due to a mismatch between vat-internal reality and vat-external reality.²⁵

Thus, according to Putnam, any sceptical doubt based on the epistemic possibility of our being 'brains in vats' is misguided. If we say "there are chairs" either we are outside the vat interacting with, and referring to, vat-external chairs, or we are inside the vat. If we are inside the vat, we speak vat-English, and so make the claim that "there are chairs*", making this claim on strength of our interaction with chair* vat elements, and successfully referring to chairs*. Basically, whatever our situation - in the absence of ordinary error - our existentially committing claims are factually secure: scepticism (or at least this breed of it) stands defeated. At least, this is Putnam's contention. As in the Moorean case, however, there's something about this conclusion that seems too strong. You might well think that the mere epistemic possibility of such sceptical scenarios does not rule out your having defeasible knowledge of your hands; but it seems odd to conclude that, vat-dweller or not, we cannot *possibly* be mistaken in our ordinary existential claims.

²⁵ If a population got suddenly transplanted from our world to the vat then they *would* on this view be wrong to continue to point at their perceptions of trees and say "Look, tree!" despite everything being phenomenologically the same. That seems like a good result. Presumably, over generations of vat dwelling, meanings of terms would shift, however.

Our answer here should be much the same as given to the Moorean. We should concede that our ordinary existential claims are certain – in so far as doubt from sources such as Brain-in-a-Vat hypotheses is concerned. However, this should be put down to the necessary plasticity in the factuality conditions of our ordinary existential claims. For when our vat-dwelling society first started out the most eligible or perspicuous referents of their chair-talk would be the vat, or phenomenal elements of it (or perhaps its physical microstructure): for when a vat dweller points and says “there's a chair*!” the vat-elements are the only things in the vicinity that exist in the Ontologese sense of existence. Yet, as the vat-dwellers' chair-talk becomes more sophisticated, it will become clear to laboratory observers that “chair*” is not intended to refer to an individual element of the vat: statements such as “I can see exactly two chairs* right now”, “this is the same chair* as I saw yesterday”, and so forth, will make this apparent. They would not be remotely true if “chair*” was meant to refer to an individual (transitory) vat-element. One is forced to conclude that either that vat-dwellers are in the habit of speaking unfactually, or that the factuality conditions of “there is a chair*” require a very complex arrangement of vat elements, such that the vat is in a state which delivers to the speaker the perception of a four-legged piece of furniture. The arrangement of vat-elements required to produce such a sensation on any given occasion will no doubt differ wildly, and it may well be impossible to provide any systematic translation of chair-talk into vat-element talk. Yet, a principle of charity demands that we interpret them in such a way as to make their chair-talk true (factual). As such, although the truth of chair-talk will supervene on the truth of certain statements about vat elements, statements such as “there is a chair*” may have no systematic, finite paraphrase into statements referring only to vat elements.

So, does this mean that the evil scientists have not managed to deceive the vat-dwellers about what exists, since all their ordinary existential claims are true? It would be strange if such elaborate mischief had really come to naught – and we can see that it has not. For, by imposing their computer generated veil between the vat-dwellers and reality the evil scientists have ensured that the vat-dwellers' idea of what exists is a very poor fit with how things really are. There is really nothing that fits with vat-dwellers' conception of a chair* besides passing phenomenal elements. It has made their talk less metaphysically perspicuous, since they are quantifying over objects that do not carry ontological commitments, since they don't exist in the Ontologese sense. However, despite this 'veil' between fundamental reality and their own perceptions, it is not impossible for vat-dwelling metaphysicians to engage in productive ontological enquiry. For they are free to ask 'what is there?' questions stipulating that they are employing a fundamental quantifier. Using the same theory choice criteria that we would from outside the vat, they may well produce contributions to the field that we would do well to read.

The lesson we should draw is twofold: Firstly, even in paradigmatically poor epistemic situations such as evil-demon and brain-in-a-vat scenarios, our ordinary talk gets to be true and meaningful. Our statements will still require the backing of certain worldly conditions in order to come out true. Yet, precisely the poorness of these epistemic situations should lead us to the realization that our ability to make true and meaningful existence claims far from guarantees that we are making *perspicuous* existence claims. Secondly, that even in these pragmatically poor epistemic situations, successful ontology investigation would be by no means impossible. So, given that we tend to think we're in a much better epistemic situation than either of the above, we shouldn't hang onto our common sense ontologies out of fear that if we abandon them we'll have "nothing else to go on". Ontological enquiry shouldn't end with our collective prejudices; it should start with setting them aside.

3. Escaping the Fundament

In the chapter I discuss the status of non-fundamental objects and why the ontologist who adopts the OS strategy shouldn't worry about being able to recover the existential claims of non-fundamental theories in her theory. §3.1 Tries to clear up a potential misunderstanding about fundamentality which threatens to render nihilism trivial on OS, while exploring some rival accounts of (non-)fundamentality. §3.2 I give an account of the status of relatively natural or partially joint carving existential discourse, such as Chemistry, which argues that non-fundamental ontological debates can be substantive, but are autonomous from the project of OS. §3.3 Examines the status of social ontology a putative example of substantive but highly unnatural or non-joint carving discourse and attempts to meet an objection by Elizabeth Barnes that the framework OS is unable to accommodate such projects or is unjustifiably dismissive of them.

3.1. *Dependence, Fundamentality and Reality*

We've seen that we can make sense of shifting the composition debate from ordinary English to Ontologese. However, it's worthwhile to consider Korman's diagnosis of what is - mistakenly, in his view - motivating OS strategies. It will, as a side-effect, provide an opportunity to discuss some rival metaontological positions which don't see the interesting ontological questions as about being about what exist but do not make appeal to a the fundamental language either. Korman points out that in this discussion (almost) everyone around the (putative!) table agrees that composite objects are not *fundamental*. By fundamental, goes the thought, we could mean a variety of things: a fundamental object is the sort of object that appears in one's fundamental physical theories, they are the basic objects on which everything else depends - *all God needed to create in order to bring about the existence of everything else*.¹

In this sense everyone can agree - says Korman - that only simples are fundamental while the things they compose, like tables and chairs are non-fundamental.² All the fundamental language ontologist is in effect, by the lights of the ordinary speaker, restricting her quantifier to quantify over all and only the fundamental things - though by

¹ See Barnes (2012).

² Again, provided the world is not gunky and thus has no smallest parts.

switching to a language with a different meaning for the unrestricted quantifier her quantifier will not count as restricted. The main problem here, if the OS strategy is to employ an unrestricted quantifier in which to carry out ontological debate which has the same domain as the ordinary language quantifier restricted to only those things that are fundamental, even if there's not anything internally problematic with OS strategies, is that it threatens to turn what were apparently interesting ontological debates boring and one-sided. That's because everyone already agrees that only simples are fundamental, so everyone will agree that in this new language of course nihilism is true, but this only tells us what we already knew: composite objects aren't fundamental.

The only reason we would have for believing that composite objects were *fundamental* is if they possessed emergent properties.³ But on the assumption that there are no such things as emergent properties – an assumption widely accepted in the debates – we're never going to have any good reason to think that composite objects are fundamental, so the debate is purely about whether we should only quantify over fundamental objects, but the substantive first-order question of which objects are fundamental is already agreed on by both sides. The upshot then is that what appeared to be an interesting first-order question about material objects turns out to be a high level meta-question about whether we should conduct debates about such issues in Ontologese or not— once we settle the metaontology all the first-order answers just fall out.

Now, I do as a matter of fact believe that *if* there are no emergent properties – and I think that's a big if – then we should say that, at least in the realm of concreta, only the 'fundamental things' in the sense above really exist. And, yes, this is informed by my metaontological views. However, I don't think that the adoption of a particular metaontological view *settles* the composition debate or makes it uninteresting: Rather, by adopting what I am advocating as the correct metaontological framework one removes certain systematic epistemically conservative biases in the old system that lead us to give undue credence to more ontologically liberal ontologies.⁴ The adoption of the new metaontology shows how many of the existing arguments for belief in composite objects rests on unfounded metaontological assumptions, while the reasons for being nihilist remain. As such, as the composition debate stands, I think nihilism is by far and away the leading first-order position where it might have been seen under more authentically Quinean metaontologies as lagging behind. What the believer in composition needs is

³ See Part II of this thesis.

⁴ See §4.

new arguments – or to reinvigorate older ones. I think that the prospect of there being ontologically emergent properties in our world are better than Korman supposes: rather, I believe arguments from emergent properties currently represent the best hope for the believer in composition to present a strong challenge to nihilism. Which is why I spend the latter half of this thesis attempting to defend nihilism against arguments from emergence and – spoilers! – why even though I believe the nihilist can adequately respond to the argument for emergence, I still have much higher credence in the truth of nihilism on the assumption that there are no emergent properties. So, if anything, I think adoption of the proposed framework serves to reinvigorate the composition debate rather than render it uninteresting.

However, it's also entirely misleading to see the metaontological framework itself as somehow forcing the conclusion that only the fundamental things really exist. This is because the aim of the OS strategy is not to quantify over only the fundamental things: anything that's fundamental must be quantified over for sure, then whether a certain kind of non-fundamental object exists must be considered on a case by case basis. We can't have a theory without the fundamental objects: the place in our theory for non-fundamental things can, in principle, be earned. As other proponents of fundamental language strategies have been at pains to point out, to exist in a *fundamental sense* (to be real) does not by itself entail being a fundamental kind of object.

To compare the OS to other strategies on the table: Schaffer (2010) tells us that existence questions are easy and should not be the primary concern of the ontologist – so far, then, we're in agreement. However, he does not offer us any direct means to distinguish existence from real existence. Instead, he thinks we should focus on questions of what grounds what, of which objects are ontologically prior to others.⁵ On the standard view, then (contra Schaffer's priority monism) the simples that compose a chair are ontologically prior to the chair and furthermore nothing is ontologically prior to the simples – the simples are ultimately prior, the real, or fundamental. So, as ontologists, we:

We should be interested in discovering what's ultimately prior.

We should be interested in discovering the relations of relative priority between everything that exists.

⁵ Compare Fine (2001, 2009) who believes that instead of priority relations holding between entities that a relation of grounding holds between facts.

On the first point, it looks like what the ontologist's job is reduced to is just the task of discovering what is fundamental. Even if it turns out not to be as uninteresting as Korman supposes: Schaffer puts forward a case for Priority Monism, the view that it is not parts that are prior to wholes but wholes that are prior to their parts, and therefore that is wholes (or given certain plausible further assumptions, the largest whole, the World) that are fundamental while the simple parts (if anything is mereologically simple) are derivative. But this priority question is a long way from the questions of which answer to the Special Composition Question is correct: if nihilism is equated with the claim that there are no *fundamental* composite objects, this just amounts to the priority question – but few defenders of the existence of composite objects in the traditional debate would have wanted to reject the claim that a chair's simple parts are fundamental, or at least that the chair is ontologically dependent on these parts. Regardless of the direction of priority, can we not still have a debate about whether there really are any chairs or not? But, not having recourse to an Ontologese quantifier it seems on this framework we're left with no other option than to make the debate about what exists according to the ordinary language quantifier – with Schaffer having made quite clear that everyday objects of our experience can obviously be taken to exist for Moorean-type reasons (this is not an existential quantifier pulled strongly by reference magnetism). This renders being a nihilist not really a credible option, since one has to violate Moorean constraints.

Is to exist in Schaffer's framework not to exist in a substantive or joint carving sense? Hard to say. It seems that Schaffer, like the Quinean, must say that existence in ordinary language is existence simpliciter: So there's a sense in which it seems on Schaffer's view real existence becomes objectionably cheap: that mere Moorean certainty is enough to ensure the *real* existence of an object. Conversely, however, Schaffer believes that for the purposes of weighing up competing theories by their ontological commitments, only the fundamental objects should count as ontological commitments.⁶ Now, if 'ontological commitments' is read as referring merely to ontological costs and not as synonymous with real existence, Schaffer's statement is not incompatible with derivative objects having real being too— but *derivative* real being that doesn't count as an additional ontological cost over and above the things it is dependent on. But it also suggests a second strategy on which only the fundamental things have real being: I don't find this suggestion as problematic, since it is more or less equivalent to OS, in that you could broadly endorse Schaffer's metaontology while also shifting debates about what is absolutely prior into Ontologese where only these objects – the ontologically committing things – are

⁶ Schaffer (2010).

quantified over. But the remaining downside, despite the rough equivalence to the view I endorse, is that this framework really would force us to accept at the outset that only the fundamental things really exist rather than leaving room for the real non-fundamental.

On the second point, this requires us to investigate priority relations which hold between fundamental things and non-fundamental things. This would be a distinctly odd thing to do if we don't take one end of the relata to have real existence, since we would be trying to posit a real relation where one end bottoms out in something that is not real.⁷ That is, unless talk of a relation here is really shorthand for the semantic fact that e.g. fundamental things like simples *make true* (or at least play their part in making true) ordinary language sentences about e.g. tables and chairs.⁸ But this would make such questions about how our ontology matches up with our ordinary talk – that is, a question of metaphysical semantics⁹ – rather than ontological questions proper. However, if we take the view that both ends of the relata do have real existence, then there does seem to be a fruitful ontological questions to be asked about what depends on what. If such dependences between real beings exist, is it not the ontologist's job to investigate these?

I think that my proposed metaontological framework can allow for an ontology with such a priority structure. As we'll explore in greater detail in the second part of this thesis, we can separate out three aspects of ontological priority views like Schaffer's:¹⁰

1. The having of perfectly natural or fundamental properties¹¹ – call this being a fundamental entity.
2. Being ontologically independent of any other entity – that is, not existing because of another entity.
3. Having real being.

To really exist is to satisfy 3. But I argue that neither 1 nor 2 need be satisfied just because 3 is satisfied. Furthermore, while satisfying 1 is plausibly sufficient but not necessary for the satisfaction of 3, 2 is neither necessary or sufficient for the satisfaction of either 1 or 3. One upshot of this is that one can be a dependent entity yet still have real

⁷ A violation of Sider's (2013) purity constraint.

⁸ As in Cameron (2010).

⁹ See Sider (2011).

¹⁰ See Barnes (2012) for a defence of separating out dependence from fundamentality.

¹¹ Or being an entity that plays a similar role in a nonminimally acceptable theory.

being— as such, the proposed metaontological framework can handle there being existential dependence relations between real beings (we may or may not be endowed with fundamental properties). No plausible purity condition is violated since the relation is holding between two real entities, even if one end of the relata is dependent and possibly non-fundamental. Now, whether to admit such dependencies into one's theory of reality hangs on whether there's any need for expressive or explanatory power that such a move might grant.

For instance, an OS theorist who wishes to avoid being deflationary about the composition debate should admit at least the coherence of postulating such dependencies. It is fine for the nihilist to say that the existence of tables depends on the real existence of simples, since it is real simples that are the truth makers for ordinary language statements about tables. But the universalist or the restricted composition theorist must say that tables have real existence, but most will still want to claim that they depend for their existence on microphysical simples and that its properties are not fundamental but inherited from the properties of its microphysical simple parts— so in this case it's not enough to fall back purely on a semantic story.¹² Of course, if you don't think composite objects are even possible, you don't have to accept any theoretical machinery concerning ontological dependence into your own theory— but at a metaontological level at least you should leave conceptual room for it. In a similar vein, we'll see in Part II that perhaps the best way of making sense of ontological emergence is as real, fundamental entities that are ontologically dependent on other real, fundamental entities.¹³ Even if we don't want ontological emergence in our own first order theories our metaontology should leave conceptual room for this live epistemic possibility. Separating out really existing from being a fundamental entity and from being ontologically dependent allows us to do this in a way that a metaontology based purely on priority relations does not: While the possibilities that such a separation allows are live options in first-order ontological debates this makes priority ontologies at risk of being objectionably non-neutral towards these positions.

¹² This is not to say that ontological dependence itself need be real relation as in more ontology: all this is committed to as things stand is that in reality x depends on y ; we can remain agonistic on whether 'depends' is joint carving or corresponds additional ontology.

¹³ See also Barnes (2012).

3.2. *Unreal and Imperfectly Natural Metaphysics (of A-Level Chemistry)*

Can there be any metaphysics of the unreal – of objects that are of interest to us, but do not fall in the domain of the Ontologese quantifier? This is probably a particularly pressing question asked within the particular metaontological framework I am proposing, especially when also considering spare ontologies such as nihilism. Let's consider the entities of chemistry to illustrate my point: Compounds, molecules... even atoms... These are not the sorts of entities that appear in our *fundamental* physical theories, and they would appear to be composed of smaller entities— molecules, atoms, subatomic particles respectively.¹⁴ As such, the nihilist will not believe in their real existence, instead quantifying only over their ultimate parts. Is the nihilist, by saying that there aren't really any compounds or molecules relegating such objects to the status of glints or unicorns? More worryingly, perhaps, is the nihilist saying there can't be any substantive metaphysical debates about, say, the exact conditions under which two hydrogen atoms and a helium atom come together to compose a water molecule?

While few if any would defend the position that the objects of chemistry have perfectly natural properties, it seems equally indefensible to say that there is not some sense in which water molecules and helium atoms and chemical compounds are natural objects— carved out by the world rather than merely our conceptual scheme. In expressing this thought, one might have recourse to Lewis'¹⁵ relative naturalness: The properties of entities dealt with by chemists may not be perfectly natural – such as, perhaps, those dealt with by particle physicists – but they are still a lot *more* natural than those dealt with by, say, economics or psychology and far more natural still than anything mentioned in discourse about furniture or alcoholic drinks. As such, the entities they deal with, and the questions they ask about when something is an F or when two Fs compose a G, and so on.

On the standard neo-Quinean account, we should expect that the nihilist has provided us with systematic paraphrases for the eliminated chemical entities.¹⁶ Now we can proceed by simply substituting talk of helium atoms with talk of electrons, protons and

¹⁴ Electrons might turn out to be fundamental— so perhaps chemistry does deal in *some* fundamental entities.

¹⁵ E.g. Lewis (1986*b*: 61).

¹⁶ This stands us in good stead if we want to assure ourselves of Chemistry's high relative naturalness using the standard Lewisian recipe of length of definitions.

neutrons¹⁷ arranged helium-atom-wise. Further, there are legitimate, substantive-seeming questions that might be asked about what it takes for there to be particles arranged helium atom-wise, or particles-arranged-atom-wise-arranged-water-molecule-wise, and so on. While in Schaffer-type entity grounding metaontologies we can at least ask what more fundamental entities ground less fundamental entities and what less fundamental entities are grounded by the entities in question.

Yet I actively shun the need to provide any systematic paraphrase scheme or account of what grounds what: In a loose sense, I'm happy to provide a paraphrase for helium by, say, replacing my talk of a particular given helium atom with talk of some individual subatomic particles, their properties and relative locations. But there's no systematic paraphrase scheme of, for instance, of the form:

Replace all instances of 'there is a helium atom' with 'there are some particles arranged helium-wise'.

Replace all instances of 'there is a water molecule' with 'there are some particles arranged hydrogen-wise and some particles arranged oxygen-wise arranger water-wise.'

Etc., etc..

As we'll see in the next chapter, this way lies spirally ideological complication and a re-arrising of the sorts of puzzles the nihilist sets out to solve: In one's fundamental theory, it just replaces Fs with things-arranged-F-wise, so is hardly in the spirit of mereological-atoms-in-the-void eliminativism. Really, there are no mereological atoms satisfying plural *arrangement* predicates, just atoms with the properties and locations they each have (any plural arrangement predicates they may or not satisfy will not be part of the fundamental ideology of the theory).

One could perhaps seek to further paraphrase away talk of F-wise-arrangement. But a systematic schema of this sort is highly unlikely to do the job:¹⁸

Replace all instances of 'water molecule' with '[two hydrogen atoms] and [an Oxygen atom] in [close proximity] under [such and such... and such and such... and such and such... etc.... conditions...]

Etc.

In any case, it would be next to impossible to go about constructing one and the result

¹⁷ Or *their* ultimate part, anyway.

¹⁸ Terms in [square brackets] to be paraphrased away first.

would be so arbitrary seeming that it would be of dubious theoretical utility. We would gain little beyond the previous admission that whenever we would otherwise have described a situation with reference to the existence, properties and locations of atoms and molecules, we should refer instead only to the properties, locations and existence of simple particles.¹⁹ This hardly invites systematic replacement of terms like 'water molecule' in chemical theory with such a so-called paraphrase: It's rather an admission of the elimination of any vestige of systematic molecule talk from our theory.

On the (admittedly both questionable) assumptions that I can help myself to a physics first epistemology and that fundamental physics has an ontology of particles, I have sufficient independent reason to believe in an ontology of subatomic particles without needing to check that speaking in this way has enough expressive power to 'recover' the entities mentioned in the false but fruitful generalisations of higher-level or special sciences. Meanwhile, since the ontology of my favourite special science is independently fruitful, there's no pressing reason to get hung up on whether the entities of that science can be grounded somehow in the entities of some putatively lower 'level'— I already suppose that these entities aren't *real*²⁰ so there's no need to have some sort of existential crisis on behalf of your theoretical entities.

Unlike Sider, I don't require that for existential debates to be *substantive* - to be about more than just our 'conceptual schemes' - that they are carried out in perfectly natural, perfect joint carving terms, or else fulfil some opaque criterion of high relative naturalness or partial joint-carving.²¹ In my view, there can be substantive metaphysical debates carried out in highly non-joint carving terms: they don't automatically lack substantivity because of this - i.e., it doesn't immediately follow that the debate is defective or that there is no correct answer given by the world - they just lack reality. That is, the terms of such a debate will not ultimately deliver you a fully perspicuous representation of the world - but that doesn't mean that the world does not provide a uniquely correct representation as the 'answer' to that debate. I think the best way to get a handle on this is to draw on hypothetical/categorical distinction: We can ask how we can best represent the world *given certain constraints or a certain project imposed by us*. In which case we will be holding up a distorted mirror to the world: The world - not us - will then determine

¹⁹ Perhaps we could also try to re-paraphrase the paraphrase, but if regress doesn't lie this way then madness surely does.

²⁰ Modulo discussion in §3.3.

²¹ Sider (2011, p.131)

the distorted reflection we get back. Or we can just ask how we should represent the world, *simpliciter*— well, freed of any other constraint we should represent the world how it *really is*.

The quantifier of chemistry discourse is the best one to use if you want to do chemistry well. Or if not the best, it certainly constitutes a local maxima: Do chemistry with a quantifier that is more or less the same but casually omits argon and helium from the domain, you're going to do worse. Similarly, if you go a bit further afield and use a quantifier that only quantifies over things visible to the naked eye, or the things talked about in theories of psychoanalysis, you can be pretty sure you'll be doing worse. You could try and conduct everything in the language of fundamental physics, but it'll take you an exhausting amount of work just to describe one hydrogen atom— God forbid (though, unfortunately, She didn't) that the hydrogen atom interact with another hydrogen atom, let alone anything more complex. Maybe the ideal transhuman cyborg chemist would and should proceed in the language of fundamental physics but, importantly, by *not* restricting herself to the quantifier of fundamental physics it's not like the human chemist does particularly badly— in fact, she does demonstrably well.

With unabashedly Carnapian sentiment, I conclude that the privilege of the - umm? - *chemical* quantifier is to a great extent rooted in deeply contingent facts about ourselves, in our practical circumstances and so forth. Its epistemic privilege is hypothetical: *if* you want to do chemistry well, with your limited knowledge and abilities in this particular sociohistorical context, you better use a quantifier that includes the sort of thing helium molecules and hydrocarbons are in its domain. Trial and error; application of the scientific method has shaped the quantifier-like concept we employ in chemistry and ensured it is fit for purpose. A hypothetically privileged quantifier is all that's needed to get a simulacrum of substantive ontological debate off the ground. Interlocutors can bring their ordinary *conception* of existence to the table in thinking about and representing the existence of chemical entities, while bound to employ a single, privileged quantifier with a sufficiently determinate domain:

What do you get if two hydrogen atoms combine with a helium atom? As the answer is fixed by a privileged quantifier, this in many ways resembles the sort of debate the ontologist takes herself to be having about real existence - there's no sense, for those involved, that we're just having a debate over how to use our words; that it's a lightweight matter. And because the same conception of existence is being employed, the phenomenology of being involved in an existential debate in chemistry is going to be substantively the same as being involved in a debate about fundamental ontology— the sets, the props and the actors have all changed, but the play is the same. In fact if, against

expectations Non-Chemical Nihilism turns out to be true and only the entities starring in chemistry textbooks really exist (nothing smaller or larger or otherwise non-chemical) then it would transpire that the interlocutors would have been doing fundamental metaphysics all along, had they prudently adopted an epistemology for their metaphysics that was first and foremost deferential to the proclamations of chemists.

The Ontology of Chemistry is an interesting and substantive project, and might properly be called ontology. However, on the nihilist's assumption that the entities of chemistry are not real then the Ontology of Chemistry is not the age old project of trying to uncover the ultimate structure of reality – that is, it is not asking the categorical representation question. As such, it is not the same project which is the preoccupation of the OS ontologist and not my project in this thesis. In the eyes of OS, there is some absolute, categorical epistemic reason to represent the world as containing the objects of the Ontologese quantifier rather than the chemical quantifier or the ordinary language quantifier: to do (for want of a better term) *fundamental* ontology, well it's a precondition that you use the quantifier that the world has categorically privileged quantifier and there's little more than can be said than this— it would therefore be wrongheaded to try and use a quantifier in such a debate that did better for various pragmatic or instrumental reasons.

Meanwhile, we can step back from a debate in which a hypothetically privileged quantifier is being employed and question the utility of speaking that way or our goals or our motives – once this is up for grabs then it does like we've entered a debate into how best to use our words, or what our conceptual scheme is like, or what have you. But there are plenty of debates in which for those internal to them it is part of the assumed background to the debate that there's good reason to be speaking in the way everyone in that debate in fact is: similarly, if you're having an idle debate in the pub about what constitutes a Martini, it might seem to everyone involved that it's just a matter about how we want to use 'Martini', but we can imagine a conversation among sufficiently invested parties with set goals in mind for which it is perfectly obvious what we should mean by a Martini and the question of exactly what constitutes one to be an important issue.²² As such, by my account, the substantivity question is topic neutral.

The thing language is powerful— it's hard to avoid speaking it.²³ So, whether you're ordering a drink, studying for a qualification in interior design, doing any science from pharmacology to fundamental physics, or limning the ultimate structure of reality, there's

²² See Sider (2011) for more on the Martini debate.

²³ Recall §0.2.

a good chance you will end up making all manner of existential claims. What's more, keeping tabs on these existential claims will matter when it comes to succeeding at the project you have embarked upon, whether that's figuring out how to obtaining satisfactory liquor or perspicuously describing the fundament. But I take a very neo-Carnapian line here, in so far as I think that interesting ontological questions are largely *internal* to each of the projects in question: If you have an epistemically fruitful project then (at the risk of immediately sounding somewhat Quinean again) the entities you have to quantify over in course of carrying out that project are a given. However, I differ from Carnap in thinking that one project you can have is to ask which way of speaking represents the world in the categorically best way.

The difference between a substantive and a non-substantive ontological debate is purely whether you're in a context in which the quantifier meaning is fixed or is up for grabs. In a debate where the ordinary language quantifier is held fixed we can do ordinary language ontology, for instance. It doesn't matter for this purpose how the debate is conducted or its aims (find out what exists according to the privileged quantifier) whether the quantifier employed carves nature at the joints or not. What makes *real* ontology distinctive is that the quantifier does carve at the joints and we take ourselves to have categorical epistemic reason to represent the world in a joint-carving-way— to do so is just to do better, even if you get all the same “facts” right as another theorist, where ‘better’ is not to do with just our goals or circumstances.

But what if your concept of existence is not perfectly joint carving? Do we need a notion of partial joint-carving or some such to distinguish the likes of chemistry, that are highly continuous with physics from trivial talk about Martinis. If we're scientific realists then it's hard not to think that speaking about molecules somehow comes close to representing the world in a perspicuous way, compared to talking about glints and ghosts say, and that is doing better according to the world, not just pragmatically. Well, I don't doubt for a moment that chemistry is closer to the fundament, if you will, *more natural* or *more joint-carving*, in Sider's terms, than fundamental physics. While our future transhuman cyborg may have no trouble conducting chemistry by only including the entities of fundamental physics in her theories, even she is going to struggle to carry out social science in the language of physics - it is a task many orders of magnitude more troublesome. If there *is* hope of systematic paraphrase of some non-fundamental disciplines into Ontologese, despite my earlier doubts, it will be in chemistry as opposed to economics.

But all I'm really doing here is expressing some knee jerk scientific realist sentiments which many other metaphysicians have shared: The subject matter of chemistry is much

closer to the scale of particle physics than sciences which take the whole of human society into account; its primitive terms are arguably less prone to semantic indecision and anthropocentrism – above all, the great *success* of chemistry should be attributed to its theories mirroring the structure of the world at least tolerably well. Yet, if none of this were true, the fact would remain that representing the world as chemistry would have us do closely corresponds to success in chemistry: So, I see no bar to substantive disputes carried out representing the world *as if* it were that way.²⁴ But since how well chemistry discourse mirrors the structure of reality is, in the first instance, an unknown quantity, it would unhelpful to both the ontology of chemistry or fundamental metaphysics to insist that the former be systematically paraphrased by or shown to be systematically grounded in the latter before either got off the ground.

None of this is to say that there isn't an interesting project of trying to uncover grounding relations between the entities or facts of one discipline and another, or trying to paraphrase talk of entities in field away by replacing it with talk of more fundamental entities. It's just to say that it's to put the cart before the horse to demand that our *more* fundamental theories can recover the *structure* of less perspicuous discourse, even indirectly – all that's required of more fundamental theories is empirical adequacy along with theoretical, predictive or explanatory virtues that lead us to the theory in the first place. Also, I'm certainly not saying that just because an entity is not mentioned in fundamental physics it *doesn't* have real being and thus that it *shouldn't* be mentioned in our fundamental ontology – I think that's entirely up for grabs. But we shouldn't let their appearing in fruitful, fairly natural discourse persuade us that either such entities must appear in our ontology or else be paraphrased away or grounded in some systematic way. The other side of the coin, as we'll see in the next subsection, is precisely that we shouldn't necessarily presuppose that putatively less natural or 'higher-level' disciplines are correspondingly less likely to be having substantive ontological debates or quantifying over real entities.

3.3. Social Ontology

Is there an unhelpful naturalist bias in the proposed metaontological framework? While I don't rule out that entities lacking perfectly natural qualitative properties could still exist in the Ontologese sense, we would have to have some good independent reason

²⁴ Perhaps this makes me a little bit of a fictionalist—but I don't think any any of the usual ways. For a discussion of resorting to non-standard fictionalism as a reconciliation strategy see Williams (2012).

for doing this— while there is a presumption that putative bearers of perfectly natural properties really exist.²⁵ As such, views on which the entities of high-level disciplines such as the social sciences are not real will tend towards the default, and thus that practitioners of the metaphysics or ontology of such disciplines cannot be said to be engaging with the same project as those metaphysicians with the self-professed aim of limning realities ultimate structure. Elizabeth Barnes²⁶ has recently presented a criticism of fundamentalist metaontologies - both of the Sider and the Schaffer sort - along these lines. The worry is that fundamentalist, physics-deferential metaontological frameworks play a gatekeeping role, unjustly determining what counts as doing genuine, substantive metaphysics. Fields like social ontology - feminist metaphysics is particular Barnes' example - which quantify over entities such as *social structures* risk being counted as unsubstantive and not part of metaphysics proper, despite the self-description of the relevant practitioners and little in the way of detailed, persuasive argument to the contrary.

Barnes gives the example of Haslangerian social constructivism about gender. As Barnes characterises the Haslangerian view, 'gender is a system of embedded hierarchies - based on normative assumptions of perceived sex characteristics and their assumed role in reproduction - within a social structure.'²⁷ Social structures are created from patterns of contingent social interactions: But once a structure is created it is self-sustained and in some sense emergent from the mere patterns of interactions which created it. They are real, but made by us.²⁸ For Barnes, many debates in Social Ontology are obviously substantive ontological debates. For instance, Katharine Jenkins²⁹ and Haslanger are engaged in a substantive ontological debate: While agreeing that gender is grounded in hierarchical social roles, Jenkins believes that a perspicuous description of a person's gender requires describing both their *gender-class* and *gender-identity*³⁰, while the Haslangerian about gender is happy to speak about gender *simpliciter*. This is an

²⁵ See §5.2, §5.3 for full discussion on this point.

²⁶ Barnes (forthcoming-a).

²⁷ Barnes (forthcoming-a). See Haslanger (2000; 2010).

²⁸ Barnes (forthcoming-a); Haslanger (2011).

²⁹ See Jenkins (MS.).

³⁰ This is to take account of the fact that a trans woman does not 'pass' as a cis woman she will not belong to the gender class of women (more or less equivalent to gender *simpliciter* on the Haslangerian account) since she will not be assigned the relevant position in the social structure; while the introduction of gender-identity allows us to account for the fact that a trans woman is still a woman, however society perceives or treats her. (See Jenkins MS.; Barnes forthcoming-a).

important difference of opinion about what gender *is* which is at the same time not a debate about perfectly or fairly natural subject matter, nor merely an investigation into how we – the folk – actually happen to think or speak about gender.

As outlined above (§3.2), contra Sider (2011), I don't think that lack of naturalness or partial joint-carving leads straightforwardly to non-substantivity or mere conceptual substantivity. Debates about whether three-legged, high-backed seats are chairs, for instance, are non-substantive because we can only be asking either whether we actually happen to call such things chair or whether as a language community we *should* call such things chairs— were little of worldly, ethical or social consequence seems to hang on what collective decision we come to (as a language community we can call three-legged high-backed seats chair or not as we please, without being at risk of making any kind of objective mistake.) Meanwhile a lot hinges – socially, ethically – on how we speak about and represent gender, race, class, and other social phenomena: how we *should* represent these phenomena if we want to understand their social and ethical impact perspicuously is going to be a function of how the world is, including our own patterns of behaviour, and how the social and ethical facts lie. Thus, we can't simply collectively decide on an arbitrary (or historically established) way of speaking about and representing these facts without making a mistake relative to the terms of an inquiry which seeks to bring out the social and ethical implications of such behaviour patterns – and perhaps also a categorical ethically normative mistake (given the normative thickness of many of the terms in which the debate is cast³¹).³² As such, it's tempting to say the debate is clearly substantive and clearly ontology proper because the terms of the project pre-select for a privileged quantifier. It doesn't matter that the entities involved are not, by the lights of OS, *real*— proponents of OS have just tended to go wrong in taking reality to be a necessary and sufficient condition for substantive ontological debate, rather than just a sufficient condition.

However, while we could leave it there, having rescued substantivity, it does leave OS open to the 'Ladies' Metaphysics'³³ objection: While proponents of OS are off supposedly limning the ultimate structure of reality, social ontologists such as, e.g., feminist metaphysicians, are left to tidy up all the non-fundamental, social-type stuff that

³¹ See Barnes (forthcoming-*a*).

³² The appeal to normative perspicuity I take to be closely aligned with Haslanger's (2007) appeal to ameliorative projects – though I fear my own characterisation fails to do justice to the social constructivist assumptions informing Haslanger's own descriptions of such projects.

³³ Elizabeth Barnes – in conversation.

the OS-metaphysicians are demonstrably not interested in. The OS-metaphysicians may be happy to admit that the debate of feminist metaphysics are properly (not just conceptually) substantive, there's explicitly two separate projects going on here, this separation being constitutive of social ontology's explicit exclusion from the 'traditional' metaphysical project of uncovering reality's ultimate structure.

Now, I take the spectre of 'Ladies' Metaphysics' to be a forceful reason to suspicious of the above reconciliation strategy, but by no means decisive one. It's a live option for serious practitioners of social ontology to voluntarily choose to proceed on a somewhat different path from that which traditional metaphysics has happened to tread: For starters, it may well be - and would perhaps be unsurprising if, for all sorts of contingent sociological reasons tied up with the very subject matter of social ontology - the remit of traditional ontology just turns out to be objectionably *narrow*. In which case, we should actively encourage ontological investigation which doesn't fall within the traditional purview of the discipline. This, I think, is made more palatable in my proposed framework since social ontology is placed in the same sphere as e.g. the ontology of chemistry or other non-fundamental physical sciences which have garnered the partially natural accolade - and thus a recognition of independent importance - from proponents of OS. The dividing line is between the unabashedly niche project of *fundamental* ontology, with its categorically privileged quantifier, and The Rest, whether partially natural or non-natural, replete with their respective hypothetically privileged quantifiers. As opposed to conferring interest and substantivity only on the perfectly natural, though vicariously on the partially natural, then having to carve out a separate conceptual and dialectical space for the sole purpose of trying to secure the relevance of social ontology.

Secondly, with separation comes autonomy. Autonomy may seem like a dubious virtue in so far as it entails separation and separation is prejudged to be a Bad Thing, for precisely the reasons that drive the initial worry. But, as I see it, this must be balanced against the fact autonomy, in and of itself, is a Good Thing. When trying to uncover patterns, say, modes of systematic subjugation or privileging in social structures emergent from complex patterns of human social interaction, it's hard to think of anything less helpful - or potentially derailing - than a nihilist like myself coming along and saying this whole project is wrong or misguided because there are really no such things as humans. I can't see that whether the clouds of human-shaped simples floating about the Earth actually compose human beings has any bearing whatever on, for instance, how many genders there are when sufficient attention is paid to the relevant social and normative facts— but without autonomy arguments in the vein that there aren't really any human being, and thus not really any patterns of human social interaction, and thus not really any

genders look like they should at least be given the time of day. But preserving certainty about the existence of, say, gender-identity against attack from radically revisionary fundamental metaphysics is even more important than preserving certainty about my having hands: Arguably, hands don't matter much – they're just a vestigial feature of an arbitrary conceptual scheme thrust on me by history and evolution. Gender identity matters, though, (for real!) – if we're going to stop believing in it, it better be because we've uncovered novel facts about, say, our patterns of social interactions, not because we've learned something about the natural entities that ultimately ground these patterns of social interaction

The above notwithstanding, however, let's assume that separation is not an attractive option. Can OS make sense of the possibility that social entities are real in just the same way that natural entities are real— and that therefore the social ontologist's project is one and the same with the OS-ontologist's project? Well, there's certainly an easy shortcut to this result: Just admit that social entities are, after all, natural entities. I think this option is more attractive and plausible than it might at first sound: On the proposed framework, there is no presumption whatever, as I've already been at pains to point out, that everything which really exists must be the sort of entity that will appear in fundamental physical theories or instantiate properties that do. A physics first epistemology for metaphysics is an optional extra in this package, not a standard part of the bundle: much less is it hardwired into the view that only the entities of direct or indirect interest to physics can possibly have real existence. Nonetheless, it is a part of the view as stated thus far – and you would be forgiven for thinking a non-negotiable part of the view – that only the objects “carved out by the world” have real existence and are proper subject for the sort of project that aims at representing reality perspicuously. On the assumption that there are compelling, if not decisive, reasons to be naturalist in the first place, would this be such a bad thing? What the assumption of naturalism does not entail is that social objects and categories don't really exist and, as such, does not entail that the study of social objects and categories is anything but genuine substantive metaphysics.

Humans and their activity are as much a part of the world as subatomic particles. Nature may carve where physics fears to tread— perhaps into the realms of the normative, the phenomenological or indeed the social. As I write this Tiny, a male Jack Russell Terrier, sits at my feet and surveys the text with an eyebrow raised in apparent scepticism: As a nihilist, it's very tempting for me to go around believing and speaking as if the only candidate truthmakers for the English sentence “Tiny and I are both male” are subatomic particles or quantum fields. But Tiny and I are both persons integrated into social structures and while there are not on my view *really* any such things as human or canine

animals or biological sexes, perhaps there really are persons and genders and species.³⁴ This is all consistent with real being meaning *natural* being, and also with social entities being ontologically dependent on subatomic particles, where the natural world turns out to contain lots of perhaps abstract-seeming social objects, along with the theoretical entities of the physical sciences, and not so many paradigm macroscopic concrete things like mushrooms, trees and rivers as you might naively have assumed. Just because entities may be *caused* by ongoing human activity by no means rules out their existing naturally as a consequence. As such, it's not a foregone on adopting the proposed framework that the entities dealt with by the social science do not have real existence, and thus that the project of social ontology is not a proper part of the project of uncovering realities ultimate structure.

And yet, this is not enough to fully meet Barnes' criticism, for reasons that are made explicit in the paper. The big problem is social constructivist approaches to social ontology would appear to be ruled out: for the social constructivist, their theoretical entities are *ex hypothesi* non-natural. The quantifier of social constructivist ontology does not carve at the joints, so the OS strategy risks dismissing social constructivist claims as concerning no real entities and thus having debates that, even if still substantive, are not aimed at providing categorically perspicuous representations of reality. For starters, it would be objectionably committal for what is supposed to be a *general* theory of what it is to do fundamental ontology to rule out any non-naturalist approach to ontology straight off the bat. Beyond this, since much work on these issues does as a matter of fact take a constructivist line, merely pointing to the supposed possibility of doing this work along *non*-constructivist lines doesn't do justice to important work that's already been done.

I think there are two ways one could go here to satisfactorily recover the compatibility of social constructivism with OS (short of abandoning one or other of them to the wolves in the face of Barnes' criticism – which I fear in either case would be to risk throwing the baby out with the bathwater):

The first is just to go back more or less back to the account of social ontology I originally proposed, casting it as dealing in non-real entities, and offering an account of social constructivism that retains much of the spirit of the view as possible but is much more concessive than Barnes allows to the framework set by OS: I'll call this concessive social constructivism. On the concessive version, there is *some* sense of existence on

³⁴ The latter, I would venture, being a much more significant feature of Tiny's experience within the anthropocentric social structures he happens to occupy than his gender.

which it's true to say that social entities exist (this much comes cheap given what's been said), and while to talk of such entities existing does not represent the world categorically perspicuously, it *is* the right way to represent the world if we want to get the best understanding possible of complex human social structures and their (ethically) normative import. While we could adopt a different conceptual scheme for furniture and alcoholic drinks without losing anything of epistemic value, the best way to represent normatively thick structures like genders and races is forced upon us by what our actual patterns of social interaction are, what the ethical facts are like, and so forth (HINT: It doesn't involve any application of Quantum Field Theory) - to learn facts about how we should represent social structures is not merely to learn what the folk think about gender, or facts about our actual conceptual scheme, but about what conceptual scheme we *should* have. It may not be fundamental ontology, but it's *revisionary* ontology in the best and strongest sense. Since, as Barnes is at pains to point out, participants in the debate do not take themselves to be engaging in mere conceptual analysis or having a disagreement simply over how to use terms such as 'gender', but are debating substantive ontological issues such as what genders there are, it seems an attractive option for all parties to make use of something like hypothetical OS to underscore this.

However, we can also accommodate - dialectically and conceptually - a view of social constructivism on which social entities are categorically real: real in exactly the same sense as an electron is real, and yet, *ex hypothesi*, not a natural or fundamental entity. Call this full constructivism. If full social constructivism in ontology is true, then everything I've said with a naturalistic presumption should be recast in constructivist terms. Thus, we end up with a concessive version of OS rather than a concessive version of social constructivism. In these terms, the single conception of existence - shared by the folk and denizens of the ontology room alike - is of existing *either* in a joint carving way or a socially constructed way, where we are not furnished with the conceptual capacity to differentiate between which side of the disjunction is being satisfied: Indeed, to exist in either of these ways is just to exist simpliciter. While we're free to define one up, we cannot have a genuine understanding of what it is to exist according to a quantifier that carves nature at all of and only its joints: this is because we are simply unable to think about and represent things in a way that is independent of the social structures we find ourselves located in.³⁵ As such, the purely naturalist representation project was doomed before it started, since while we can draw up a list of all and only the natural entities with

³⁵ 'We are in reality, we must be in our theories [...] theories describe agential realities - which is just what we are interested in (we don't live in a transcendent reality).' - Barad (1996).

real existence, this is just to arbitrarily exclude non-natural entities which exist in exactly the same way according to the only conception of existence we have.

This would require far reaching changes to the way OS is usually understood and employed, but would not consign the general framework to the scrapheap, and I'm relaxed about this possibility. Full social constructivism is a bold and interesting claim not just about what exists but about what it is to have real existence: it is thus at least on a par with most of the metaontological postulates we've been examining in this thesis so far and should have similar important, far-reaching methodological consequences. For example, if a 'levelled', dependence based metaontology turns out to be correct this will obviously have big implications for work carried out in a 'flat' ontological framework without relative ontological priority: but it wouldn't mean all work done assuming that framework would immediately be rendered worthless, just in need of careful re-evaluation. The same would go for naturalism if a constructivist rather than a naturalist ontology turned out to be right (and, importantly for meeting Barnes criticism, *vice versa*). If full constructivism is ultimately proved false, perhaps a retreat to something like concessive constructivism would be needed: but we would be extremely unlikely to find ourselves in the position of having to consign any (let alone all) work assuming full constructivism to the proverbial flames. Again, in the spirit of metaontology being just more ontology: If you think the world is *really* full of non-natural social structures then adopt a full social constructivist metaontology— if diehard naturalists fail to see the value of adopting such a framework, that's just their loss.

4. Dismissing Dismissivism— the nihilist triumphant!

In this chapter I defend the composition debate against a certain sort of dismissive attitude: Not deflationary views such as quantifier variance, the view that (certain) metaphysical questions have easy or trivial answers, or are somehow non-substantive. Rather, I focus on what Karen Bennett¹ calls ‘epistemic dismissivism’. According to the epistemic dismissivist, it may well be that the suspect metaphysical questions have a single best answer but, far from being easy or trivial, these answers are exceedingly hard to discover – perhaps impossible – and are therefore not worthy of our efforts. Using composition debate as a case study, I show that the arguments in the literature adduced in favour of epistemic dismissivism fail. The theory choice criteria by which the dismissivist judges the competing options in the composition debate to be equally well (or, rather, equally badly) supported turn out – given the proposed metaontological framework – to in fact favour nihilism.

§4.1 introduces epistemic dismissivism while §4.2 explains Bennett’s argument for dismissivism from parsimony and §4.3 explains the argument from puzzle solving. In §4.4 I respond to the parsimony argument on behalf of the nihilist, while in §4.5 I respond to the puzzle solving argument. §4.6 shows in detail, through application of the general response I offer to epistemic dismissivism, that the argument that composite objects are causally redundant is not overcome by Bennett’s counterargument. Finally, in §4.7 I conclude this chapter and this first half of the thesis by arguing that as all reasons for rejecting nihilism or the composition debate in its entirety have failed, that the believer is in need of a new theoretical role for composite objects to play if her theory is to remain plausible.

4.1. *Epistemic Dismissivism*

Is there something wrong with the composition debate? Is it in some way silly, futile or misguided to be concerned with questions of when – under what conditions – do some

¹ Bennett (2009)

things compose some further thing;² This is a common suspicion:³ the composition debate has been a frequent target of such attacks in the recent metaontology literature.⁴ The sorts of arguments that go on in the mereology literature can seem fundamentally wrongheaded: For instance, arguing over whether there are tables composed of simples arranged ‘table-wise’, or just simples arranged table-wise and no extra thing, a table, that they compose. The general thought is that that this debate is in bad standing, that engaging with it on its own terms is not worth the time and effort, and that we would be intellectually better off just to forget about the questions raised by it. We shall call this general attitude ‘dismissivism’. Other metaphysical debates besides composition have been the target of dismissive attitudes: the material constitution debate, for instance, over whether, for example, the statue is identical to the lump of clay that constitutes it.⁵

Deflationism is just one attempt to cash out exactly what justifies taking a dismissive attitude to the composition debate and other suspect metaphysical debates. According to the deflationist, the composition debate is futile because there really is no sensible debate to be had: once you see things as the deflationist would have you see them – when you go along with her metaontological assumptions – it turns out that there are no sensible options on the table to choose between. There is no language independent fact of the matter, according to the deflationist, concerning whether or not composite objects exist. Whether or not it’s true to say that there are tables – and, thus, whether there are tables – just comes down to a more or less arbitrary or accidental choice of conceptual scheme.⁶ The target of this chapter will not be deflationism – I give an extensive reply in §2. Here, I’ll be defending the composition debate against a form of dismissivism that is a relative newcomer to the literature, but which has undoubtedly lurked in the minds of some sceptics for almost as long as the debate itself has been going on. We’ll call this ‘epistemic dismissivism’ (ED), following the terminology of its main advocate, **Karen Bennett**.⁷

The epistemic dismissivist is willing to accept that there may well be a single correct

² See van Inwagen (1990) for a definitive introduction to this debate and a detailed survey of the various options on the table.

³ Thomasson (2009) argues that it is a very widespread even within metaphysics.

⁴ In Chalmers, et al. (2009) alone examples include: Chalmers (2009), Hirsch (2009), Hofweber (2009), Price (2009), Thomasson (2009)

⁵ See Bennett (2009).

⁶ Recall §1, §2.

⁷ Bennett (2009).

answer to the Special Composition Question (SCQ): *When do some things compose some further thing?*⁸ But far from it being a trivial matter that is settled by just understanding how we, as a linguistic community, use our language, the epistemic dismissivist doubts that we are in a position to *ever* settle – or even make progress on – such questions: We’re just not in conducive enough epistemic circumstances to do so, and have no reasonable expectation of ever being in better circumstances with respect of SCQ. Once we have come to terms with our inability to make any progress on SCQ surely we really would be wrongheaded or misguided to continue to ask it and expend reams of paper debating the possible answers.

Bennett’s dismissivism concerning the composition debate is inspired by the observation that the nihilist seems unable to gain any ground against the believer in composite objects according to the theory choice criteria that tend to be employed in the debate. She identifies two main categories of theory choice criteria:⁹

Simplicity - This concerns minimizing the number of entities, or types of entities, your theory posits (its ontology) and the number or complexity of theoretical notions needed to state the theory (its ideology).

Puzzle Solving - This is a ragbag of assorted dilemmas and difficulties associated with belief in composite objects, which the nihilist alleges can all be swept away by rejecting an ontology of composites.

Bennett’s observations of how nihilism has tended to fare against the above theory choice criteria lead her to formulate her pessimistic ‘bump... under the carpet’ (p.65) hypothesis (BUC). She notes that for any putative gain the nihilist secures for herself over the believer in composites, she invariably incurs some analogous cost of a comparable magnitude elsewhere— from an epistemic standpoint, we appear to be just ‘riding a see-saw’ (p.65) rather than making any significant progress in the debate. According to Bennett we should infer from this pattern that our epistemic toolbox is just not up to the job: We will never be in a position to choose one theory over another on their individual merits, since this debate suffers from an extreme case of ‘*underdetermination of theory by evidence.*’ (p.71)

⁸ Again, see van Inwagen (1990).

⁹ I discuss more or less identical criteria in detail in §0.5, §0.6, §0.7.

4.2. Competing Simplicity Considerations

Bennett argues that competing simplicity considerations effectively balance out when comparing the nihilist's position with the believer's position,¹⁰ giving us no reason prefer one theory over the other. Ockham's razor gives us one simplicity consideration which can be used for evaluating competing theories: it tells us not to multiply entities beyond necessity. Thus, all else being equal, we should prefer theories which posit fewer things, or fewer types of thing.¹¹ Such theories will be more ontologically parsimonious: they'll posit less ontology than their rivals. Bennett accepts ontological parsimony as an important tool for deciding between competing theories that fit the available evidence, and concedes that nihilism is the most ontologically parsimonious of the different options on the table in the composition debate: for the believer is committed both to mereological simples and composite objects, while the nihilist gets away with positing just the mereological simples. The believer has an extra type of thing in her ontology, and for every table or water molecule or living being she says exists, her ontology grows larger relative to the nihilist's.¹² When comparing ontology, then, we should say that the nihilist has a simpler theory than the believer. However, a theory's ontology is not the only relevant consideration in determining how simple it is. We should also be interested in a theory's ideology— in the theoretical machinery it needs to posit, in the vocabulary it takes to be primitive. The more ideology a theory needs to posit, the less ideologically parsimonious it will be.

Now, neither ontological nor ideological parsimony can help out a theory that fails to account properly for the observational data— a theory must be *empirically adequate* for simplicity considerations to come into play. Thus, only if two competing theories fit the known facts equally well can simplicity considerations help us choose between them. Bennett argues that in order for the nihilist's theory to be made adequate, she has to massively expand her ideology relative to the believer in composite objects. As a result, says Bennett, one's theory is made no simpler by denying composite objects: any gains in

¹⁰ By believer I mean (following Bennett's usage) a believer in composite objects: this is inclusive of both the universalist, who believes that any collection of objects composes some further object, and restricted composition theorists, who place some restriction on when composition occurs, but allow that there are some composite objects.

¹¹ Nothing here will turn on the distinction between positing fewer things and positing fewer types of things, since the nihilist succeeds in doing both relative to her rivals.

¹² I leave aside for now which, if either, sort of parsimony we should care about most, or whether either genuinely applies in the case of nihilism (recall discussion in §0.7).

ontological parsimony will have to be paid for, like for like, in loss of ideological parsimony. As discussed in §1, the problem for the nihilist is that denying the existence of composite objects seems to put a lot of everyday statements that seem obviously true on a par with completely crazy statements: it's wrong to say there's three-headed tiger on a hover-bike zooming around my living room, but if there are no composite objects it is also false to say there is a chair in my living room, which seems obviously true. The statement 'there is a chair in my living room' is getting at something right about the way the world is, whereas 'there is a three headed tiger on a hover-bike in my living room' is completely wide of the mark. This can be brought out by considering the pragmatic implications of the two statements: Given 'there is a chair in the living room', I wouldn't be crazy to try and sit half a meter off the floor, despite knowing that I don't possess any powers of levitation. This is surely right: I appear to be doing this right now. However, 'there is a three headed tiger riding a hover-bike in the living room', suggests I should dive to the floor, roll for cover and grab the elephant gun. A terse denial of composite objects fails to preserve these differences in pragmatic content: the nihilist could not be taken seriously if she went around making recommendations such as 'there are no toasters; revise your breakfast plans' (Bennett 2009, p.58).

The nihilist (following van Inwagen 1990) usually gets around this by paraphrasing statements making putative reference to composite objects into statements which make reference only to mereological simples arranged in certain ways.¹³ Thus, since "there are chairs" can be paraphrased as 'there are simples arranged tablewise', which the nihilist and the believer both agree is true, the nihilist can treat 'there are chairs' in English to be in some sense correct, since it entails the true proposition "there are simples arranged tablewise", while 'there are no three-headed tigers' entails the true claim that there are no simples arranged three-headed tigerwise. The nihilist who follows this strategy rejects claims that commit to composite objects, but accepts the truth of the object-free paraphrases. Since no observational evidence could distinguish between there being a chair and there merely being some simples arranged chairwise, and the two eventualities do not differ in their pragmatic implications, the nihilist's denial of composite objects cannot be dismissed as obviously crazy. Bennett argues that this comes at a considerable ideological cost: The nihilist paraphrases ordinary sentences such as 'there are tables' by introducing *plural quantification*, allowing her to quantify over 'groups' of simples without committing her to any objects corresponding to these 'groups'. But what if the nihilist

¹³ See §1, esp. §1.3.

wants to paraphrase sentences such as ‘the chairs arranged in rows’ (p.59). Here the nihilist seems to have ‘used up’ plural quantification and requires some new tool – perhaps *perplural* quantification, allowing her to quantify over groups of groups of simples without incurring any ontological commitment to anything but the simples. The more complicated the talk of composites to paraphrase, the more tools the nihilist seems to need in her ideology. Thus, to capture the difference between talk of composite objects that gets things right in the important way that ‘there are chairs’ gets things right, and talk such as ‘there are three-headed tigers riding hover-bikes’ gets things obviously wrong (at least when the domain of quantification is restricted to the living room), it seems the nihilist theory needs to be greatly complicated.

4.3. Solve a Problem, Make a Problem

Having presented her argument for thinking that the nihilist has to complicate her ideology with perplural quantification and a notion of “being arranged F-wise”, Bennett shows us how she thinks having this commitment undermines another *prima facie* attraction of nihilism. One might be attracted to nihilism because it seems to dissolve several puzzles to do with belief in composite objects— however, Bennett argues that the nihilists belief in “things arranged F-wise” causes the puzzles to instantly re-arise in a subtly different guise. Bennett looks at four such puzzles: the question of what it takes some things compose an F, the problem of the many, paradoxes of co-location (specifically ‘statue and lump’ cases) and the causal exclusion argument (which I will set aside for now and discuss at length in 2.4).

Related to the Special Composition Question (SCQ), which asks when some objects compose a further object, Bennett thinks there’s an important question for any believer in composite objects (see Bennett 2009, p.66):

(SFQ): When, if ever, do some things compose an F (where F is a sortal kind term)?

This is a tough question, for most Fs. Consider: “When, if ever, do some things compose a cat?” To see that this issue is separable from the question of how things have to be to compose *some* further object, let’s assume universalism for time being. How do we go about find out what collections of simples count as cats? Conceptual analysis? Should the organism’s biological origins be taken into account? Our concept of a cat is almost certainly not a fully determinate one— as a community of speakers we’ve not decided to impose any sharp cut-off between things that are cats and things that are excellent feline impersonators. Furthermore, once we’ve done the conceptual analysis on our intuitive or biological concept of a cat, how do we know that our concept of a cat corresponds to worldly cats? The world may draw a very natural distinction between cats

and non-cats (a sharp one, or one involving ontic vagueness) which acts as a reference magnet, making ‘cat’ in English mean something that doesn’t quite correspond with either our intuitive or our biological conception of what it is to be a cat.¹⁴ The nihilist appears to have an easy way out: she doesn’t believe anything composes anything, so she simply answers ‘never’. Conditions can never be such that there is a cat, or any other composite F. However, Bennett points to a very similar question in the vicinity which *does* look problematic in the same sort of way for the nihilist:

Special Arrangement Question (SAQ): ‘when, if ever, are there some things arranged F-wise?’¹⁵

On the face of it, it looks just as difficult to answer the question ‘When are there some simples arranged cat-wise?’ as to answer ‘When do some simples compose a cat?’ The nihilist’s insistence that all expressions about composites can be paraphrased into expressions about the arrangement of simples means that the problematic question just gets paraphrased into the nihilist’s preferred way of speaking, and appears to remain as difficult to answer as ever.

Bennett shows how this difficulty seems to block easy nihilist solutions to certain well-known puzzles. Let’s take statue-and-lump cases of co-location:

There’s a lot of intuitive appeal to the general principle that two material objects can’t occupy exactly the same location. Perhaps two subatomic particles could interpenetrate under certain conditions, or that matter at other possible worlds with strange properties could, but not everyday macroscopic objects such as tables and sheep and rocks.¹⁶ Furthermore, in interpenetration cases, the co-located objects do not have all their parts in common, so they can be at least individuated in this respect: there’s a lot of intuitive pressure to deny that two objects could share all the same parts. However, intuitions in specific cases tend to conflict, or at least come into tension, with the general principle. For instance, in cases where a lump of clay is formed into a statue: here we want to say, for instance, that the lump could survive being squashed but the statue couldn’t— the statue and the lump would therefore have different persistence conditions, so according to the Indiscernibility of Identicals must be different objects. Much has been written on the best

¹⁴ See Lewis 1983, pp.171-2

¹⁵ Bennett 2009, p.66

¹⁶ See §0.5, also Sider 2001, p.141.

way to reject or embrace this conclusion, but Bennett's point¹⁷ is that the nihilist can't avoid taking sides by simply professing disbelief in statues and lumps and the like.

Why is there no easy way out for the nihilist according to Bennett? Well, after the nihilist has paraphrased composite object talk into simples-arranged-F-wise talk, the nihilist wants to be able to recover everyday claims about such objects. This is going to include claims about the persistence conditions of everyday objects like statues and lumps of clay. In order to pull this off, it seems the nihilist would have to attribute incompatible persistence conditions to the simples arranged statue-(and-lump)-wise: that they are arranged such that both 'would-survive-being-squashed-wise and... would-not-survive-being-squashed-wise.' (Bennett 2009, p.70) The nihilist can perhaps avail herself of one of the proposed solutions - properly adapted - offered by either side of the debate. But there would appear to be no distinctly nihilistic solution to statue-and-lump case: and thus no advantage to adopting nihilism with respect to this puzzle.

The pattern evident in the nihilist's response to worries of this kind is that in her commitment to properties of the sort "being arranged F-wise", and her insistence that talking in terms of simples being arranged F-wise is just as expressive as composite-object talk, the nihilist ends up with puzzles about which things are arranged F-wise. As such, Bennett concludes, it looks as if nihilism will never be a cheap or cost-free way of getting out of *prima facie* paradoxes involving coincident or composite objects.

4.4. A Nihilist Response: Parsimony

So, we can sum up Bennett's challenge to the nihilist as being something along these lines:

Doesn't having "simples arranged tablewise" in your theory throw up all the same sorts of metaphysical puzzles as having simply having "tables"? Doesn't it complicate you overall theory just as much?

This is, on the face of it, a legitimate concern. If the nihilist is content just to deny that there are tables, but affirm that there are "simples arranged tablewise" instead, it's not clear how this mere change of description makes the phenomenon in question any simpler or less puzzling. However, the nihilist doesn't have to follow van Inwagen's¹⁸ advice and commit to there being an available object-free paraphrase for every sentence

¹⁷ Drawing on McGrath 2005.

¹⁸ Inwagen (1990, p.108).

involving composite objects.¹⁹ Taking this route doesn't force the nihilist to deny there is something correct, or even true, about "there are toasters". If the proposed OS account²⁰ is adopted, provided the nihilist has an empirically adequate fundamental theory, she can take the truth of ordinary language claims about ontology as given, while not being required to provide a systematic translation schema to bridge correct sentences of ordinary language and veridical sentences of her fundamental theory.

One immediate benefit of taking an OS approach here is that the nihilist can dispense not only with claims about composites of the form "there are Fs", but also their problematic paraphrases of the form "there are simples arranged F-wise". It may be true, for some F, that there are Fs, or that there are simples arranged F-wise: but in the former case, the nihilist is not ontologically committed to Fs, and in the latter case is not committed to any fundamental property of "being arranged F-wise" that is collectively instantiated by some simples. All that's required to make either of these claims true is that there are some simples with certain intrinsic and spatial properties: some simples being certain ways individually (having charge and mass, let's say), and having particular locations, is enough to ground the truth both of "there are tables" in English, and the object-free paraphrase "there are simples arranged table-wise". Thus, the nihilist's fundamental theory of the world need make no mention either of tables or simples arranged tablewise: all the nihilist is ultimately committed to are some simples with certain attributes at certain locations. It just so happens that English speakers, confronted with some simples that are a certain way, are willing to assent to statements such as "there is a table" and (perhaps) "there are some simples arranged tablewise". In view of these patterns of assent, the nihilist can admit the truth of such statements in English; but in doing so nothing need be added to her fundamental account of how the world is.

So, with respect to whether belief in nihilism makes one's overall theory more complicated - through a decrease in ideological parsimony - we can say that no ideological cost is incurred on the proposed framework. Our nihilist says a description of the world which wore both its ontological and ideological commitments on its sleeve, could get by equally well without a property of 'being arranged F-wise' as it could without quantifying over composites. Nihilism, then, is at least as ideologically parsimonious as universalism or restricted theories of composition. There is no reason to suppose that nihilism is simply pushing a bump under the carpet when it minimizes its ontology: all the

¹⁹ See §1.5, §3.2.

²⁰ See §1, §2 and §3.

evidence suggests it achieves a genuine overall gain in simplicity. Non-fundamental ontological and ideological commitments are not counted, since we don't count such theories as categorically perspicuous representations of the world.²¹

There's also reason to think that nihilist has a *more* ideologically parsimonious theory than the believer. In addition to posing the Special Composition Question, van Inwagen's (1990) also poses the General Composition Question: what is (in non-mereological terms) composition? Van Inwagen (p.51) is '...inclined to think there is no way of answering...' this question, since:

... the concepts of "part", "sum", and "compose" form what (by analogy to "the modal circle" or "the moral circle") one might call 'the mereological circle, a closed family of concepts.

To employ mereological concepts in our theories, it seems we have to take one concept in the mereological circle - "part", say - as a primitive term from which the others can be defined. The primitive term is an addition to our fundamental ideology— it cannot be analysed in terms of more basic concepts we already have a grasp of outside of the mereological domain. However, the nihilist doesn't require a primitive notion of "part" (or "compose") in her ideology. The nihilist thinks there are no wholes, no parts— what's more, she thinks that there couldn't have been any. Thus, the nihilist's fundamental theory doesn't require a parthood primitive. A fundamental theory has to say how the world is structured. To do this, it doesn't have to include notions employed by theories that mistakenly postulate extra structure. A physicalist's fundamental theory won't include a dualist notion of phenomenal consciousness, for instance. The atheist needs no coherent concept of omnipotence to state her fundamental theory. The theorist who thinks that everything that there is exists in the same sense does not need to build in a distinction between different ways of being into her theory and assert that everything that exists only exists in one of these ways. Likewise, the nihilist does not have to have a primitive notion of parthood in her theory in order to disagree with the believer— she just has to say that when the believer employs terms like "part" and "whole", they fail to refer to any fundamental feature of reality. Thus, the nihilist should declare that mereological terms such as "part" are 'semantically defective'²² predicates, classing it with predicates of defunct scientific theories, such as 'Phlogiston', and myth, such as 'unicorn'. Since the nihilist gets by with one less fundamental concept, nihilism is the more ideologically

²¹ See §3.2.

²² See Dorr (2005).

parsimonious theory.²³

4.5. A Nihilist Response: Puzzle solving

So much for BUC when it comes to parsimony. What should the nihilist say about the puzzle-solving utility of her theory relative to the believer's? Well, simply that problems concerning SFQ (when is some composite an F?) and SAQ (when are there some things arranged F-wise?) simply don't arise on her theory. It's not the aim of fundamental ontology to clear up any confusion about the appropriate application of such terms as "there are tables". Why do we say that there is a table, or that the simples are arranged tablewise, only given certain distributions of simples in the vicinity, and not others? The nihilist's answer is that our choice about whether to say there is a chair, or simples arranged chairwise, tracks nothing of any worldly significance:²⁴ fundamentally, there are no chairs or simples arranged chairwise.

Accordingly, there can be no real ontological puzzles concerning when something is a chair or not, or how many instantiations of chair-ness or being-arranged-chairwise there are. Whether or not we say "there is a chair" when confronted with a given distribution of simples is a purely conventional matter, and the question must be settled looking at when English speakers actually find it appropriate to assent to the statement. The conditions under which English speakers are prepared to say "there is a chair" may well be vague and contentious, but it is not metaphysically puzzling for the nihilist— for whenever English speakers *do* decide that a certain distribution of simples would count as a case of there being a chair, it will be the simples being distributed in this way in any particular case that make it true that there is chair. For the believer, on the other hand, genuine metaphysical puzzles threaten— the believer (or the universalist at least) holds that when we say "there is a chair here", this carries an ontological commitment to a chair. Thus, the truth or falsity of "there is chair here" makes a real difference to the world that cannot be cashed out in terms of the arranging of simples. Thus, a believer with an incomplete or unsatisfactory account of when it is true that there is a chair has an incomplete or unsatisfactory account of what the world is really like.

Consider the statue-and-lump case again. The nihilist puts forwards an account of the fundamental nature of the world in which there is nothing metaphysically puzzling about such a case: we start with some simples (the English speaker looks at the simples and says

²³ This is the argument of Sider (2013).

²⁴ Or, for that matter, instrumental (§3.2) or ethical (§3.3) significance.

“there is a lump here”); next we move the simples around, changing the location of some of them (the English speaker now says “there is now a statue here”, and may also maintain that there is still a lump); finally we move the simples about some more (the English speaker says “there is no statue here anymore”). All that’s really happened, according to the nihilist, is that some simples have changed their spatial properties— nothing is at any point co-located, nothing has gone into or out or out of existence. The only potentially puzzling features of the scenario are all to do with the idiosyncrasies of the English language: at exactly what point is it true in English that there is a statue, or that there isn’t a statue anymore? Is it true in English that there is still a lump after the creation of the statue? But the nihilist will just take this to show that we have a rather confused and unnatural way of describing reality— of describing the distribution of the simples. The nihilist solves the putative metaphysical puzzle by showing that reality itself is not as confused and complicated as our English descriptions of it.

McGrath (2005, p.482) worries about the adequacy of the sort of account given in the previous section. If the nihilist doesn’t think there’s fundamentally any fact of the matter about, e.g., persistence claims in statue and lump cases, he warns, she ‘...will be forced to say the same about nearly every ordinary claim about persistence and destruction.’ (p482) Our failure to commit to an answer in the problematic cases will seep corrosively into all our ordinary talk:

‘Bush lost Rhode Island in the 2004 election’... is just as factual as ‘the sky instantiates blueness’... But our objector will have to disagree because it entails ‘Bush existed in 2004’ which, for him, is not determinately factual.

But the nihilist is not ontologically committed to Bush. ‘Bush’, in English, is taken by the nihilist to refer, with some considerable degree of semantic vagueness, to some simples (which simples exactly differing wildly depending on the time). Thus, when the nihilist translates ‘Bush lost Rhode island’ into the language of her theory, its truth will be *just as determinate or indeterminate* as the original English utterance: the nihilist can say that it’s determinately the case that there are some simples which an English speaker refers to when using the name ‘Bush’ (even if its indeterminate *which* simples are being referred to). The nihilist can still consistently hold that the English speaker does not say anything determinately true or false when she says that the statue is destroyed when the lump is squashed. The nihilist is still perfectly clear on what’s really going on in statue and lump cases: some simples are being rearranged (big deal). It’s the English speaker she takes to be confused, and thus the English speaker who has to worry about whether semantic indeterminacy in statue-and-lump cases might creep into statements about the political success of Republican politicians— *except*, of course, the ordinary English

speaker gets on just fine talking about these things. Semantic decisions have clearly been made along the way, collectively, which mean we're able to keep track of the identity conditions of the ordinary objects we quantify over. These rules may be confused in the sense that they are likely very inconsistent and gerrymandered from the perspective of fundamental ontology: But English speakers are likely successfully tracking Bush's career by using equally gerrymandered criterion that nonetheless seem very obvious and natural to us— e.g. the collection of simples that “looks like Bush”. In the final analysis, though, if there are any problem cases of genuine contradictions in the conceptual schemes of ordinary speakers, we could in principle solve them by collective stipulation, without treading on any deeply held metaphysical principles: we don't tend to allow for, e.g., co-located objects in our ordinary conceptual scheme, but since nothing much hangs on the matter we could simply stipulate that this is how we will “clean up” our conceptual scheme without needing to check that this corresponds with reality.

Yet, since we have lots of (often conflicting) intuitions about composites, including how to count them and how to destroy them, there is a danger that an ontology which takes such objects seriously will implicitly incorporate such intuitions, and thus be vulnerable to the puzzles that arise because of them: once incorporated into fundamental ontology they can't simply be stipulated away, and *prima facie* arbitrary or gerrymandered solutions will need justifying. As such, the believer has to show that her theory can escape these puzzles (see §0.5), whereas nihilism simply dissolves them. The nihilist contention, then, is that hers is the best unified solution to the puzzles available— Bennett and McGrath's worries do not tell against this. Contra Bennett, then, it seems that nihilism can be a powerful puzzle-solving tool in ontology.

4.6. *Causal Exclusion: A case study*

Let's apply what has been argued for in the previous two sections to another puzzle Bennett discusses: the results here will tell strongly in favour of nihilism's puzzle-solving ability. Bennett considers a version of Kim's (1993) exclusion problem, adapted by Merricks (2001) to target composite objects in general: the problem suggests that if there are any composite objects, they are mere epiphenomena with no causal powers— all causal work is done at the microlevel by mereological simples. This threatens to make composites into strange, redundant postulations on the part of the believer: if our intuitive idea of medium-sized objects bumping into each other and doing causal work is completely wrong, why bother holding the belief at all and positing all this extra ontology? The argument parallels arguments against epiphenomenal qualia in the consciousness debate: there's some intuitive force behind postulating non-physical phenomenal

properties; but if these non-physical sensations and desires play no part in guiding our actions, or even our thought processes – and are thus so unlike our pre-theoretical conception of them – it’s doubtful that the intuitive impetus for postulating them survives in any significant way.²⁵ Here I’ll assume there is a significant cost for the believer taking this horn of the dilemma.

We can construct a causal exclusion argument against composite objects in the following way: Assume that every physical event has a sufficient microphysical cause. Now consider a composite object – a ball, say – striking a window. Was the ball the cause of the window’s breaking? Well, the breaking of the window had a sufficient microphysical cause: the simples constituting the ball. So, if the ball *did* cause the window to break, the breaking of the window was overdetermined, since both the ball and its microphysical parts are supposed to be sufficient causes of the window breaking:




Figure iii. Overdetermination

We are familiar with causal overdetermination as an occasional or accidental phenomenon: if you were to throw two balls at a window, it’s natural to say the breaking was overdetermined, since either ball hitting it would have been sufficient to cause the breaking. Now, it’s intuitively false to say that in the one ball case, and other standard cases involving what we take to be composite objects, are overdetermined. Furthermore, it seems very strange indeed to suppose that world conspires to ensure *all* causation involving putative macro objects are overdetermined in this way. The believer must pay a significant cost if she embraces this horn of the dilemma also.

Bennett responds on behalf of the believer by presenting what we might call a microlevel exclusion problem (inspired by Hudson 2003). If overdetermination is a problem, she argues, then it’s just as much a problem for the nihilist as the believer. If

²⁵ Further, some might worry that we can only make sense of something being a concrete object if it has at least some causal powers (see Merricks 2001, Hudson 2003).

there's systematic overdetermination at the microlevel, the reasoning employed in the exclusion problem pushes us to say that microlevel phenomena are also epiphenomenal, pulling the rug out from under the nihilist's feet— are we to say that nothing at all has causal efficacy? The obvious way out of such absurdity is to deny that systematic overdetermination is a problem: which means we shouldn't think there's anything problematic about composites systematically overdetermining the effects of their microlevel constituents.

So, what is the microlevel exclusion problem supposed to be exactly? The thought is that (due to the same sorts of considerations at play in formulating the problem of the many) when some simples arranged ball-wise are hurtling towards a window, there are many different collections of partially overlapping simples, *each* arranged ball-wise, and *each* on its own sufficient to break the window:




Figure iv. Many baseballs

Thus, when the simples hit the window, there is 'casual competition *among the simples*' (Bennett 2009, p.69) when it comes to determining the cause of its breaking. Each of the many overlapping groupings would have been sufficient on its own to break the window, so it seems the window's breaking was overdetermined by various microlevel causes:




Figure v. Micro-overdetermination

The mistake is to assume that there is a single event, e , “the window smashing” or even “the simples which were arranged window-wise becoming arranged smashed-wise”, which is overdetermined. Or rather, there is really no such event: an Ontologese description of what’s going on here would not mention any such occurrence. We take “the window smashing” to be taking place if any of a diverse range of possible microlevel events (e_1 , e_2 , e_3 , ...) is occurring. The specific fine-grained microlevel event, e_i that in fact occurred (we might say “the window breaking *in a particular way*” in English) when baseball-arranged simples collided with the window-arranged simples was not overdetermined. Had any of the simples hurtled at the window been absent, a slightly different microlevel event (e_2 or e_3 or e_4 ...) would have occurred: thus it took *all* the candidate simples acting in concert to bring about e_i . To see this, consider that “the window breaking” is a complex microlevel interaction *between simples*.²⁶

²⁶ The diagram suggests a very simplistic “billiard ball” model of microlevel causation: but however messy, and perhaps metaphysically indeterminate, actual microlevel causation may be, the nihilist’s picture of fundamental causation should still be one of highly specific, non-overdetermined events. (Talk of events here does not have to be taken too seriously— Bennet (2009, p,68) and I agree that it doesn’t seem to matter for the purposes of the argument whether events or objects are taken as causal relata.)




Figure vi. Simple breaking

As no individual interaction is causally overdetermined, no mere aggregate of these individual interactions will be overdetermined. The most perspicuous description of e_i then, is simply a list of all these individual - not overdetermined - interactions.

The picture is complicated by the fact that there are many distinct events *actually taking place* when the simples collide that are candidates for being picked out as the window-breaking event: for instance, we could restrict our attention to just the simples near the epicenter of the collision, or cast our net very widely to consider all the simples in the local vicinity. So when we simply speak of the precise way in which the window smashed, e_i , there are a number of very precisely described microlevel events that could be meant (e'_i , e''_i , e'''_i ...). Thus, there is semantic vagueness concerning exactly which simples were involved in bringing about e_i , but only because what event we mean to pick out by “ e_i ” is vague: for any of the admissible precisifications of e_i (e'_i , e''_i , e'''_i ...), e_i is not overdetermined. It is, then, determinately not overdetermined. The believer may complain that all this explication of what I take to be going on this putative case of overdetermination comes to nothing: after all, I’m not denying that e_i , “the window breaking”, is overdetermined— I’m just making the less controversial claim that e'_i , “the window breaking in a particular way”, is not overdetermined. True enough: but the important point here is that, fundamentally, the world contains no events like e_i ; only events like e'_i - there are no breaking windows, just simples changing their intrinsic and spatial properties.

So, once we adopt the metaontological framework proposed in this thesis, only the *fundamental* overdetermination of effects matters - and we’ve seen there is no fundamental overdetermination in the case of the simples-arranged-ball-wise smashing the

window. But once we say that only fundamental overdetermination matters, might this get the believer off the hook also? Why not say wholes overdetermining the effects of their parts is also an innocent form of overdetermination: can't the believer agree that the micro-constituents of the ball are the fundamental cause of the window breaking, while maintaining that the ball itself is still a cause? This line of reasoning only looks plausible given an equivocation in the use of the word 'fundamental'.²⁷ Sure, even for the believer it's uncontroversial that the only causation involving fundamental entities is the causation involving the ball's parts, if by 'fundamental entities' we mean the very small entities talked about in fundamental physics, or the bottom of the mereological hierarchy— but this is not what is at issue. If one takes both simples and composites to be fundamental, in the sense that both are real, that we are ontologically committed to both – that is, if one rejects nihilism – then there is the question of whether one sort of fundamental entities (composites) systematically overdetermine the effects of another sort (simples). Which entities are mention by physics, or at the bottom of the mereological hierarchy, is besides the point: nature would still be conspiring to have two (or many more) objects overdetermine the effects of pretty much every event that takes place. Only by denying there are really any composites in the sense that the nihilist does can the situation be rescued.

4.7. End of Part One

In previous chapters I said that on the proposed metaontological framework we are able to make sense of the project of fundamental – and non-fundamental²⁸ – ontology, that our Moorean intuitions can be preserved and that the deflationary objection of the quantifier varientist can be met.²⁹ In this chapter, we saw the fruits of abandoning the van-Ormwagean paraphrase schema:³⁰ Minimal ontologies, such as nihilism, can achieve overall parsimony gains compared to more profligate ontologies, and proffer systematic eliminative solutions to puzzles which plague richer ontologies, without these puzzles simply re-arising as Bennett describes.

With the proposed metaontological framework in place, things are not looking so good for the believer in composition. The believer has a more ontologically and ideologically

²⁷ See §3.1.

²⁸ §1.4; §2.

²⁹ See §2.

³⁰ A strategy first mooted in §1.5 and further justified in §3.2.

costly theory which comes along with the baggage of puzzles including the problem of the many and the paradoxes of co-location, which at the very least will require further significant theoretical commitments to overcome.³¹ And what does all this extra commitment buy the believer? Causally redundant objects which inherit their properties from their parts: it's hard to see what useful theoretical role such entities could play. I suspect the most common reason for keeping such objects around has to preserve an ontology that fits closely with our ordinary conceptual scheme³² and to maintain Moorean truths: However, we've seen that the nihilist can do this without having to resort to additional ontology, since the truth of our everyday existential claims can be maintained without requiring one-one correspondence between the things we ordinarily quantify over and our ontological posits.

Could there ever be any good reason for belief in the reality of composite objects? Or does the adoption of the proposed framework make the move to nihilism inevitable? Well, I think that there still could be good reason to have composites in your ontology if they provided necessary expressive power. In particular, if wholes turned out to have novel properties “over and above” their parts, this would make them much harder to eliminate successfully from fundamental ontology. I speak of ontological emergence: a concept with a long history but which has tended to be treated with suspicion on both sides of the composition debate, and which many have even branded unintelligible or contradictory. As we see in the next part of this thesis, the believer can fight back by availing herself of her own plausible metaontological principles: principles which make the notion of ontological emergence intelligible and fit to do explanatory work. Given this, putatively emergent phenomena from quantum mechanics or consciousness, say, might be most straightforwardly explained by positing genuine ontological emergence. This presents a poignant challenge to the nihilist: can she maintain the theoretical benefits of her austere ontology while competing with the expressive power of emergent composites - is emergence even possible without composition? It is this challenge which I will spend most of the rest of this thesis attempting to meet.

³¹ See §0.5, §0.6.

³² Either by placing a restriction of composition that closely mirrors our ordinary conceptual scheme, or through universalism which provides more enough objects and recovers our ordinary conceptual scheme through quantifier restriction.

*PART II: NIHILISM AND
EMERGENCE.*



Figure vii. Cthulhu emerges: Father of Ferrets; prior to all*

5. Emergence for Nihilists

In the first part of this thesis I defended the substantivity of the composition debate and one particular position in that debate, mereological nihilism, from various forms of dismissivism. In the context of an independently motivated metaontological framework, we saw that radically revisionary ontological claims such as nihilism are possible, intelligible and don't conflict with common sense (contra Moore, Hirsch, Thomasson, Korman, Hofweber, etc.). Also, that within the proposed metaontology there is the possibility of making progress on the composition debate via appeal to the theory choice criteria usually appealed to - e.g. puzzle solving; simplicity (contra Bennett (2009)). Furthermore, we saw that employing these criteria leads us to the conclusion that nihilism is an attractive answer to SCQ.

In this second part of the thesis, I see whether there is still hope for the believer in composition if she adopts a similar strategy to the one I have recommended to the nihilist - if she fights metaontology with more metaontology. Through defensible modifications to the proposed framework, the believer can put forward an account of ontological emergence that can provide an attractive solution to outstanding issues in the metaphysics of consciousness and the ontology of quantum mechanics, but which cannot be straightforwardly employed by the nihilist. I then proceed by exploring a number of ways the nihilist might account for ontological emergence, and thus resist the believer's argument for composition from emergence. In a similar manner, I then argue that *pluralistic* nihilism (the belief that there are many simple things, e.g. subatomic particles or spacetime points, rather than just one giant simple— the entire world) is not substantially threatened by a similar argument from ontological emergence - contra Schaffer (2010). I conclude with some remarks in favour of parsimony - which I have drawn on heavily in both parts of the thesis - as a theory choice criterion in the composition debate.

This chapter defends mereological nihilism, the view that there are no composite objects, against a challenge from ontological emergence, the view that some things have properties that are 'something over and above' the properties of their parts. As the nihilist does not believe in composite wholes, there is nothing in the nihilist's ontology to instantiate emergent properties - or so the challenge goes. However, I argue that some simples (taken together) can collectively instantiate an emergent property, so the nihilist's ontology can in fact accommodate emergent properties.

§5.1 briefly explains ontological emergence, then outlines the argument against nihilism from the existence of emergent properties (the No Bearers argument). I discuss putative examples of ontological emergence from the philosophy of mind and quantum mechanics, giving two reasons the phenomenon cannot be dismissed out of hand. §5.2 sets out some metaontological assumptions that will be helpful for understanding emergence, which include introducing and clarifying the notions of fundamentality and ontological dependence. §5.3 formulates the No Bearers argument in more detail and discusses the different ways the nihilist might respond to it. I argue that the best solution is for the nihilist to employ plural instantiation, holding that emergent properties are collectively instantiated by multiple simples. §5.4 weighs up the ontological and ideological cost of introducing plural quantification – necessary for making sense of plural instantiation – into the nihilist’s theory. I conclude that the cost is negligible and so does not negate any of the claimed ontological or ideological parsimony advantages of nihilism.

5.1. *The Argument from Emergence*

Mereological nihilism is the view that there are no composite objects: that when you have some things (maybe some subatomic particles), they never compose some further thing (such as a cat).¹ Here I’ll be defending mereological nihilism against a challenge from ontological emergence. According to the believer in ontological emergence, it’s at least sometimes the case that when you have some things and you arrange them in a certain way, you not only get some further object, but some further property – an *emergent property* – instantiated by that further thing. Emergent properties are supposed to be something genuinely in addition to, or not metaphysically reducible to, the properties of the things you started out with: the whole is greater than the sum of its parts, having at least one property that is neither inherited from those parts, nor somehow derivative of the properties those parts.² The most often discussed (putative) example of this phenomenon hails from the Philosophy of Mind: When you have some neurons and arrange them in a certain very special way – brain-wise, that is – not only do you have some further object, a brain, but some further property (or properties) – namely, phenomenal consciousness (or qualia) – that is something genuinely *over and above*, or not metaphysically reducible to, all the properties of the individual neurons and the

¹ For more details see... well, most of the previous half of this thesis, but especially §0.4.

² At this stage I’m deliberately characterising the phenomenon loosely using vague terms of arts – we’ll be working towards a more rigorous formulation of emergence (or at least better precisified terms of art) in the next section.

spatial relations between them taken together.

The challenge to the nihilist is clear: If you don't believe in composite objects, how can you account for the existence of emergent properties? Assuming we have independent reason to believe in the existence of emergent properties, we have reason to posit composite objects as their bearers. Suppose a believer in composite objects grants, for sake of argument, that we have no evidence for the existence of chairs over and above the evidence we have for believing in simples arranged chair-wise: That there appears to be a four-legged wooden object in region *R*, and that when I sit down in region *R* I don't fall on the floor, is just evidence for there being simples arranged chair-wise in region *R* - it doesn't tell between there being some further object, a chair, in region *R* over and above the simples so arranged, on the one hand, and their being *merely* some simples arranged chair-wise on the other. However, our believer might go on to say, the fact that I have (say) the subjective experience of diverse yet unified phenomenal feels - i.e. the existence of emergent consciousness - is evidence for the existence of some thing (a mind or brain) over and above the simples arranged brain-wise in the vicinity of my head.

The crucial thought here is that emergent properties such as consciousness require the existence of a bearer, and the nihilist doesn't provide any viable candidates to be the bearers of the properties. The nihilist only believes in simples: tiny (perhaps point-sized) mereological atoms that - very plausibly - lack the rich internal structure needed to be thinking, feeling beings. Van Inwagen gave an extended defence of nihilism, and came partway to embracing it: He was ultimately deterred by the worry that the nihilist's ontology contained no bearer for the property of consciousness (see especially van Inwagen 1990, p.118). Merricks (2001) follows roughly the same path, rejecting pure nihilism because of a belief that a given human being, by virtue of being conscious, possesses at least one property "not, of metaphysical necessity, implied by the existence and intrinsic properties of, and spatiotemporal and causal interrelations among [its] constituent atoms." (p.89) Indeed, the prevailing assumption in the literature would appear to be that a nihilist ontology cannot accommodate emergent properties. Many of those disposed favourably towards nihilism for independent reasons are impressed by this enough to ultimately reject the theory. Yet little work has been done to defend the nihilist against the challenge. The nihilist can, of course, simply reject the premise that we should believe in the existence of emergent properties. Whether we have good reason to believe in emergent properties is a debate I largely want to side-step in this paper, but I do want to point out that, given some recent developments in the literature, two reasons one might have had for dismissing the existence of emergent properties out of hand are unsatisfactory.

The first unsatisfactory reason for dismissing emergence is that many of the putative cases of emergence in the literature concern properties such as consciousness or qualia, or mental properties more generally (as with the examples we have discussed so far). Because of this it might be assumed that taking a hard-line physicalist or eliminative approach to the mind removes the need to posit emergent properties. If an instance of pain is nothing more than some neurons firing in a particular pattern, then why posit some further thing, a mind, with novel properties? But the merits or otherwise of such an approach aside, it has recently become clear that there is at least one other domain that generates putative examples of emergence: namely, Quantum Theory. In particular, quantum entangled systems appear to possess properties that are not reducible to the intrinsic or spatiotemporal properties of their constituent subsystems.³ There have been attempts to provide the nihilist with alternative resources to account for this phenomenon - to *explain away* the putative emergence - which we won't be discussing in this paper.⁴ However, the point is that these new cases can't easily be swept under the rug, and if putative examples of emergence can be found in two very different domains, perhaps they will be found in others. A general expressive limitation in the nihilist's theory has been uncovered. Unless she can find a way to incorporate the required expressive resources in her theory, that theory is hostage to fortune.

Secondly, it's often assumed that the very idea of an emergent property is somehow confused or incoherent (see, e.g., Kim 1999). When combined with a popular thesis we'll call *micro-determination*, it becomes difficult to make sense of the idea that mereologically complex entities could exhibit properties that are 'novel' or which 'transcend the properties of their constituent parts'. The core of the micro-determination thesis is that fixing the facts about the smallest objects in a world fixes all the facts about that world: On one natural reading of this, the arrangement and properties of point-sized particles determines the locations and properties of all macroscopic objects. The thesis often comes along with the metaontological assumption that it's only the smallest entities and their properties we should really care about when evaluating competing ontologies, and that macroscopic entities are 'ontologically innocent' since they supervene on the arrangement of micro-entities.

³ See, for example, Maudlin (2007) or Schaffer (2010).

⁴ For instance, Shaffer (2010) examines - though ultimately rejects - using 'entanglement relations' to explain entangled states. Barnes (forthcoming-*b*) proposes an analogue of this strategy in which the ideology of metaphysical indeterminacy in effect stands in for or replaces the ontology of entanglement relations.

Such metaontological assumptions push you to the conclusion that, were there emergent properties of macroscopic objects, they would be redundant. According to micro-determination, the existence of ‘higher-level’ properties supervenes on the configuration of the microphysical properties. So, rather than endow a macro-object with genuine causal powers, over and above the powers of the micro-base, we can always attribute these powers to the base itself. Rather than saying that configuration-X of the base causes E to emerge, and E has power P, why not just say that the base in configuration-X has power P? With the related assumption in place that macroscopic objects are ontologically innocent – that they come for free once you’ve fixed the properties of the micro-objects – we’ve gotten into a position where it’s very hard to make any sense at all of the claim that some macro-objects could have properties that are something over and above their microscopic parts: *Ex hypothesi*, for any given macro-object, its properties are fully determined by the properties of its parts, it can make no causal contribution that fails to be over-determined by its parts, and it is in some sense a light weight ontological posit, since it comes for free if the micro-objects and their properties are already included in your theory.

Formulating an interesting definition of emergence requires giving up micro-determination. However, on first analysis, this comes at a high theoretical cost: For a start, the truth of micro-determination is often thought to be necessary for the truth of physicalism.⁵ Furthermore, the believer in emergence wants to maintain that for many ordinary macroscopic objects – for medium-sized dry goods such as tables and chairs – *their* properties are nothing over and above the properties of their parts. Perhaps worse, if the believer in emergence gives up micro-determination, in what sense could she hold on to the idea that at certain levels of complexity the sorts of properties she’s interested in emerge *from* the microphysical base? Without saying something more to explain the relationship between the micro-base properties and emergent properties, the believer in emergence is left with an ontology of autonomous ‘levels’ that float free of one another, rather than one in which a complex micro-structure can give rise to novel macro-level properties.⁶ However, Barnes (2012) shows us that giving up micro-determination needn’t mean incurring any of the above costs, provided we replace it with a weaker thesis, which I’ll call micro-dependence. Many of the metaontological concepts of the micro-determination picture carry over to the micro-dependence picture, but the way that

⁵ To list just a few influential works where something like this assumption is made: Chalmers (1996), Kim (1989), Jackson (1994), Lewis (1986*a*).

⁶ See Kim (1999) and Merricks (2001, p.90).

these various bits of apparatus interact with each other changes: on Barnes' framework, we make sure to keep apart the notions of fundamentally, *coming for free*, and being *nothing over and above*, on the one hand, and notions such as dependence and supervenience (i.e. mere modal co-variance), on the other.

5.2. Understanding Emergence

In the Barnes framework, we can think of micro-determination as the conjunction of two entirely independent metaontological claims: First, that the microphysical determines the macroscopic in the sense that there can be no change in properties of a macroscopic thing like a brain without some corresponding change in the properties (intrinsic or relational) of the microscopic particles that compose it (at least when considering worlds where actual laws of nature hold).⁷ This is a supervenience thesis stripped of any of the usual metaontological baggage: it does not by itself imply that the supervening entities are ontologically innocent or lightweight; it is a claim of mere modal co-variance. This thesis is recast by Barnes in terms of *ontological dependence*: the properties of the whole cannot exist without (at least some of) the properties of the parts.⁸ Here ontological dependence is, similarly, stripped of any implication that the dependent entities are somehow derivative from or less real or less ontologically committing than the things they depend on. The idea is that my brain would have none of the properties it in fact has without its actual microphysical parts – or qualitative duplicates of those parts – standing in the same or similar relations to those they actually stand in. This leaves open whether we want a relation of ontological dependence to be a part of our fundamental ontology or ideology, in order to *explain* or ground the posited modal-covariance, or whether we want the intuitive notion of ontological dependence to be itself explained in terms of, or reduced to, the type of modal covariance just described.

An advantage of the move to talking about ontological dependence that is not made explicit in Barnes (*forthcoming*) is that the form of modal co-variance it delivers us is weaker than strict supervenience in a useful way: it may be that strict supervenience – i.e. no change in the properties of a supervening *a* without a change in the properties of the

⁷ e.g. Chalmers (1996), Kim (1999), Lewis (1983*b*).

⁸ This is not exactly how Barnes herself prefers to cash out ontological dependence – she opts for the absolute notion of something being a *dependent entity*, over specifying precise conditions for *x* being dependent on *y*. However, I don't think this wrinkle need concern us here.

subvening bs ⁹ has to be given up in the face of the examples of emergence from quantum mechanics that we have previously discussed. It seems to be the case that for a given entangled system c , there's a physically possible world at which its spin is different from its actual spin, yet the qualitative and relational properties of its component particles a and b are the same.¹⁰ Ontological dependence weakens supervenience just enough to accommodate these cases without giving up the spirit of the original micro-determination thesis: For it's the case that actual spin properties of c could not have been instantiated if it had not been composed of two particles with such-and-such properties. Thus, the properties of c are appropriately constrained by the properties of its parts— c could not have been a spin up if both its component particles were spin down, for instance, or (given some plausible assumptions) if it had been composed of ectoplasm instead of fermions or bosons. However, this constraint still allows God to strip c of its actual spin properties without changing the properties of a and b , and to add in new spin properties that are similarly dependent on a and b .

Weakening supervenience in this way does open the door to zombies: for all we've said, it's a physical possibility that there exists an exact duplicate of me that lacks consciousness (Chalmers 1996). However, one can rule out objectionable physical possibilities on a case by case basis: even in the quantum mechanical case there will be nomic constraints on what properties God can simply subtract from objects with a wave of the divine hand—for example, let's suppose system C must have its actual mass if its parts have their actual mass, even if it's consistent with ontological dependence that c 's mass might not have existed at all, despite the presence of a and b 's mass. If the quantum mechanical case teaches us anything it's that we can't hope to get a systematic theory of what properties are subtractable in this way simply from the truth of supervenience physicalism or a general microphysical dependence thesis—we have to get out of our armchair and study the actual laws.

The second, distinct, claim of microphysical determination is that macroscopic objects – that is, dependent mereological wholes – are ontologically innocent. That is, we get the composite objects and their properties 'for free' once we have fixed what microphysical things there are and what properties they have: the intuition is that these entities are somehow derivative and should count against the ontological parsimony of our theory

⁹ See, e.g., Chalmers (1996) and Jackson (1994).

¹⁰ I.e. the system as a whole could be in either a correlated or anti-correlated spin state while each individual particle that composes it is merely in an indeterminate spin state (for more detailed descriptions of cases see Maudlin 2007; Schaffer 2010; Barnes *draft*).

since they are nothing ‘over and above’ their parts. Barnes cashes out these intuitions in terms of fundamentality: some entities are fundamental while others are derivative (non-fundamental). The fundamental entities are ontological commitments in their own right, while derivative entities are ontologically innocent. Importantly, fundamentality is binary rather than degreed on this framework: if x is *more fundamental* than y , then y is not fundamental *simpliciter*; an entity is either fundamental or it isn’t.¹¹ A natural way of understanding what’s going on here is by employing Sider’s natural quantifier: the fundamental things are what exist according to the perfectly natural quantifier— the quantifier that carves nature at its joints.¹² Non-fundamental entities only exist according to less joint-carving quantifiers, such as perhaps the quantifier of ordinary English.¹³ We should only care about the domain of the perfectly natural quantifier when doing serious ontology, which explains why derivative entities aren’t ontologically committing. But that’s just one way to explicate the fundamental / derivative distinction: one could adopt a metaontology more like Schaffer’s (2010) on which fundamental entities exist in the same sense as non-fundamental entities. The important point is that on neither metaontology does positing a derivative entity incur any additional ontological commitments: As Barnes (*forthcoming*) likes to put it, God fixes what the fundamental entities are and gets the derivative entities for free.¹⁴

It’s easy to see why the fundamental / derivative distinction has tended to be run together with the supervenience thesis. Lewis (1986*b*, p.14) introduces the very idea of supervenience by asking us to consider the relationship between a digitally produced picture and the pixels that compose it: Concluding that the compositional properties of the picture are “nothing over and above” the properties of its pixels, and that when you fix the colour and arrangement of the pixels you get those compositional properties for free,

¹¹ This doesn’t prevent you from giving entities some sort of ordering in terms of, say, metaphysical priority if you find that useful, but such a notion won’t concern us here and shouldn’t be confused with the notion of fundamentality being proffered.

¹² See e.g. Sider (2001) and (2011) for detailed explanation.

¹³ This sort of view is put forward by Sider (e.g. 2013, 2011) as a sort fall-back position if reference magnetism doesn’t do the job of making ordinary English determiners correspond to the perfectly natural quantifier; a more wholehearted endorsement of this sort of view can be found in Cameron (2010*a*).

¹⁴ Note that once we have a notion of fundamentality that’s not directly tied to supervenience or mereological notions, the nihilist might – and I think should – follow Cameron (2010*a*) and Sider (2013) in saying that composite objects do *exist*, it’s just that they’re not fundamental. But I’m going to put this to one side, given that nothing in this Chapter hangs on this issue and we’re already knee deep in metaontology.

is almost irresistible. But while it may be that many common or garden properties of medium-sized dry goods are indeed both dependent – i.e. supervenient – and derivative, by making conceptual room for the possibility that something could be dependent yet fundamental we are given the tools to make sense of emergence. What does it mean for a genuinely new property – such as consciousness – to emerge from a particular arrangement of micro-particles? Well, when you have some micro-particles with certain intrinsic and relational properties, a further property becomes instantiated: This property is fundamental – thus, a genuine addition to one’s ontology over and above the properties already instantiated. Yet it is also dependent on those microphysical particles for its continued existence: hence it emerges *from* them in a meaningful sense, rather than floating free of microphysical events.¹⁵

The moral is that by endorsing micro-dependence, we can retain the spirit of the original micro-determination thesis while being able to accommodate the coherence and epistemic possibility of emergent properties. Of course opponents of emergence are at liberty to deny that dependence and fundamentality come apart, but they would need to adduce non-question begging reasons for believing that if they wish to use it as an argument against emergence. As I see it, the case for adopting the metaontological framework goes hand-in-hand with the case for emergence: We need some way to explain or explain away putative cases of emergence – whether in the philosophy of mind or quantum mechanics or elsewhere. If positing genuine emergent properties turns out to be one of the better ways of doing this, we have good reason to adopt the framework, in addition to any independent reasons we may have for preferring it. Dialectically, the circle is a virtuous one: the coherence of emergence cannot be dismissed out of hand because of the availability of the Barnes framework, and the truth of the framework cannot be dismissed out of hand because of the theoretical utility of being able to make sense of emergence.

5.3. *The Plural Instantiation Strategy*

So, with a metaontological framework in place to make sense of what’s meant by an emergent property, let’s give a more precise formulation of the argument against nihilism

¹⁵Importantly, this strategy does not require you to be a realist about properties. The nominalist can just apply whatever translation scheme she usually employs when confronted with ‘property talk’ to the mention of properties here, provided she accepts the spirit of the fundamental/non-fundamental distinction. For example (following Sider 2011) a nominalist might take the distinction made here between fundamental and non-fundamental properties to be the distinction between natural and unnatural sets.

from emergence:

The 'No Bearers' argument -

P1. If a property is instantiated in the world, there exists a suitable bearer.

P2. There are emergent properties instantiated in the world.

P3. There exist suitable bearers for emergent properties (from P1 + P2).

P4. If nihilism is true, the only property bearers are mereological simples.

P5. Mereological simples are not suitable bearers for emergent properties.

P6. If nihilism is true, there are no suitable bearers for emergent properties (from P4 + P5).

C1. *Therefore:* Nihilism is false (from P3 + P6).

Potential replies to this modified No Bearers argument on behalf of the nihilist will fall broadly into two camps: On the first, grant for the sake of argument the assumption that there are emergent properties and show that the nihilist can provide suitable bearers for them (rejecting P6 via P5). In defending a such line of response, the nihilist need not be concerned, dialectally, with trying to provide their own detailed account of emergent properties or giving reason to think appeal to such properties is the best way explaining the target phenomenon: the nihilist will be able to say whatever her opponent says about emergent properties and why they're applicable to the target phenomena. If her opponent is doomed to fail in this task, the nihilist is more or less off the hook: If P2 is false, there is no direct argument against nihilism from the existence of putatively emergent phenomena. The other style of response will deny that emergent properties provide the best account of putatively emergent phenomena such as phenomenal consciousness or quantum entanglement. Once we have an alternative account, it will be argued, we'll see that we already have a ready explanation of the problematic phenomena, and so positing emergent properties would be gratuitous. If there are no emergent properties, there need not exist suitable bearers (undercutting support for P3 via the denial of P2).

Examples of this latter sort of response in the quantum mechanical case include the strategy of positing entanglement relations, and employing the ideology of metaphysical indeterminacy as an ontologically innocent stand-in for entanglement relations.¹⁶ In the

¹⁶ Again, see Schaffer (2010) and Barnes (forthcoming-b) respectively. I explore these strategies further in §6.

philosophy of mind, any theory which seeks to explain consciousness as anything other than an emergent property – e.g. eliminativism, reductive physicalism, Cartesian dualism – falls into this camp. No such strategy will be discussed in detail in this chapter. Instead, we'll be interested in the question of whether, if we grant for sake of argument that there are emergent properties, the nihilist can accommodate them in her ontology. That is, whether a suitable non-composite bearer can be found.

In the search for a non-composite bearer for emergent properties, this chapter grants for the sake of argument the assumption that a lone simple is not a suitable bearer for an emergent property. It's worth flagging up, however, that there are some interesting options worth considering before giving up on lone simples as bearers of emergent properties. After all, individual *microphysical particles* aren't suitable bearers for properties such as phenomenal consciousness or entangled state properties, but why assume all simples are microphysical particles? For instance, why couldn't whatever object in fact instantiates my consciousness – which I take for granted isn't a microphysical particle – be mereologically simple? The most persuasive objection, I think, is that simples lack the requisite internal structure to instantiate a rich, complicated macro-property such as consciousness, which is presumed to be the product of the interactions of countless micro-particles in very particular arrangements.¹⁷ That is, you might think that simples necessarily lack internal heterogeneity by virtue of their simplicity: A simple can't have a positive electric charge *here*, and a negative electrical charger *there*, for instance, because this requires having some proper part that is positively charged *simpliciter*, and some distinct proper part that is negatively charged *simpliciter*, and simples by definition have no proper parts (see McDaniel 2007). However, there have been some compelling arguments put forward in the literature to show that simples could be heterogeneous and spatially extended by possessing fundamental distributional properties: e.g. the property of being positively-charged-here-and-negatively-charged-there, where this property is not had by virtue of it or its parts possessing any non-distributional charge properties.¹⁸ If this is right – if mereologically simple things could have limitless internal heterogeneity by instantiating irreducible distributional properties – the main objection to simples not being suitable bearers for emergent properties is undercut.

So, could the object (my brain?) that instantiates my consciousness be a heterogeneous

¹⁷ See Markosian (1998) for an in depth discussion of the nature of simples – especially whether they are necessarily point-sized.

¹⁸ See Parsons (2004); McDaniel (2007); Cameron (2011).

extended simple? One that emerges from the microphysical base in much the same way as my consciousness is being supposed to emerge, as effectively a purpose-built bearer for my emergent consciousness? Although this may seem outlandish, I would need convincing that – when the proposal is properly chewed over and digested – it is any more strange or counterintuitive than the emergence of the property of consciousness itself. However, I'm postponing full discussion of this option until §6.1, in order to first pursue a possibility I take to be less outlandish on first blush and ultimately more promising: that perhaps numerous simples, acting together, could instantiate an emergent property. That is, perhaps there are emergent property instances that, instead of having a single bearer, have multiple bearers, which jointly bring about the conditions necessary for the emergent property to be instantiated. Let's call this 'plural instantiation'. As I'll be arguing, plural instantiation is a possibility we should take seriously.

Given plural instantiation, I see no reason why a plurality of simples could not jointly give rise to a single instance of a given emergent property. The factors that might be taken to disqualify individual simples from being suitable bearers of emergent properties are not present when we consider a plurality of such entities: The individual simple is not (let's grant) spatially extended, but the plurality is dispersed across the same region of space that any emergent property they collectively instantiate would be extended over. An unextended simple cannot be qualitatively heterogeneous, but the plurality can be simply in virtue of different members of the plurality possessing differing (qualitatively homogeneous) properties.

The nihilist's opponent assumes that that when we have a case of emergence, we have some simples that give rise to *some further particular* that instantiates the emergent property in question.¹⁹ On the Plural Instantiation strategy, by contrast, no further entity is posited – there are just the simples, which jointly instantiate the emergent property. They cut out the middleman, so to speak. Instantiation here is not taken to be a one-one relation between some object 'the plurality' and the emergent property being instantiated; rather it is a many-one relation between the many simples, on the one hand, and the emergent property on the other: we're not employing run of the mill one-one instantiation here. This is what allows the proponent of this strategy to posit some simples arranged brain-wise, say, that together instantiate the emergent property 'being conscious', without thereby composing some further object, a brain. It's also what makes this strategy somewhat revisionary. Indeed, this whole approach hinges on the plausibility of

¹⁹ This is assuming that she believes in a world that contains simples rather than one that is 'gunky'.

introducing plural instantiation. Do I have an argument for plural instantiation? Only that there is a dearth of arguments against it and no obvious cost involved in employing it in our theorizing. In what follows we'll try to get clear on what exactly the view entails and I'll try to deal with any potential worries or misunderstandings that arise along the way.

From the outset plural instantiation should be distinguished from multiple instantiation.²⁰ A property is multiply instantiated if it has many separate instances: Red is multiply instantiated since there are many red things. Each of these red things is a separate instance of (one and the same) redness.²¹ Plural instantiation, on the other hand, occurs when there is more than one object contributing to the very same instance of a property. For example, the fleas collectively engulf the cat, but no flea individually engulfs him: there is just one instance of the property of engulfing a cat in the vicinity of the cat. Call any property which is plurally instantiated a plural property. A property that is not plural is singular. It might help here to distinguish a collective predicate from a distributive predicate: If a distributive predicate *F* applies to some things then those things are *F* because each of those things is individually *F*. However, if a collective predicate *G* applies to some things then those things are *G* but none of those things is, by itself, *G*. If our language is wearing its metaphysics on its sleeve, then distributive predicates will predicate singular properties of things, while collective predicates will predicate plural properties of things. Being a plural property is not the same thing as being a relation, even though both involve many things participating in bringing about the one property instance. As we shall see, the *monadic property / relation* distinction crosscuts the *plural / singular* distinction, allowing for monadic plural properties and singular relations. For example, if the chairs are three metres away from the stage, it is true of each of the chairs individually that they are three metres away from the stage: this is a straightforward singular relation (and 'three metres from' is a distributive, as opposed to a collective, predicate).

On the other hand, consider the following couple of examples: First, the citizens of a democracy collectively enjoy sovereignty, while no individual citizen is sovereign. Intuitively, although we might sometimes speak of the 'The People' as if it were some singular composite entity that alone enjoys sovereignty, we want to say that it's the

²⁰ In drawing this distinction I follow Yi (2002). Elsewhere in the literature the term 'plural instantiation' is sometimes used to refer to multiple instantiation.

²¹ This may be one way of characterizing a difference between universalists & trope theorists: pluralists think that one and the same universal is instantiated many times over, whereas a trope is 'non-repeatable' in the sense that it is *not* multiply-instantiated in this manner.

individual people considered jointly that have sovereignty, not some hypothetical mereological sum of those people. Second, Jess looks up at the night sky and declares that the stars are beautiful: Asked if any individual star is beautiful on its own, she's inclined to respond in the negative – sure, an individual star can be quite pretty, but it is after all just a tiny pin prick of light. Asked if it's the *arrangement* of the stars which is beautiful, she's also inclined to answer negatively: their being arranged in a certain way in the sky contributes to *the stars themselves* being beautiful, but the arrangement itself would be nothing special if the stars were substituted for salt crystal or tea leaves. No, according to Jess it's the stars that – jointly – instantiate the monadic property of beauty. It doesn't seem like Jess's opinion is crazy or incoherent.

You might quibble about whether sovereignty or beauty are genuinely monadic properties. What we have in each case is a monadic, collective predicate— but plausibly in these cases our language *doesn't* wear its metaphysics on its sleeve, and there is no genuinely plural property being predicated. But whatever the ultimate metaphysical structure of the cases, we can make sense of 'is sovereign' and 'is beautiful' as being monadic, collective predicates in ordinary language. The claim that there are plural properties is just the claim that at least some of the time these sorts of predications correspond directly with the underlying metaphysics. That there is any place for plural instantiation in serious metaphysics, or fundamental reality, is precisely what I'm arguing in the paper as a whole, with respect to emergent properties; but you should at least be convinced that the idea of a jointly instantiated monadic property is coherent.

To help get clearer on the various distinctions at work, we can consider the number of argument places had by a predicate standing for a given property or relation. Predicates of monadic properties have just one argument place: For example '___ is beautiful' has one space to be filled. Whereas a two-placed relation, such as '___ loves ___' has two spaces to be filled. A singular property or relation has each space filled by a single object: for instance, 'The Morning Star is beautiful' and 'Troilus loves Cressida', respectively. A plural property or relation, on the other hand, has at least one argument place filled by more than one object. For example, 'the (individually unremarkable) stars are (collectively) beautiful' and 'the (individually tiny) fleas (collectively) engulf the cat'. The intuitive point is that each argument place in a relation assigns a distinctive role to the object or objects that fill that place: while a case of plural instantiation is where many objects collaborate to play the one role and so collectively fill the one argument place. To reiterate, all the putative examples of plural instantiation above have been provided just to help the reader get a fix on what plural instantiation is supposed to be: it doesn't matter for these purposes if you don't agree that any of these are genuine cases of plural

instantiation, provided you understand what is being claimed. It seems fairly obvious to me that the English language has irreducibly collective predicates such as ‘surround’ – and, indeed, ‘compose’ – but that’s also orthogonal to the debate here. What’s at issue is whether the nihilist is free to introduce plurally instantiated properties into her ontology in order to meet the challenge from emergence.

Given all this, I see no conceptual barrier to the nihilist holding that the many simples are the collective bearers of a single, monadic emergent property. Perhaps one outstanding worry is that properties such as consciousness are *intrinsic* – that is, not possessed in virtue of the relations that one stands in to other things, or something that could be had by a lone object.²² You might grant that we can coherently have monadic (i.e. non-relational) instances of plural instantiation, but hold that these are necessarily *extrinsic* properties. Plausibly, the examples I gave above of plural monadic properties such as sovereignty and beauty might somehow be construed as extrinsic. However, the objection is founded on two unwarranted assumptions— that we have good reason to think consciousness is intrinsic, and that we have any reason at all to suppose intrinsic properties can’t be plurally instantiated. Is consciousness intrinsic? When doing serious or fundamental ontology, there’s no reason to think our intuitions are any sort of reliable guide to telling us what properties are intrinsic or extrinsic— that’s to suppose we have some privileged pre-theoretical insight into the ultimate nature of things.²³ Given the nihilist has good theoretical reason to posit plurally instantiated consciousness, then added to the (unwarranted) assumption that intrinsic properties can’t be plurally instantiated, this is just reason to suppose that consciousness is extrinsic. Our pre-theoretical hunches count for nothing in the face of evident theoretical utility. Even then, what would be lost by having to say that consciousness is extrinsic? You might worry that *paradigmatically* extrinsic properties such as ‘being such that one is six-feet from a hamster’ are not genuine properties, or not ontologically heavyweight, while consciousness is. But we already have an explanation, from the previous section, of why properties like ‘being such that one is six-feet from a hamster’ are ontologically suspect: such properties are not fundamental— they come for free when the spatial relations between things are fixed. But, *ex hypothesi*, emergent consciousness doesn’t come for free when the spatial relations between things are fixed even if it depends on such relations holding: whether we call it intrinsic or extrinsic, it is nonetheless fundamental.

²² See Lewis (1983a) for a detailed characterisation of the distinction.

²³ For a detailed argument for this point see Sider’s (2001) discussion of the problem of temporary intrinsics.

All we have here is a reason for thinking the intrinsic/extrinsic distinction comes apart from the fundamental/derivative distinction, just as the believer in emergence must say is the case with the independent/dependent distinction.

Can intrinsic properties be plurally instantiated? If you take an intrinsic property to be one that can be instantiated by a lone object, this is objectionably arbitrary when plural instantiation is brought into the picture, since it just begs the question of whether plural properties are intrinsic: To be principled, we should amend the definition to say that an intrinsic property is one that an object *or objects* that bear the property in question can have when otherwise unaccompanied— that is, if everything apart from the bearer(s) of the property were removed from the world, the bearer(s) could still instantiate that property. This is a criterion plurally instantiated consciousness can meet (while paradigmatically extrinsic monadic properties like ‘being such that one is six feet from a hamster’ would not meet it): Only the simples in my brain are required for me to be conscious, and these are the things that (collectively) *are* conscious— nothing else need exist.²⁴ Much the same goes if you cash being intrinsic out in terms of not being had in virtue of one’s relations to other things: It’s usually assumed that a thing’s *parts* are to be excluded from consideration here, which is how the believer in emergence can hang on to the idea consciousness is intrinsic despite, by her own lights, it depending on the brain’s relations to certain microphysical particles. But if the parts of a bearer should be excluded from the analysis, then fellow co-bearers certainly should be: The relations to the parts are excluded because they’re assumed to contribute to the intrinsic nature of the bearer; but in a case of plural instantiation the intrinsic character we’re interested in is the intrinsic natures of all the co-bearers *considered collectively*. The collective nature of the simples, *qua* co-operative enterprise, can’t be fully grasped without considering all the properties of the things co-operating, including the relations they bear to each other— while relations they bear to things not co-operating in bringing about the property in question can be legitimately excluded.²⁵

Finally, you might worry that including plural properties in your ontology requires

²⁴ Given somewhat idealized circumstances that the believer in composition must assume too if she wants to hold consciousness is intrinsic (i.e. that brains / persons / simples-so-arranged can survive unaccompanied).

²⁵ I’ve also been asked: “*Where is* the intrinsic character of the plurality, given that the plurality does not compose some further object?” To this I merely reply that the intrinsic character of the simples is where the simples are, given that the intrinsic character of the ‘plurality’ just is the intrinsic character of the simples taken collectively.

significant revision to your logic. However, as Byeong-Uk Yi (1999, 2002, 2005, 2006) and others²⁶ have demonstrated, we can accommodate plural properties simply with the introduction of plural quantifiers and plural predicates to elementary logic (PFO+). Given the relative ease with such predicates can be accommodated, I don't think there's any sustainable objection to the effect that the very idea of plural properties is confused, or that formulating a theory in a language that contains plural predicates would be overly arduous or unattractively revisionary. That is, I don't think a metaphysics containing fundamentally collective properties or plural instantiation can be dismissed out of hand.²⁷

5.4. *The Cost of Plural Instantiation*

We've seen that the nihilist can accommodate emergent properties in her theory if they are collectively instantiated properties. But just because the nihilist is at liberty to introduce plural properties into her theory without falling into conceptual confusion or straying beyond the philosophical pale, that doesn't necessarily mean it would be prudent to do so. Such a move may still come at a cost: the price of introducing plural instantiation into your theory may be one that the nihilist would be unwilling, or unwise, to pay. How might that price be exacted? Plausibly through reducing the ontological or ideological simplicity of the nihilist's theory. We'll assess each possibility in turn. In each case, we'll see that at the centre of the controversy is the strategy's reliance on plural quantification. The nihilist needs to be able to quantify plurally over *some things* in order to say of them that they satisfy a given collective predicate and thus plurally instantiate an emergent property. For, *ex hypothesi*, it's not the case that any individual thing instantiates an emergent property, and so it's not true of any individual thing that might be picked out by a singular existential quantifier that it satisfies a collective predicate.

Many have worried that plural quantification is not ontologically innocent.²⁸ The worry is that our best metaphysical theory is not only ontologically committed to each thing in the domain of its singular existential quantifier, but also to every 'plurality' in the domain of its plural quantifier, should the theory contain one. The thought is we should extend the Quinean criterion of ontological commitment²⁹ to cover not only singular

²⁶ E.g. Morton (1975); Mundy (1989); Rayo (MS).

²⁷ Byeong-Uk Yi provides arguments for adopting plural instantiation which primarily concern its ability to give an attractive metaphysical account of numbers.

²⁸ For a long list of doubters see Linnebo (2004).

²⁹ See Quine (1948).

variables bound by the singular existential quantifier, but plural variables bound by the plural quantifier. I take it the primary motivation here is the feeling that it would somehow be unprincipled or arbitrary to refuse to extend the Quinean criterion to new quantifiers you introduce into your language: after all the same thing is going on in both cases; variables being bound by existential quantifiers.³⁰ However, this puts the cart before the horse as far as ontological commitment is concerned. If anything is going to be a primitive in ontology, it's going to be the notion of existence itself. Being is not going to be analysed away: to be is not *literally* to be the value of a variable. What I take to be attractive about the Quinean criterion is that it picks out all those entities that play some role in our fundamental theory and are therefore indispensable to it: If we felt the need to predicate something of an object, then it clearly plays a role in our theory. On the other hand, if we didn't predicate anything of it, it's not clear what role it could possibly play in the theory, and therefore why the theory should acknowledge its existence. The criterion does a nice job of counting up all the things we would intuitively suppose need to exist in order for our theory to be true.

However, things get more complicated when we bring plural quantification into the picture. Thinking for a moment about Yi's (1999) account of numbers as plurally instantiated properties might help to bring this out. Given plural properties, incurring a plural commitment to *some things that are (that is, have the property of being) seven* is – intuitively, at least – to incur a commitment to seven things, not to one thing, the 'some things' that are seven, or to eight things – the seven thing plus the 'some things'. This is supposed to be what distinguishes the plural strategy from various surrogate strategies that attribute the singular property of seven-ness to some further entity, in virtue of its relationship to the seven things, such as having them as parts or members. Seven-ness, on the plural view, is something that can only be instantiated jointly by seven entities. But if we're already committed to the seven things for independent reasons – for instance, because our theory says each of them possesses certain singular properties – then we shouldn't count them *again* (plurally this time, rather than individually) when we attribute to them, jointly, the property of being seven. The point is not that the plural quantifier isn't ontologically committing— I think it is. If you only had a plural quantifier in your ideology you could still read your ontological commitments off that quantifier alone: being ontologically committed to a plurality of seven things is not *less* ontologically committing than being ontological committed to seven things via singular quantification— so, strictly speaking it's wrong to say the plural quantifier is ontologically innocent. Rather, the point

³⁰ Compare Boolos (1984).

is that if you already have the singular quantifier in your language, you don't incur extra ontological commitments just by adding another quantifier: both are quantifying over the same beings.

I think the really pressing challenge the nihilist needs to face is not that plural quantification complicates her ontology, but that it complicates her ideology: Ideological commitments are incurred by adding primitive vocabulary to the language of your fundamental theory, and the plural existential quantifier needed by languages containing plural instantiation is just such a primitive— it is interdefinable with the plural universal quantifier, just as the singular existential quantifier is interdefinable with the singular universal quantifier, but is not definable in terms of the singular quantifiers (see Yi 2002).³¹ Given a growing acceptance in the recent metaphysics literature³² that ideological commitments are on a par with ontological commitments, it would be hard for the nihilist to argue that any real advantage is gained through absorbing the cost in her ideology rather than her ontology. If the nihilist has to expand her ideology in order to employ the Plural Instantiation Strategy, this cost must be weighed against the gains the nihilist claims for eliminating parthood from her ideology.³³

If we're to believe the argument of Bennett (2009), it's a cost that's not really worth paying, because it leaves the nihilist in a position that is no more attractive in terms of theoretical virtues than that of the believer in composition: the nihilist is simply pushing around a bump under the carpet. For a form of the Plural Instantiation strategy, along with its reliance on plural quantification, has been around as an option for the nihilist for as long as the Special Composition Question in its present form, in the guise of van Inwagen's paraphrase strategy. When van Inwagen recommends the nihilist paraphrase

³¹ An alternative, arguably more perspicuous way, to look at the extra cost involved here is that the plural quantifier is more internally complex than the singular quantifier: as well as having a primitive logical predicate of *identity* the believer in plural quantification needs the primitive predicate *is one of*, to denote that a given object belongs to a plurality. So, the real cost of plural quantification is adding more identity-like structure to the theory (the 'is one of' relation), rather than more quantificational structure as such (for once one has the more expressive plural quantifier arguably one does not need the singular quantifier too). However, I think what follows will play out much the same regardless of your view on the technicalities of where precisely the extra commitment lies.

Alternatively, it might well be argued that the believer herself is already committed to plural quantification simply through belief in composition, which is defined by van Inwagen (1990) using plural quantification: I set this aside for the sake of charity to my opponent, space being limited, but this cheap dialectical line may well be worth pressing.

³² E.g. Barnes (forthcoming-b), Bennett (2009), Sider (2011), Turner (2011).

³³ Sider (2013) argues for the view that eliminating parthood is an ideological economy.

the sentence ‘there is a table’ as ‘there are some things arranged tablewise’, the upshot is to reject the surface logical form of ‘there is a table’ – i.e. that there is a thing which instantiates the singular predicate ‘is a table’ – in favour of a regimentation under which one plurally quantifies over *some things* and says of those that they together satisfy the collective predicate ‘are arranged table-wise’ (van Inwagen 1990). If one takes the nihilist’s fundamental theory to consist of sentences in the paraphrased form, then the nihilist incurs a commitment to plural quantification. However, according to recent defenders of nihilism such as Cameron (2010a) and Sider (2013), and the position defended earlier on, the nihilist need not provide a systematic paraphrase of every true English sentence putatively committed to composite objects— Bennett’s objection can thereby be avoided with respect to plural quantification. Yet, if the nihilist wants to avail herself of the Plural Instantiation strategy to account for emergent properties, then it seems that she needs to re-introduce the apparatus of plural quantification into her fundamental theory and be vulnerable to the same ‘bump-pushing’ objection that Bennett levels against the paraphrasing nihilist.³⁴

One way for the nihilist to respond to the argument of pushing the bump around is just to concede that there’s no ideological parsimony gain to be had by swapping the parthood relation for plural quantification, but this doesn’t diminish the importance of the ontological parsimony gained by eliminating composite objects from her ontology: Belief in composition comes with both an ontological commitment to composites and an ideological commitment to parthood, whereas nihilism with plural quantification at least reduces our ontological commitments – since, as I argued above, adding a plural quantifier to a language with singular quantification is ontologically innocent, even if the plural quantifier itself is ontologically committing. Now, I think this nihilist rightly points out an important gain in ontological parsimony, but is needlessly concessive concerning the ideological cost of plural quantification. The ideology objection all rests on what, exactly, the ideological cost of introducing the apparatus of plural quantification into your theory is, and I think the nihilist should say that the cost – if any – is negligible when you already have singular quantification in your ideology. The argument is structurally similar to the one I ran above concerning ontology: By having singular quantification in your ontology is already to incur the cost of having quantificational apparatus in your ontology, so although plural quantification is a significant ideological cost, it is not a significant

³⁴ Why didn’t van Inwagen himself use the ideology of plural quantification to account for consciousness? Ultimately, I think, because he was impressed by the objection – dealt with in the previous section – that intrinsic properties can’t be plurally instantiated (see Inwagen 1990, p.118).

additional cost for the nihilist relative to the believer in composition.

As Sider (2013) points out, and Goodman (1951) before him, not all ideological costs are created equal: However, if the nihilist wants to say that the cost of introducing a plural quantifier into her ideology is somehow minimal or negligible the burden of proof is on her. Such arguments can be tricky, given the many seemingly intractable issues surrounding theory choice methodology in general, but I think the intuitive case that can be made at least manages to shift the initial burden of proof away from the nihilist and onto her opponent: Let's say you had no quantificational apparatus in your theory at all: adding some quantificational apparatus into your theory would greatly enhance its power for a corresponding reduction in simplicity. However, let's say you're like Kris McDaniel (2010) or Jason Turner (2010), a fan of there being many ways of being— e.g. you think numbers exist in a different way from chairs (abstractly vs. concretely), or spacetime points in a different way from tropes (absolutely vs. relative to a spatial location), where this is not merely down to these things having different properties, but in literally existing in a different way. Turner and McDaniel both cash this ways of being talk out in term of there being different fundamental quantifiers in one's ideology corresponding to each way that something could exist: thus, rather than having just the one fundamental quantifier, \exists , that ranges over everything, they have two quantifiers, \exists_c and \exists_a , say, that range over concrete and abstract things respectively (\exists can be defined in terms of \exists_c and \exists_a , but it is not – or does not have to be – part of their fundamental ideology). Adding a quantifier with a subscript shouldn't count as a significant addition to your ideology when you already have a quantifier in your theory: You already have quantificational apparatus in your language, you're just re-tasking that apparatus in a slightly different way. I would say it still *is* a complication to your ideology compared to the person with just one quantifier, but it is not of the same order as adding in quantificational apparatus to a theory that lacks it entirely: By adding another quantifier when you already have one is not to *double* your ideological commitments related to quantification.

One way of further motivating this thought – at least in the particular cases under discussion – is that we're multiplying tokens, not types. Many have followed Lewis³⁵ in thinking that, when it comes to ontology, qualitative parsimony matters, but quantitative parsimony does not— adding a new *type* of entity to your theory counts against it when judging its simplicity, while adding more entities of a kind already mentioned in your theory doesn't make it any less parsimonious. One way in which you might apply the

³⁵ e.g. Lewis (1973, p.87) – though compare Nolan 1997.

qualitative / quantitative distinction found in the discussion on ontological parsimony to ideology³⁶ is to say that adding a new bit of ideology to your theory is analogous to adding a new type of thing to your ontology, whereas the number of times you employ that ideology in your theory is analogous to adding a new thing of the same type to your ontology – e.g. the difference between adding a primitive modal operator or not, on the one hand, and saying of a thousand and one things rather than of a thousand things that they are possible, on the other. There maybe some mileage in this analogy, but here’s another, perhaps equally useful one: Adding a bit of ideology that is structurally different is adding something of a new type, whereas, adding a structurally identical bit of ideology is just adding more of the same rather than adding something of a new type. By structurally identical, I mean in the way – for instance – a distribution might be said to be structurally analogous. If I have three cakes and you have one, that distribution is structurally identical to the distribution in which you have three cakes and I have one: the form or type of distribution is just the same, the only difference is in what individuals fill the roles in that distribution. A new quantifier is just more quantifiers: it doesn’t matter whether the old one ranged over concreta but the new ranges over abstracta, or – indeed – whether the old one ranges over individuals and the new one ranges over pluralities.³⁷

So, is plural quantification as big a cost as having composition in your ideology?³⁸ As I’ve argued that the cost of introducing plural quantification into your ideology is negligible if you already have singular quantification in your ideology, the cost of having parthood in your ideology rather than plural quantification will be significant, unless the cost of adding parthood to your ideology is itself negligible. But there’s no reason to think that having parthood in your ideology is negligible – in fact, we have good reason to think that the cost is significant. The nihilist doesn’t have *any* mereological primitives in her language (Sider 2013), whereas the believer in composition needs at least one, and the jump from none to one is the most ideologically costly: it’s the difference between having

³⁶ A view I’ve come across a number of times in conversation.

³⁷ A very different way of arguing for this conclusion – which I’m not recommending, but is worth drawing attention to – is to motivate a retreat from the very austere logical realism of Sider (2011). One could still think that many ideological commitments are a theoretical cost – such as a primitive modal or tense operators, say – but hold the more general background features of are logical, such as identity and quantification, are not *wordly* features as such, but just necessary artefacts of theorising or thought, and therefore innocent.

³⁸ I’m assuming here that ‘compose’ is not a primitive in the believer’s theory – since composition itself is a plural relation of the parts to the whole – and that it can be explained in terms of a singular relation like parthood.

no mereological structure in your theory at all and having a whole domain of mereological facts involving this primitive. It's the equivalent of adding your first quantifier to your language to create a theory capable of expressing rich quantificational structure; or adding a primitive necessity or tense operator to your theory, in order to accommodate a domain of modal or temporal facts: once you have one primitive in a domain the expressive benefits of adding more are limited, and the ideological cost correspondingly more modest. If the believer in composition thinks their favoured mereological primitive is ideologically innocent it's up to her to say why— and she can't appeal to the same reasoning that the nihilist has available to show that plural quantification is no additional ideological cost.³⁹

To sum up, then, the plural instantiation strategy provides the nihilist with a way to answer the challenge from emergence that doesn't rely on the empirically vulnerable premise that there is no such thing as emergence. By accommodating emergent properties in her ontology, this strategy also allows the nihilist to enjoy any of the theoretical benefits claimed by the believer in emergence. The cost of plural instantiation is its reliance on plural quantification: But I've argued that if you already have singular existential quantification in your language you don't incur any extra commitments - ontological or ideological - by introducing plural quantification. So, in so far as there is an ideological or ontological advantage to being a nihilist (which is not for us to adjudicate here), this is not eroded or cancelled by employing plural instantiation in order to account for emergent properties.

In the next chapter, I explore three potential alternatives to the plural instantiation strategy— emergent simples, entanglement relations and primitive metaphysical indeterminacy.

³⁹ I expand a bit further on my view of ideological parsimony, including responding to another potential objection, in §8.2.

6. Extended Simples, Entanglement Relations and Indeterminacy.

In the last chapter I argued that the Barnes (2012) framework for ontological emergence provides an attractive account of putatively emergent phenomena in, for example, quantum mechanics (§5.2). According to the No Bearers Argument (see §5.3), the nihilist can't accommodate ontologically emergent properties in her ontology, since they are taken to require mereologically complex bearers (which the nihilist eschews): I demonstrated, however, that the nihilist could respond by positing ontologically emergent properties plurally instantiated by multiple simple bearers.

In this chapter I explore three alternative strategies for accounting for emergence given a nihilist ontology. The first (§6.1) – extended simples – is akin to the plural instantiation strategy in that it finds a substitute bearer for genuinely emergent properties, however it posits additional ontology rather than providing a non-standard account of instantiation. The next two strategies – entanglement relations (§6.2) and a novel approach employing the ideology of primitive metaphysical indeterminacy (§6.3) – take a more reductive approach to the putative emergent phenomena in quantum mechanics, explaining them by means of alternative ontology or ideology, respectively, rather than by positing genuinely emergent properties.¹

Each view could have been given a chapter to itself, but for reasons of space I simply give a brief overview of each strategy, following up by responding to what I take to be the major objections to the strategy in question. As a consequence the coming discussion may seem a bit piecemeal: I try to offset that in my closing remarks by making some comparative comments on my preferred strategy of plural instantiation and the three alternative strategies proposed here. In the final analysis, the more viable ways the nihilist has to account for putative ontological emergence, the stronger her dialectical position will be – which strategy the nihilist chooses to employ will likely depend on her wider

¹ It remains the assumption of these approaches, and indeed an assumption of Part II of this thesis, that *some* extra expressive resources are needed to explain putatively emergent phenomena. If this is not the case – if, once properly understood, it turns out nothing out of the ordinary is going on in these cases after all – then the nihilist can rest easy with the dialectical situation at the end of Part I.

theoretical commitments.

6.1. Emergent Simples

The emergent simples strategy² does what it says on the tin: It posits emergent simples to be the bearers of emergent properties. This approach takes what the proponent of the No Bearers argument assumes about the need to posit ontologically emergent properties at face value. According to the view, whenever the activity of some simple micro-particles gives rise to an emergent property, they also give rise to a suitable bearer for this property. Importantly, this emergent entity does not have the micro-particles as parts: it is also mereologically simple and therefore its existence is compatible with the truth of nihilism.

At this point, it would be legitimate to wonder how the nihilist's positing yet more simples is going to help matters. The original problem³ was precisely that simples aren't suitable bearers for the sorts of emergent properties we're considering. However, the proponent of emergent simples is going to put pressure on this assumption: Sure, individual *microphysical particles* aren't suitable bearers for properties such as phenomenal consciousness or entangled state properties, but who said that all simples are microphysical particles?

The proposed emergent simples are very much like the composites that they're standing in for: A given emergent simple takes up at least as much space as the mereological fusion of the microphysical particles that give rise to it would, and has all the same physical qualities as the fusion would— making it an ideal candidate to instantiate the relevant emergent property instead of the mereological sum of the microphysical particles. The only really important difference is that the emergent simple is, well, simple. But why should the mere lack of proper parts disqualify an object from being a suitable bearer for an emergent property? It's obvious that there's a big difference between an entangled state of two particles, and just one of those particles taken on its own— plausibly a deep enough difference to exclude the latter, but not the former, from instantiating emergent entanglement properties. But why should it matter if the 'entangled system' is a composite object or a heterogeneous extended simple? If A and B are qualitatively indiscernible, is the fact that A is mereologically simple and B composite really relevant in determining what further properties A and B could each instantiate? It certainly seems that if, in a given case, B was barred from instantiating a property that A could instantiate,

² First discussed in §5.3.

³ See §5.1 - §5.3.

we'd want to be given a detailed explanation as to why. With this in mind, let's look at a few objections to using extended simples as bearers of emergent properties.

6.1.1. *Objection 1: Simples can't be extended*

If emergent simples are going to stand in for composite objects or systems, as bearers of emergent properties, then presumably they must be located in the same place as the composites are according to the Composites strategy. But if you hold a certain intuitive view about what it takes to be mereologically simple, emergent simples are not going to be able to fulfil this requirement:

*The pointy view of simples: x is simple iff it isn't spatially extended.*⁴

If the pointy view is right, emergent simples can't stand in for composite objects as bearers of emergent properties, since they can't fill the same space as those composites would. If an emergent system e is dependent on particles a , b , c and d the believer in composition can say the system is partly located where each of its parts are. What is the nihilist to say? That e is arbitrarily co-located with one of a , b , c or d , or arbitrarily located at some point between them? Or, incredibly, that e is multiply co-located, wholly present where each of a , b , c and d are?

I grant that things are looking messy, given the pointy-view of simples. But why hold it in the first place? Take the basic objects in the nihilist's ontology - those simples that are fundamental and independent, rather than emergent. There's really nothing in the nihilist's core commitments that force her to say that basic simples are point-like— for all nihilism tells us, every simple could be spatially extended.⁵ Therefore, one shouldn't see this as an *ad hoc* loosening of a firmly established principle, on behalf of any nihilist who would adopt the emergent simples solution. Sure, those nihilists who *do* hold the pointy simples view would have to abandon one direction of the bi-conditional to adopt this response. But what it takes, in non-mereological terms, to be a mereological simple is already a controversial issue that crosscuts, at least to some extent, the range of answers to the Special Composition Question.⁶ If extended simples are needed to get a nihilist ontology off the ground and the nihilist already takes herself to have independent reason to believe in a nihilist ontology, then she has a dialectically legitimate reason to reject the

⁴ See Markosian (1998) for the original and more careful formulation.

⁵ For proponents of extended simples see, e.g., Markosian (2004), Simons (2004), Miller (2006), Cameron (2011).

⁶ Again, see Markosian (1998).

pointy view of simples.⁷

6.1.2. *Objection 2: Heterogeneous properties*

Even if the coherence of extended simples is taken as a given, however, there is an obvious worry about the suitability of extended simples as bearers of emergent properties. Extended or not, by definition an extended simple has no parts. And having parts may be thought to be a pre-requisite for instantiating the sorts of properties that a simple substitute for a mereologically complex system would have to instantiate. Specifically, when an extended object exhibits heterogeneous features – such as being charged in one region but not in another – we are used to explaining these features in terms of the homogeneous properties of its parts.

For instance, take a hypothetical emergent system, Chess Board. Chess Board is ontologically dependent on thirty-two white squares and thirty-two black squares. Chess Board has some notable features: not only is it black in some places and white in others, but displays a distinctive black-white chequered pattern in regions larger than a few squares. There is an easy explanation available as to how and why the system displays these qualitative features if we consider the chessboard to be a composite object composed of the sixty-four squares. The chessboard is white in some regions and black in others because it has black parts and white parts: its blackness and whiteness is distributed in a black-white chequered pattern because of how these parts are arranged relative to each other. On the other hand, if we reject the Composite strategy and hold that Chess Board is mereologically simple, how are we to account for its black-and-white chequered nature?

We could posit fundamental distributional properties. That is, instead of holding that all heterogeneous property distributions are merely derivative, reducible to arrangements of homogeneous property instances, we hold that for some heterogeneous property distributions there is a corresponding fundamental property. This fundamental distributional property alone explains why objects that instantiate it have the qualitative features they do, without reference to the qualitative features of any proper parts or sub-regions of that object. If we're willing to buy into fundamental distributional properties, then extended systems can exhibit heterogeneous features without needing to be mereologically complex: given this, it is hard to see what other principled objection there could be to emergent simples being the bearers of emergent properties in the stead of

⁷ See e.g. McDaniel (2007) for a defence of extended simples.

composite wholes.⁸

Of course, it might be objected by the nihilist's opponent that belief in fundamental distributional properties is implausible, incomprehensible, or at the very least a cost to one's theory. With regard to the former two charges, I'm not aware of any strong arguments to support them. With regard to the relative cost of belief in fundamental distributional properties for the nihilist, the really important thing to note is that the proponent of the composites strategy is probably going to need them too, simply by being committed to the existence of emergent properties.

Putting your foot down and refusing to acknowledge fundamental distributional properties when you rely on the existence of emergent properties looks, dialectically, like a very uncomfortable position indeed. For starters, you've already accepted the principle that objects can have properties that are not reducible to the properties of their parts—probably the hardest thing to swallow about fundamental distributional properties. Further, if you want these emergent properties to be heterogeneously distributed across composite systems, you need to employ fundamental distributional properties: Why? Well, since the emergent property itself is not a mere derivative of the properties of the system's parts, it can't be that having an emergent property E distributed D-ly across region R merely consists in having parts with such-and-such properties D-wise arranged within R: If that were the case, there would be no need to posit a further fundamental property, E, in the first place, since the object's being E would entirely consist in it having certain parts arranged in certain ways—the system would get to be E derivatively, in virtue of the fundamental properties of its parts, which is *ex hypothesi* impossible.

So, it turns out the nihilist who endorses the Emergent Simples strategy is on exactly the same footing as the believer who endorses the Composites strategy. If the believer in composites thinks that emergent properties can be spatially heterogeneous, they'll need the ideology of fundamental distributional properties. In which case, the nihilist can co-opt this ideology to explain how emergent simples can have any spatially heterogeneous fundamental properties that may be required. If, by contrast, the believer thinks emergent properties are never spatially heterogeneous the nihilist is no worse off than the believer for saying the same. Granted, perhaps, e.g., having consciousness is not a genuinely spatial property, instantiated at a specific location — maybe your emergent consciousness is

⁸ It's suggested that fundamental distributional properties could be instantiated by extended simples in McDaniel (2007), Schaffer (2010) and Cameron (2011). Parsons (2004) provides a detailed explication and argument for irreducible distributional properties.

supposed to float free of spacetime in the realm of Cartesian souls, rather than being distributed across the matter that gives rise to it. In that case, the nihilist only needs only to account for how extended simples can have heterogeneous *derivative* properties: but however she chooses to fill in the details of this account, it shouldn't require any complication of her *fundamental* ontology or ideology, since by definition she's giving an account of something non-fundamental. The nihilist can't tell the story of how the system gets to be red, derivatively, in virtue of it having red parts. But what would be wrong, in principle, in telling a story where the system got to be red, derivatively, in virtue of depending on ("emerging from") a bunch of smaller red objects under such-and-such conditions?

6.1.3. Objection 3: Why not just have composition?

Is the nihilist response here skirting dangerously close to belief in composite objects, just by another name — so close, in fact, that any supposed benefit of adopting nihilism over belief in composition will be negated?

The short answer to the first part of the worry: Whatever's going on with emergent simples it simply *isn't* composition. While composition obeys a fairly well established core set of axioms, and has many intuitively familiar features, I take it that many (metaphysically and epistemically) possible instances of emergence will diverge radically from this paradigm. For instance, it may be that larger entities beget smaller entities, that one entity gives rise to many, or that an emergence base and its dependents do not overlap or share any of their properties in common.

Further, all these parameters will vary depending on the particular emergence principles at work. In familiar modal terms, the basic thought is this: The rules governing composition are necessary, whereas the rules governing the emergence of entities are highly contingent. Thus, unlike composition, these principles are not in the "metaphysical bedrock", they are something more analogous to physical laws. Though whether something akin to this distinction can hold up if you don't think modality is fundamental,⁹ or if you think composition is contingent,¹⁰ I'm not sure. But if you believe either or both of these things then your views on composition are pretty far from the orthodoxy, which holds that the answer to the special composition is necessary,¹¹

⁹ e.g. Sider (2012)

¹⁰ e.g. Cameron (2007)

¹¹ For instance van Inwagen (1990), among many others.

and which treats this as a substantive – rather than merely conventional – claim. My real target in this chapter is the more orthodox notion of composition: what to say about whether belief in emergent objects amounts to belief in composition *in some non-orthodox sense* is perhaps a question for another time.

Arguably, though, the real issue is not whether or not adopting this strategy *amounts* to belief in composition, but whether adopting it is just as costly, in terms of ontological and ideological commitments, as a belief in composition. If you take parthood to be a fundamental piece of ideology, embedded deep in the world's structure,¹² then I don't think the believer in emergent simples is committed to anything like that— 'being emergent from' need not be defined in one's fundamental ideology for the strategy to get off the ground – following Barnes, *this* notion is to be understood merely by reference to fundamentality and dependence.¹³

6.2. Entanglement Relations

The Entanglement Relations strategy is in a similar vein to the Joint instantiation strategy, in the sense that it also posits fundamental qualities with multiple, ontologically independent bearers as their instantiation base – however, instead of plurally instantiated monadic properties, variably polyadic external relations holding between simples are employed, wherever there is a putative case of genuine emergence. As usually set out,¹⁴ however, the Entanglement Relations strategy is not a straightforward response to the No Bearers problem, simply employing a polyadic relation between the parts where the believer would employ an emergent monadic property of the whole. Instead, it uses the new external relations it posits to increase the expressiveness of the theory in such a way as to provide a more reductive account of emergent properties in terms of accounting specifically for the quantum correlations which give rise to putative cases of emergence in

¹² As Sider (2012) thinks it would be, if composition were actual.

¹³ A final objection, which doesn't get its own sub-subsection, but will in a sense get its own chapter, is that once we admit an ontology of emergent simples traditional pluralistic nihilism will quickly collapse (e.g. for reasons of quantitative ontological simplicity) into existence monism – the view that the world is just one big heterogeneous extended simple. While still technically nihilism, this view might strike some readers as implausible— however, it receives an extensive defence in Cornell (MS.). See §7 of this thesis for a detailed discussion of the relative merits of monism and pluralism.

¹⁴ As in, e.g., Schaffer (2010).

QM.¹⁵ In this sense it is a kindred spirit of the Indeterminacy strategy discussed in the next subsection.

The upside of taking a deflationary approach is that one does not have to be on board with the metaontological approach set out by the believer for understanding emergence—in this sense it is more neutral. The downside of taking a deflationary approach is that it gives the proponent of entanglement relations more work to do than the proponent of joint instantiation: not only does she need to defend her new external relations against the objections considered below, she also needs to provide a workable deflationary account of emergence, rather than just piggybacking on whatever the believer in composition says.

How, then, is the Entanglement Relation strategy supposed to work? The case for emergence gets purchase in the quantum case (at least arguably) because some systems involving multiple particles exhibit behaviour that demonstrably fail to supervene on the intrinsic properties of the participating particles and the spatiotemporal relations between them. From this, it is determined that the system as a whole must have some fundamental property, perhaps independent from – though more likely partially dependent on – the properties of the particles. This follows pretty irresistibly if we assume that all the fundamental facts supervene on the fundamental objects, what fundamental properties they have and the fundamental relations they stand in: If there are facts about the larger system that are not fixed by the properties of simples, their intrinsic properties and the spatial relations they stand in, then it must be some fundamental intrinsic property of the larger system that is at work in grounding the additional facts. That is, unless some fundamental external relation is at work between the simples besides spatiotemporal relations – which is the route adopted by the believer in entanglement relations.

For illustration take two particles, *a* and *b*. On measurement at *t*, it will be the case that only one of the particles will display property F, but at *t* it is indeterminate which will be F at *t*: either *a* will be F at *t* or *b* will be F at *t*, but they won't both be F at *t*. We might say that the system of particles *a+b* has an “anti-correlation” property – its two components, *a* and *b*, won't have the same property on measurement. Thus, this simple model is supposed to capture the spirit of quantum entanglement cases and the problem they are supposed to pose for believers in micro-determination, pushed by – for example –

¹⁵ As such, it does nothing to defuse the challenges from emergent consciousness (though some may find such challenges less pressing) and it becomes an open interpretational matter whether entanglement relations, through the strategy of explaining problematic correlations, can account for *all* putative cases of emergence in QM.

Jonathan Schaffer,¹⁶ who says:

No matter how far apart the particles are, a spin measurement on one will immediately set the spin state of the other to the opposite (since the spins are anticorrelated). Entangled particles seem as if telepathic. They act as a unit.

The point of positing fundamental entanglement relations here is to provide an ontological ground for this seeming ‘telepathic’ link. On the view that the only fundamental relations are spatial relations, that the two particles could act as a unit across great distances without any intervening causal mediation is hard to explain: is it just a gigantic cosmic conspiracy that these two particles are so reliably correlated in their observable properties? But if we are prepared to admit that the two particles could be – so to speak – *directly linked* by some fundamental external relation, we have a *prima facie* explanation of how the correlation comes about (albeit a very sketchy story that would no doubt come with a promissory note of further details to be filled in later). The basic idea is that if a particle *a* is connected by a fundamental anti-correlating relation with respect to F and G to another particle *b*, it will always – of physical necessity – be F when *b* is G, and be G when *b* is F. Entanglement relations are, in effect, playing the role of truthmakers for quantum anti-correlation facts.

Schaffer’s Objections

The following three objections are put forward in Schaffer (2010). Since the target of all three objections is the idea that emergent systems can be accounted for by using a plurality of basic simples as bearers, some of the objections may also be applicable to the Joint Instantiation strategy that was the focus of the previous chapter (what I say in response here shall also go for the Joint Instantiation strategy):

6.2.1. Objection 1: No smallest particles

Schaffer points out that there are some indications that future physics – or the best interpretations of our current physics – that may abandon particles altogether. In which case, the argument goes, there will be no simples at the basic level to participate in entanglement relations or the joint instantiation of emergent properties. Obviously, this objection cuts a lot closer to the bone for the nihilist than for your run-of-the-mill pluralist: the nihilist needs there to be simples, because she doesn’t believe there’s (really) anything else. Still, although the objection may have much wider and more worrying

¹⁶ Schaffer (2010, p.50).

implications for the nihilist's theory, here is as good a place as any to try and tackle it.

The evidence Schaffer adduces is primarily from attempts in the philosophy of physics literature to interpret the ontological significance of Quantum Field Theory (QFT). In particular, he argues that in QFT there is 'no fact' of the matter concerning the number of particles that exist at any one time. Now, the interpretation and metaphysical significance of QFT is too vast a topic for us to tackle here. So I just want briefly flag up a number of ways in which it seems open to the nihilist to respond: Fleshing out any one of these responses would be a lengthy and substantive project itself, and would take us too far afield from the main thrust of this chapter.

For starters, I don't think it's an especially bad thing for the nihilist to admit that her theory is hostage to empirical fortune - outside of metaphysics this is usually considered a straight-out virtue. Does the fact that our current best science may be hostile to ultimate particles give us some warrant to be sceptical about their existence? Yes, I think so. But whatever the strength of this warrant, it's not decisive: Firstly, we don't know that QFT will resemble final theory in this respect. But, more importantly, the extent to which the science in question is 'hostile' to fundamental particles is murky - so this is far from the hubris of brazenly contradicting what our best science is patently telling us. As such, in so far as we take disbelief in smallest particles to entail the falsity of nihilism, this consideration needs to be weighed against the independent reasons we have for believing nihilism in the first place.

Furthermore, while it is perhaps true that interpreters of QFT tend to favour views on which particles do not appear fundamental, we shouldn't pretend there is a consensus—e.g. Teller (1995) argues that Quantum Fields are most plausibly seen as *constraining* what is physically possible rather than *grounding* the physically contingent facts, in which case Schaffer's suggestion¹⁷ that we abandon particles in favour of 'treating worldwide fields as fundamental' (p.54) would arguably be untenable, absent some alternative way to ground the contingent physical facts.

Speaking more to the general worry about QFT undermining nihilism, it seems we could bring in resources employed by other strategies in this chapter to attempt to defuse the worry. For instance, if we consider the extended heterogeneous simples of the previous subsection, it's far from clear that 'mereological simple' need be synonymous with 'smallest particle'. Could the Quantum Fields themselves be mereological simples?

¹⁷ Following D. H. Zeh (2003).

Might the world really consist of spatially extended – perhaps overlapping – quantum systems that are mereologically simple yet possess the quantum properties of multi-particle systems?¹⁸ It's hard to evaluate these suggestions without getting mired in the highly technical and controversial interpretive issues we're trying to side-step, but Schaffer hasn't given us any reason to doubt their plausibility or compatibility with QFT. The same goes for other ways of recasting the nihilist's ontology: What if, for instance, we take the supersubstantialist approach and identify simples with points in spacetime possessing field-values?¹⁹

What's more, the resources afforded by fundamental metaphysical indeterminacy²⁰ (discussed in §2.4) would allow us to undercut the main reason Schaffer gives for considering particles non-fundamental in QFT: 'How could particles be fundamental if there is no fact about how many the system has?'(p54.) The resources of fundamental metaphysical indeterminacy would allow us to concede that there is no *determinate* fact of the matter about how many particles there are without this giving us reason to think that they therefore need to be banished from fundamental ontology.²¹

6.2.2. *Objection 2: Losing the unity of properties*

Schaffer argues that by endorsing entanglement relations we lose the 'unity' (p.54) of quantum properties. For instance, in the case of spin, we can't attribute the very same basic spin property to any arbitrary system of particles that has, say, spin $\frac{1}{2}$. Rather, single particle systems will have the basic spin property; while a two particle system will have its spin in virtue of the two basic spin properties of its components and an entanglement relation; and a three particle system in virtue of the three basic spin properties of its components and three entanglement relations (or a three-placed entanglement relation); and so on.²² Is this an objectionable loss of simplicity or theoretical unity?

My only comment is that for this to be a 'loss' of simplicity or unity, there would have to be a rival theory that managed to maintain the alleged unity of quantum properties. On

¹⁸ See Williams (2006).

¹⁹ See Sider (2013).

²⁰ e.g. those afforded by the theory explicated in Barnes & Williams (2012).

²¹ Count indeterminacy will plausibly entail indeterminate existence, and thus might be thought to fall foul of the Sider-Lewis argument discussed in chapter 7. See Barnes (2010) for a defence of metaphysically indeterminate existence against that argument.

²² This objection has some obvious affinity with the variable polyadicity objection to joint instantiation, though the two are entirely independent of each other (the same goes for my response).

standard theories of micro-determination it looks like you don't want to say that the fundamental physical quantities had by composite systems – whether fundamental or non-fundamental – are *the very same* as those had by their basic constituents. Consider a 4Kg composite *c* made up of four (extremely heavy!) 1Kg simples, and a 4Kg (even heavier!) simple *s* going about its business several miles from *c*. In the region occupied by *c* you really don't want to say that there are four instances of being 1Kg and, in addition, an instance of being 4Kg which is the very same property being instantiated by *s*. If this were the case, the matter in the region occupied by *s* would weigh 4Kg, whereas the matter in the region occupied by *c* would weigh 8Kg! The instances of mass in region *c* have got to add together like this, else if we took one of the simples making up *c* and put it on the scale with *s* (which instantiates the *very same* mass property as *c* don't forget), they wouldn't come out as weighing 5Kg in total. So, the way in which *c* is 4Kg can't be the exactly the way that *s* is 4Kg. It doesn't automatically follow that a composite object couldn't have genuine causal powers in virtue of weighing 4Kg, or even that its mass property isn't fundamental, it's just that its having mass is evidently not exactly the same thing as its basic parts having mass. To prevent needlessly multiplying fundamental properties and to avoid overdetermination worries, I suspect the best thing *is* to say that the object's basic parts are doing the causal work and that the object itself is 4Kg merely derivatively, but this is just a tidying-up move given that unity already has to be given up on. Where we suspect genuine emergence, we avoid making this tidying up move, but are probably already taking the non-unity assumption for granted: Even setting aside the above issues, the emergent property is a dependent property, which already sets it apart from the properties in the base.

Similarly, on Schaffer's Priority of the Whole it would surely be only 'the World' – the 'one big entangled system' – that instantiated a basic spin property. Perhaps sub-world objects have fundamental intrinsic spin properties (though this hardly looks attractive), but they can't have spin in exactly the way that the World has spin: sub-world objects have a dependent sort of spin which can't be straightforwardly aggregated in the same way as independent spin without running into the kind of trouble outlined above. You can have unity if you really want it: embrace the radical view that composition is identity²³ and say that the mass trope had by the 4Kg composite is *numerically identical* to the four 1Kg mass tropes had by its basic parts. But this solution is not going to have broad appeal and

²³ See Baxter (1988a, 1988b) for canonical formulation; compare Lewis (1991). See, e.g., Wallace (*draft*, 2011a, 2011b) for recent discussion and defence. Here I assume the falsity of composition as identity (CAI).

also collapses the distinction between composite systems and their parts – on which the objection relies – in the strongest way possible.

In any case, the nihilist is going to take it as a given that the properties of composite systems are going to be derivative since composites are themselves derivative— that’s just par for the course. So even if there is something to Schaffer’s objection when applied to pluralist strategies generally, it’s already a core commitment of the nihilist which she’s unlikely to acknowledge as a genuine cost: discriminating between simples and composites by banishing the latter to the realms of the derivative is what nihilism’s all about, and getting rid of composites from her fundamental ontology is sold as a virtue of the theory. If unity considerations count for anything, it is disunity in one’s *fundamental* ontology, but nihilism – with or without entanglement relations – is not committed to this. The only sense in which there is disunity when entanglement relations are brought in is that the “spin phenomena” – the observed collective behaviour of fundamental particles correlated with what spin properties they have – is now explained by two fundamental properties rather than one: spin *and* entanglement relations. But to point out this is to reiterate what the nihilist has already conceded at this point: intrinsic properties of simples and their spatial arrangements aren’t enough to explain their collective behaviour, which must instead be some function of their intrinsic properties, spatial relation *plus* an additional relation that holds between them. Sure, there is a cost to buying this extra relation, which must be factored into the final weighing of costs, but there is no additional worry raised by ‘unity’ considerations as far as I can see.

6.2.3. Objection 3: Priority Monism & the Impossibility of Submergence

Schaffer believes that the best solution to the problem of the emergent behaviour of quantum systems is to adopt Priority Monism: If the world is the only independent entity, then there is no mystery as to why the whole sometimes fails to supervene on its parts. The properties of sub-world parts are always determined by the properties of the whole world— never the other way around. This is compatible with mereological nihilism as understood here: We can hold that the world is not just the only independent entity, but also, fundamentally, the only entity (i.e. it is the only real entity) and thus that, fundamentally, it can have no proper parts (its proper parts are not real). Indeed, this seems to be straightforwardly entailed by Schaffer’s own metaontology.

He believes in relative degrees of fundamentality, contrary to what is assumed in this

chapter—that either an object exists in the fundamental sense (is real), or it doesn't.²⁴ But it will be good enough for our purposes to treat Schaffer's ascriptions of absolute fundamentality – of being more fundamental than any other object – as ascriptions of existence in the fundamental sense, in our framework. For Schaffer, every object that is not absolutely fundamental is ontologically dependent on at least one other object that is more fundamental than it. As such, Schaffer's Monism entails that the World is the only independent object, with all other objects – the world's parts – being ontologically dependent. The competing position, given Schaffer's metaontology, is that the mereological simples are absolutely fundamental and ontologically independent, with the macroscopic objects they compose being dependent. Thus, Schaffer takes for granted a tight connection between fundamentality and independence, just as most pluralists do.

Now for the objection:

Intuitively, just as emergence involves, in some sense, having wholes that are more than the sum of their parts, submergence would involve having wholes that are less than the sum of their parts. Schaffer thinks that Monism allows for emergence since he seems to assume that only the most fundamental thing can have fundamental properties, and therefore that the properties of sub-world objects will be derivative and determined by the fundamental properties of the world. As such, the whole world can have properties that are more interesting than its sub-world parts, but those parts cannot have properties that are more interesting than any wholes they compose.

The obvious thing to do to enable pluralists to account for emergence is to relax the constraint that only the simples can have fundamental properties – or to pursue one of the nihilist-friendly strategies discussed in this chapter. But is it a cost to pluralism that it is permissive of 'submergence', and an advantage of monism that it is able to account for seemingly emergent wholes without relaxing the oft-presumed link between fundamentality and dependence? To the first, I think the answer is a definite no: As I'll argue below, in so far as we have good reason to want our ontological theories to be able to deal with emergence, we should want them to deal with submergence, so if anything the ability of Schaffer's brand of monism to deal with submergence is a weakness of monism rather than a strength. To the second, if you're strongly attached to the idea that fundamentality and independence are (at the very least) coextensive across all worlds, you

²⁴ See, esp., §3.1.

might find the monist way of accounting for emergence tempting:²⁵ but the risk is that you'll soon find yourself having to account for putative cases of *submergence*, which may force you to abandon the link between fundamentality and independence after all and thereby undercut your primary reason for endorsing monism in the first place.

Before continuing it is important to note the terminological divergence between myself and Schaffer concerning the use of the term 'emergence'. When emergence is construed as a fundamental-dependent entity having a fundamental property, there is no intuitive distinction between emergence and submergence: It's true that most of the putative cases of emergence - and thus also the intuitive characterisation of the phenomenon - will be of wholes having properties that are something *over and above* the properties of their parts. But this is classifiable as emergence only when a pluralist ontology is assumed that makes simples basic and wholes derivative. When we make the World the only basic object, all the classic arguments for emergence are turned on their head:

Now we must ask whether there are any putative cases of *parts* having fundamental properties, despite being dependent entities. And, of course, the intuitive answer that presents itself given the prevalence of pluralism in our thinking is that there is a multitude of examples: A particle having a -1 charge isn't, intuitively, a derivative property had in virtue of, say, that particle being a part of a five-particle system that has the fundamental property "having a net charge of +3". Wherever it seems theoretically useful to suppose that wholes get their properties in virtue of the properties of their parts - and not the other way around - we have a putative case of 'submergence'. Sure, the priority monist has already bitten this bullet by asserting that all the properties of sub-world objects are derivative and fixed by the World's fundamental distributional properties— though usually only a very sketchy promissory-note of a story is given concerning how this happens.²⁶ So, if pluralists should worry about how simples could give rise to entangled quantum systems or consciousness, monists should worry about how the "World properties" give rise to apparently local or pseudo-local properties: and for sure they *do*, but they must be careful not to sell a lacuna of their view as an advantage. Priority monism can't account for situations where we may feel the best explanation of a particular phenomenon is to ascribe a fundamental property to a sub-world object - unless it loosens the link it assumes between fundamentality and dependence - and in so far as it is unable to do this, it is a potential cost to the theory.

²⁵ For a defence of this see §7.2.

²⁶ Though for some attempts to give more detail see Sider (2008); Cornell (MS.).

However, Schaffer prefers a supervenience characterisation of emergence/submergence, in which emergence is the failure of the properties of the whole to supervene on the properties of its parts and submergence the failure of the properties of the parts to supervene on the properties of the whole. I've already given reason to be suspicious of the utility of such characterisations,²⁷ but let's run with it for the time being in order to assess Schaffer's (2010, p.56) main argument for the impossibility of submergence:

An underlying mereological asymmetry comes to light: the asymmetry of supervenience. The asymmetry is that the proper parts must supervene on their whole [...] but the whole need not supervene on its proper parts. In other words, though emergence is metaphysically possible, submergence—the converse of emergence—is metaphysically impossible. For submergence, the intrinsic properties of the proper parts, along with the fundamental relations between these parts, must fail to supervene on the intrinsic properties of the whole. This is impossible because (i) any intrinsic property of the proper parts ipso facto correlates to an intrinsic property of the whole, namely, the property of having-a-part-with-such-and-such-intrinsic-property, and (ii) any relations between the parts also correlates with an intrinsic property of the whole, namely, the property of having-parts-thus-and-so-related. Fix the whole, and all of its parts are fixed.

The difficulty with cashing out this argument in terms of supervenience is that it glosses over crucial questions of dependence: If the properties of wholes are ontologically dependent on the properties of their parts, then you simply can't fix the mereological properties of the whole without *first* determining the properties of the parts. In which case it would make no sense to speak of fixing the properties of the parts *by* fixing the properties of the whole. The conceptual ordering is crucial here: God has free reign to make the parts in whatever image pleases her, and she thereby fixes the mereological properties of the whole. Even in her omnipotence, she has no *independent* means of settling the mereological properties of the whole which would thereby 'fix' the properties of the parts – unless she were to also make it that priority monism were true.

But what God can do, if she's willing and able to violate micro-macro supervenience, is independently fix the non-mereological properties of the whole in such a way that the whole lacks a certain type of property, X, that its parts all have, and to allow that the parts can vary their X properties (modally or temporally) without there being a corresponding change in the whole's X properties, or indeed any of its non-mereological properties. Yes, there will be a change in the whole's mereological properties relating to which of its parts

²⁷ See §5.1, §5.2.

have X-properties, but this is hardly of consequence given that the whole itself has no X properties. Here's an example of mine to illustrate just such a case:

Frank is the mereological fusion of me and Roland, my pet rabbit. Let us suppose Frank's physical and functional properties supervene on the physical and functional properties of Roland and I. Further, let's suppose that there is 'something it is like' to be me, and something (very different but just as vivid) it is like to be Roland the Rabbit— we both have an inner mental lives, filled with rich and irreducible experiential qualities. Alas, however, there is nothing it is like to be the discontinuous mereological fusion of me and my pet rabbit: So Frank has no inner mental life, no phenomenal properties. Despite being composed of two fully-fledged conscious individuals, Frank is nothing but a zombie rabbit-man.

Sure, there are many reasons you may have for doubting the possibility of this scenario, even if you're prepared to buy into qualia: You may think phenomenal properties supervene of necessity on physical or functional properties, or you may deny Frank's existence in virtue of his being a spatially discontinuous or radically heterogeneous object. My purpose in constructing the example is purely to illustrate what seems like a conceivable, coherent and plausibly motivated example of Schaffer's submergence. For the Frank case patently deserves the label 'submergence' label - Schaffer's own gloss notwithstanding - if the converse cases ever deserved the label 'emergence':

Here we have two parts with much more interesting and rich properties than the property of the whole they compose. And, crucially, these properties can plausibly vary independently of any of the physical, functional or mental properties of the whole: If Roland and I were suddenly subject to spectrum inversion that was not the result of any physical change in our brains (perhaps due to the meddling of a mad scientist's ghost!) - tomatoes now appearing to us the colour we used to see the sky - this need not correspond to any change in Frank's non-mereological properties. The fact that Frank has the trivial mereological property 'having a part that experiences red qualia when he looks at a tomato' seems utterly uninteresting: It is certainly not *in virtue* of him having this property that I experience red qualia when I look at a tomato, and there's certainly no remotely plausible way in which Frank could be said to be phenomenally conscious merely in virtue of having this property, given the assumption of this scenario that there is no tight modal connection between the phenomenal properties of parts and wholes.

So, fix all the interesting non-mereological properties of Frank - his mental, physical and functional properties - and you will fail to fix the properties of his parts, at least in one very important respect. Insist on fixing his mereological properties and you will indeed succeed in fixing all the properties of his parts, but that's because to fix the

relevant subset of his mereological properties *just is* to fix the properties of his parts.²⁸

Should we endorse priority monism to rule out these cases across the board? No, these cases are perfectly coherent and should be assessed on their own merits. Feel free to rule them out if you have *independent* reason to endorse the priority of the whole – but this could be seen as a cost to the view in just the same way as ruling out emergence is an alleged cost of strict priority pluralism. Schaffer writes ‘any intrinsic property of the proper parts ipso facto correlates to an intrinsic property of the whole, namely, the property of having-a-part-with-such-and-such-intrinsic-property’. But we see now that these ‘ipso facto correlates’ are an empty, irrelevant formality: Having phenomenal consciousness, for instance, is plausibly nothing like merely having a part that is conscious. And the same goes for physical quantities: *a*’s having a part that is charged, massive, spin-up, or whatever, has no physical significance unless there is some reliable modal correlation between the physical properties had by parts and the physical properties had by wholes. But it is exactly these sorts of correlations that putative cases of *emergence* – which Shaffer readily endorses – threaten to provide counterexamples to.

It would be ludicrous to change one’s theory to rule out emergence – or to declare the putative cases of emergence not genuine emergence – purely because whenever we find an example of a whole that is something “over and above” its parts, we note that there is an *ipso facto* correlate of that emergent property had by each of its parts, namely “being part of something that is thus and so”. In this case, the mereological properties ignored are not intrinsic, but my contention is that Schaffer’s appeal to intrinsicity here is a red-herring: We ignore them because they are uninteresting mereological properties – paralleled in the case of submergence (rather than because they are not intrinsic properties), which is irrelevant and gives an illusion of interesting emergence/submergence asymmetry. The properties of wholes, on the pluralist picture, are fixed by the intrinsic properties of their wholes *plus* the relations between them— so only considering the *intrinsic* properties of a whole’s supervenience base is just to disregard part of the picture without principled reason.

Just as with emergence, if we’re interested in the possibility of submergence at all, we can’t be interested in it in this super-strict sense, on which it is obviously false (i.e. whether there is supervenience failure between a property and its trivial mereological correlate): But rather in the possibility of more interesting, non-trivial violations, such as emergent

²⁸ Or to determine which of the things in existence – their non-mereological properties independently settled – he has as parts.

minds (on Team Emergence) or Zombie Rabbit-Men (on Team Submergence). It's a *prima facie* cost to rule out either possibility without independent reason. So, strict priority pluralism and strict priority monism both incur a cost for ruling out *interesting* supervenience failure in one direction: there is no corresponding gain to either theory for ruling out the super-strict and obviously impossible variety of supervenience failure in the same direction.

6.3. Discussion of Indeterminacy

This next strategy has its origins in Barnes (forthcoming-*b*) and is explained and defended in some detail there. As such, I'll try to outline the view concisely:

Let's take our two particles, *a* and *b*, familiar from our discussion of entanglement relations. Given how we set up the scenario, recall that there's some pressure to say - given the properties they will display on measurement - that "the system *a+b*" (be it composite or simple), or *a+b* jointly, instantiate the fundamental property of being F and G anti-correlated with respect to each other. The Indeterminacy Strategy avoids posting this extra ontology by wheeling in the ideological machinery of fundamental metaphysical indeterminacy.

The basic gloss is: Indeterminacy is *metaphysical* when it is a product of how the world is, fundamentally, rather than a product of our thoughts, epistemic state, or how we use our language (i.e. when the source of the indeterminacy is wordly, rather than linguistic or epistemic). It's also *fundamental* because the metaphysical indeterminacy being employed here can't be reduced to or explained in terms of more primitive metaphysical vocabulary: The indeterminacy operators employed are *ex hypothesi* unanalyzable and will be appear in the "fundamental language" in which we will state our final ontological theory. How does metaphysical indeterminacy help solve the puzzle? Barnes tells us that if we employ the ideology of fundamental indeterminacy, we can say that either *a* is F or *b* is F and its indeterminate which, but determinately only one thing is F— and we can say the same for G.²⁹ This is enough to ground the anti-correlation fact: no extra ontology is needed.

An interesting thing about this strategy is that it pays for the additional expressiveness the nihilist requires with the coin of ideology, rather than ontology. For some this might seem like a raw deal, but for others it might be attractive: The nihilist who has come to the

²⁹ Barnes (forthcoming-*b*).

theory via Sider's (*draft*) argument from ideological parsimony, having just discarded one ideological primitive in the form of 'parthood', would perhaps be unwilling to throw away this gain by endorsing a new ideological primitive in the form of a fundamental indeterminacy operator. However, a nihilist who is led via arguments from ontological parsimony might have no qualms in embracing a little extra ideology in order to have a more expressive theory. In any case, even the believer in composition needs to pay one way or the other for adding fundamental emergence to her theory: The point is, the nihilist has lots of ways to accommodate emergence and can pick the one that best fits with her broader ontological theory. Fundamental indeterminacy is one way to go which is pretty distinct from the other options on the table, and may be something you have reason to incorporate into your overall theory in any case - potentially giving you a completely non-*ad hoc* solution to emergence worries at no extra cost. If you've already got fundamental indeterminacy as a wider theoretical commitment, for instance to account for stochastic indeterminacy in quantum laws (to pick an example close to the current discussion) - or to account for the open future,³⁰ say, (to pick one further afield) -then you should have no qualms applying it here.

6.3.1. Objection 1: Against metaphysical indeterminacy

Worry— Isn't metaphysical indeterminacy mysterious? Crazy even?

Well, it's worth noting that these are the sorts of concerns that have been levelled against emergence, so - at the very worst - we're just trading one "crazy or mysterious" notion for another by explaining away putative emergence in terms of metaphysical indeterminacy. Nothing to lose!

Yes, but we have at least one well worked out theory of emergence, by your lights: the one you're employing in this chapter. Can the same be said of metaphysical indeterminacy?

Yes. In fact the theory of metaphysical indeterminacy I'll be assuming here is presently better worked out, in the sense that Barnes & Williams (2011) provide a detailed logic and semantics for it. The basic idea is that reality is represented by one or more abstract possible worlds that are semantically and metaphysically precise. If the world is completely metaphysically determinate, there will be only one world that doesn't determinately misrepresent reality (the 'determinately' here is primitive). When there is metaphysical indeterminacy, however, there will be multiple worlds that are candidates to

³⁰ See Barnes & Cameron (2009), (2011).

accurately represent reality: determinately, only one of those worlds gets it right, but it's indeterminate which world that is.

Fine. But I'm still suspicious of metaphysical indeterminacy!

Tough. You don't have to buy into fundamental indeterminacy, of course, but Barnes & Williams provide an intuitive characterisation of the notion and detailed working logic and semantics. Coupled with the case made for its theoretical utility in Barnes (forthcoming-b) and elsewhere, and the defence against arguments challenging the notion in Barnes (2010), it should be considered a serious option on the table. After all, what more could you ask for? As fundamental indeterminacy is being treated as metaphysically primitive an analysis of it is out of the question. And if it's just that you find the very idea of the world being indeterminate deeply unpalatable, there's likely nothing anyone can say that's going to relieve that uneasiness completely— but this doesn't constitute a principled reason for dismissing the view.

6.3.2. *Objection 2: The right tool for the job?*

The claim that *a* being *F* hinges on *b* being *G*, and that this connection is only wrought by bringing metaphysical indeterminacy into the world is, on the face of it, highly mysterious. If we're assuming that in a determinate world *a* can instantiate *F* entirely independently of *b*, and that *b* is similarly independent with regard to *F*, how is it that simply making it unsettled whether each particle instantiates *F* brings about a co-dependence in *a* and *b*'s *F* properties? If the scenario were simply described as one in which it's indeterminate whether *a* is *F* or *G* and indeterminate whether *b* is *F* or *G*, I think the reaction of most would be to assume that *a* and *b* both being *F* was a possible way things could get settled, and equally that *a* being *F* and *b* being *G* was also a possible way things could end up. It would not occur to us that the scenario had been under-described in some way, except perhaps we would like to know whether there is supposed to be any relation – any *ontological* connection – between *a* and *b*. On hearing that there was not, our resolve that both correlated and anti-correlated property distributions are ways things could get settled would only strengthen.

What does this show? If anything, only that intuitively we don't see indeterminacy as something that could compromise the autonomy of two *prima facie* independent entities. But to insist on this risks begging the question against the proponent of the indeterminacy strategy: Her thesis is precisely that fundamental indeterminacy offers us a way to ground instantiation correlations between separate individuals without the need to posit an ontological connection. If her opponent is going to insist that indeterminacy is not an appropriate tool for this task then he can't just rely on intuition. Rather a substantive

argument should be offered against using fundamental indeterminacy to ground such connections.

That's not to gainsay the intuition: Just because we introduce expressive resources into our theory that allows us to say x doesn't mean we have to hold that x is possible. Say we introduced expressive resources to allow us to state that the future is open, but these same resources allow us to state that the past is open. We can still have independent reasons for thinking that the past is closed and the future open, even if we buy into ideology that allows us to say that both are open. But since the expressive resources of our theory allow us to state that the past is open and opponent would be well within her rights to demand to know what the "independent reasons" are for ruling out the possibility— if they're not compelling, it's rationally open to her to hold the opposite.

What reason could we have for denying the possibility, even granting the availability of the requisite expressive resources? The following reason is suggestive: just as the Humean fears necessary connections between distinct existents, perhaps in the same spirit we should have a dictum against penumbral connections between ontologically independent entities— I pursue the consequences of this for the indeterminacy strategy further in §6.4. However, while I feel the dictum has the air of plausibility about it, I don't believe it constitutes a non-question begging reason to reject the indeterminacy strategy.

6.3.3. *Objection 3: Deep indeterminacy*

Skow (2010) objects to applying the Barnes-Williams model of fundamental indeterminacy to quantum mechanics, since quantum mechanics displays what Skow calls 'deep' metaphysical indeterminacy,³¹ which he believes the Barnes-Williams model is incapable of adequately modelling. As the primary example given in this chapter of how metaphysical indeterminacy might be used to explain away a putative case of emergence comes from quantum mechanics, this is potentially a big worry.

A theory exhibits deep metaphysical indeterminacy, on Skow's characterisation, if it says that when it's metaphysically indeterminate how the world is, all the ways it says the world might be resist precise representation (for worldly rather than semantic reasons). For instance, say an electron has a separate spin-value for each of three axes – x , y and z . For each spin direction, our theory tells us it's metaphysically indeterminate what spin

³¹ At least on orthodox interpretations— but Skow (2010) points out that on the sorts of non-orthodox interpretations, such as hidden-variable theories, which the Barnes-Williams approach can model straightforwardly, there may be less reason to posit metaphysical indeterminacy in the first place.

value our electron has. But suppose it also tells us that, determinately, the electron can never have a determinate spin value for x, y and z all at the same time. The Kochen-Specker Theorem tells us that examples of roughly this sort will abound in quantum mechanics.³²

How to model this? The worry is that, since the Barnes-Williams model works by specifying which perfectly precise worlds are candidates to be actualised, it will tell us that there are no physically possible worlds that are candidates to be the actual world— which is absurd.³³ This is because all the representational worlds in the Barnes-Williams semantics will assign determinate spin values to our electron for all three axes, since the representational worlds need to be free of any indeterminacy themselves if they are to successfully model the phenomenon in the way Barnes & Williams suggest. Yet, the physics tells us that it's physically impossible for an electron to have a determinate spin for all three axes. So, all the “candidates” to represent the actual world are physically impossible: but since actuality entails physical possibility, it's determinately not the case that any of the supposed candidate worlds could be actual.³⁴ Here Skow takes the Barnes-Williams model to break down.

I think the above description of the problem goes too fast, however. If it's physically impossible for the world to be fully determinate, then – yes – every ontic precisification of the state of the actual world will come out as physically impossible. But that's no bar to it being the case that one and only one of these worlds is actual *but that it's indeterminate which* – it's this indeterminacy, we may suppose, that *gives rise* to the phenomenon of quantum indeterminacy, given the actual laws. To clarify:

If we're going to distinguish between “deep” and “shallow” metaphysical indeterminacy, I think it helpful to distinguish between deep and shallow *physical possibility*. According to the shallow sense of physical possibility, the set of physically possible worlds is just a restriction on the set of metaphysically possible worlds in which the actual physical laws hold: In this shallow sense of “physically possible world” there is no incompatibility between a world being perfectly precise and it being subject to laws that rule out there being a perfectly precise world. All this means is that any such worlds, while representing themselves as instantiating the One True property distribution, will – of physical necessity – always “see” other worlds, with different property distributions, as

³² Skow (2010: 6-9)

³³ Skow (2010: 8)

³⁴ Skow (2010: 8)

competing candidates for actualisation. It's fine for there to be a candidate world that gives precise x, y and z spin values for our electron, provided our theory says that whenever it is a candidate there will always be at least one rival candidate that disagrees on at least one of these values.³⁵

There is, intelligibly, a “deep” sense of physical impossibility, under which it is indeed true that none of the candidate worlds are physically possible: the laws rule out any of them being *determinately* the One True description of reality – they could never be unrivalled candidates for actualisation. But that's just to be expected when you have laws which state the world will always be metaphysically indeterminate! You only run into trouble if your theory says there are no physically possible worlds in the shallow sense: this would mean the actual world was unrepresentable in the model – that no ontically precisified worlds are even *candidate* representations of the actual world. Here the Barnes-Williams model would ostensibly break down, but there's no reason to suppose such degenerate cases are genuine possibilities. That is, we shouldn't see “physically possible” as a restriction on the set of metaphysically possible worlds. Instead, they are sets of *sets* of not determinately incorrect worlds – in the terms of Barnes & Williams, haloes. There are no physically possible *precise* worlds: not because the theory has broken down, but because it is modelling deep indeterminacy in the actual laws! So-called physically possible “worlds” will in fact be sets of metaphysically possible worlds on this view, with each member being a co-candidate for actualisation.

Some objections:

“The model of indeterminacy relies on it being the case that determinately there is only one world that represents actuality – it's just indeterminate which world that is. But Quantum Mechanics rules out there being one world that determinately represents reality!”

This is just a confusion over the scope of the determinacy operator. I've argued that – absent further evidence – we should say that quantum mechanics is compatible with it being determinate that there is only one precise way to represent the world. Quantum mechanics has only been shown to be incompatible with there being – even after collapse of the superposition – one precise representational world which is itself determinately the One True way to represent reality. This is no bar to these worlds representing *themselves*

³⁵ It becoming settled (e.g. on observation / collapse) that our electron is x-spin up is just to move to restrict our group of candidate worlds such that all represent the electrons x-spin as ‘up’, but disagree about its y-spin and z-spin.

as being the One True representations – that’s just part and parcel what it is to be a world on this picture. In other words, QM tells us that there is no precise world such that it is determinately actual, *not* that determinately, there’s no precise actual world.

Another complaint that might be levelled against this way of incorporating quantum indeterminacy into the Barnes-Williams framework is that none of the candidate actual worlds can be coherently represented in quantum mechanical state space:

In addition, then, to the fact that quantum mechanics does not assign a determinate value to every one of the observables, the Kochen-Specker theorem shows that such a thing cannot be done – it is prevented simply by the vector space structure of observables.³⁶

This is only a problem if you take the state space to be metaphysically real – if you think that *what it is* to instantiate a certain set of quantum properties is to occupy a particular position in state space. In that case, it would not only make no physical sense to speak of a pattern of co-instantiation of quantum properties un-representable in the state space, but no metaphysical or ontological sense either. But the primary function of the state space is to represent the physically possible distributions of observable quantum properties: So in a world of “deep indeterminacy” the state space wouldn’t be doing its job if there were coherent representations in the state space of physically impossible (i.e. maximally ontically precisified) property distributions. We can take this on board without being barred from saying that tropes or universals corresponding to the observed properties such as spin, position, etc., are fundamental, and that the state space is just a useful mathematical representation of the possible patterns of determinate co-instantiation among these properties (combinations of properties un-representable in the statespace will never be determinately co-instantiated).

I can see two potential pitfalls in refusing to acknowledge the state space as real: It could be seen as discarding an explanation as to *why* some patterns of determinate co-instantiation are permissible and others are not. But unless we have some deep and satisfying explanation of why the state space is the way it is, it seems we have just shifted the target explanation from the co-instantiation patterns to the state space itself without gaining anything (perhaps we do have such an explanation available – but I leave this to the judgement of those more knowledgeable than I on the foundations of Quantum Mechanics).

The second is that this whole strategy represents a retreat from scientific realism: If

³⁶ Darby (2010).

quantum mechanics can make no physical sense of *any* of my candidate metaphysical descriptions of it, then in what sense am I taking any cue from science in constructing my ontology? Well, the list of observable properties I describe as fundamental will be taken from our best science, as well as much else, and I'm not second-guessing any of its empirical claims. Rather, it's my contention that in a world of deep indeterminacy it's no surprise that our best theories mandate that all physically possible states of the world are immune to fully precisified representation - we should not take this to *undermine* the Barnes-Williams model of indeterminacy, but rather see the Barnes-Williams model to be providing an attractive metaphysical analysis or representation of these indeterminate "physically possible states" in terms of there being multiple precise worlds being indeterminately realised.

I think the case remains to be made that the Barnes-Williams model of indeterminacy is inadequate for capturing the "deep" metaphysical indeterminacy of quantum mechanics. Again, there may be a cost incurred here to be considered in the final reckoning, but no knockdown objection.

6.4. A Quick Comparison of the Four Strategies

It's not clear to me - and in any case beyond the scope of our investigation - whether all cases of putative emergence in QM can be explained away simply by providing ontological or ideological underpinnings for the correlations observed in simple cases of quantum entanglement. Notwithstanding Schaffer's worries above, I do have residual doubts about the ability of entanglement relations to *fully* account for the phenomena of quantum entanglement, since quantum entanglement might be thought to have *two* distinctive features:

a) that restriction on certain collective qualitative possibilities for entangled particles, for instance that one particle can no longer be spin-up unless the other is in a spin-down state - i.e. loss of independence,

b) the fact that components of a given entangled system can't be assigned an individual qualitative state— only the system as a whole (e.g. we can say that the system is in a state as if having one particle spin-up and one particle spin down, but we can't say of either particle whether it is spin-up or spin down).

In a nutshell, the worry is that the nihilist can employ entanglement relations to explain the first feature of entanglement, but not the second: Given feature *b*, the properties of the components cannot ground the properties of the entangled whole, even if entanglement relations explain the entanglement correlations— because the individual

components have no qualitative properties themselves. One response might be to claim that, in an entangled state, the relations just *are* the qualitative properties of the parts, rather than any intrinsic monadic properties— at this point, though, I'd recommend abandoning the relational strategy and go all out for plurally instantiated (monadic) qualitative properties.

Another potential solution is to *combine* the entanglement relations strategy with Barnes's indeterminacy strategy. I take the indeterminacy strategy to be very well equipped to deal with feature *b*, since it can say that, while QM tells us that the particles in an entangled sub-system don't have a given set of monadic qualitative properties *determinately*, they do have some none the less, it's just *indeterminate* which they have. The indeterminately instantiated properties of the parts can ground the (probably also indeterminate) properties of the whole. Meanwhile, the indeterminacy strategy can technically solve *a*, but in a way which (as highlighted in §6.3.2) is going to make some queasy: That *merely* by adding primitive unsettledness to the world, God can suddenly make it the case that there are penumbral connections between distinct, otherwise independent existents, may be too much for some to countenance: The Barnes-Williams framework *allows* us to say that the world is unsettled between these two possibilities, but perhaps our Humean-style sentiments will push us to independently restrict which worlds can be can occupy the halo together as co-candidates for actualisation. Hence, one might employ Barnes-Williams to account for the putative indeterminacy in the qualitative properties of the parts, but insist on entanglement relations as ontological grounds for the existence of the penumbral connections.³⁷

Beyond adequacy worries in the quantum case, the potential to employ emergent properties in solving other problems — e.g. consciousness or as-yet unthought-of applications - promises greater theoretical utility. For these reasons, as things stand, I prefer the non-reductive strategies of plural instantiation and emergent simples to the more reductive strategies offered by entanglement relations and metaphysical indeterminacy. Of the two, I'm more inclined towards the plural instantiation strategy as it is less obviously reliant on additional resources - such as fundamental distributional properties - and the resultant theory is more in the spirit of traditional pluralistic nihilism by not providing 1:1 surrogates for composite bearers. However, as outlined above, I believe the proponent of the emergent simples strategy has promising ways to respond to

³⁷ This isn't any more theoretically expensive than employing entanglement relations alone if, as argued, we probably already have good reason to buy into the machinery of fundamental indeterminacy in wider theory.

each of these worries: So, when all is said and done, the two strategies may end up on a par.

However, if they can be shown to be adequate, the entanglement relations and indeterminacy strategies have the potential to offer powerful solutions to the problem of putative emergence in the relevant quantum mechanical cases. While the former explains quantum correlations by appeal to ontology, the latter pursues much the same approach by appeal to additional ideology. On first analysis the entanglement relations strategy may look more attractive, since it just posits a new type of fundamental relation (to go along with spatiotemporal relations) rather than an entirely novel ideological primitive. An additional external relation comes at a more minimal cost than the believer's need to appeal to both emergent properties *and* the ideological and ontological complication of belief in parthood in the first place. Especially if you come to the party impressed by Sider's (2013) argument for nihilism via ideological parsimony, the choice might seem clear cut. But even ignoring how difficult it is to sensibly weigh ontological and ideological parsimony off against each other directly, wider theoretical commitments may turn the tables: It's very plausible, as just argued, that one might be committed to fundamental indeterminacy to account for other phenomena, since the apparatus has such wide application. In this case, it can be wheeled in to provide a nihilist account of emergence in quantum mechanics for free.

In the final analysis, the more ways the would-be-nihilist has for dealing with emergent properties, the better dialectical position she finds herself in: as such, while I've indicated my own preference for plural instantiation, there's no need to pick a single winner. I believe each constitutes - or at least has the potential to facilitate - an adequate response to the No Bearers argument against nihilism, without thereby undercutting the argument for nihilism from simplicity.

7. Priority and Pluralism

Are wholes ontologically dependent on their parts, or are parts dependent on the wholes they compose— and should we expect the same answer to hold for all actual parts and wholes, let alone necessarily (as Schaffer 2010 assumes)? Does positing wholes that are prior to their parts help us account for putative emergent properties of wholes, e.g. from quantum mechanics? And why, as nihilists, should we care about this debate anyway?

The starting assumption of this chapter is that we reject the understanding of emergence, defended in the previous two chapters, as dependent novelty. After briefly motivating such a rejection, I put forward a tentative *non-uniform* answer to the priority question: Sometimes parts are prior to the wholes they compose, while sometimes wholes are prior to their parts (call this ‘Local Holism’). Specifically, I argue that the view that the whole is prior to its parts iff the parts form an entangled system provides a more attractive ontology than either priority monism or pluralism, given putative cases of ontological emergence from quantum mechanics. Most importantly (for our purposes), this serves to motivate a more parsimonious *nihilist* reconstruction of Local Holism - one that appeals to the notion of *merging*. The view is then defended from an analogue of the Sider-Lewis argument against restricted composition.

§7.1 introduces the notions of ontological dependence and priority, drawing attention to some important background assumptions of the debate that follows— most significantly, that ontological independence and reality mutually entail one another, thus that dependent things are never absolutely fundamental (in keeping with what I take to be the assumption of much previous debate on this issue, but contra the discussion in the last chapter). §7.2 Gives some reasons for thinking these background assumptions attractive, notwithstanding the case for fundamental dependence made earlier. §7.3 examines putative advantages of taking wholes to always be dependent on their parts - focusing on an argument from Ted Sider which seeks to demonstrate that pluralists can enjoy the theoretical benefits of employing combinatorial possibility while the monist (who takes the whole always to be prior and composition to be unrestricted) cannot. I conclude that although the monist has a convincing way to respond to Sider’s original argument, the pluralist accrues certain advantages of theoretical simplicity. However, in §7.4 I argue that the aforementioned simplicity advantages only accrue to the pluralist given certain assumptions about the basic structure of the world that are in *prima facie* conflict with

uniquely quantum phenomena such as entanglement— reminiscent of Schaffer’s (2010) argument that pluralism is ill-equipped to deal with quantum entanglement. §7.5 is concerned with whether monism or pluralism has the best hope of giving a metaphysics for an entangled world – or, at least, an entangled toy Screen World. The result is that, provided the lessons from the simplified Screen Worlds carry over, *neither* monism nor pluralism may be particularly well suited to accounting for quantum entanglement. However, a tentative case can be made that a promising quantum entanglement friendly ontology is one on which sometimes the parts depend on the whole, while sometimes wholes depend on their parts. That is, a sort of localized ontological holism (LH) is outlined.

With §7.6 we move to a discussion of mereological nihilism: It is argued that mereological nihilism fits best as a package with the metaontological assumptions outlined in the first subsection and employed in the paper and goes on to develop in more detail a nihilist reconstrual of LH by appeal to the notion of ‘merging’. In §7.7 I defend this version of LH against a worry that it has a bloated ideology compared to monism or traditional nihilism. §7.8 defends the view - *qua* “restricted” (i.e. non-uniform) answer to a question about merging analogous to SCQ – against an analogue of the Sider-Lewis argument against restricted composition, arguing against the no sharp cut-offs and no worldly vagueness assumptions respectively.

7.1. Priority, Parthood and our Intuitions.

Is the whole prior to the parts, or are the parts prior to the whole? Schaffer (2010) posed this question, defending an answer that had previously fallen out of favour in analytic metaphysics – that the whole is prior to the parts. He gives us the example of a circle and its component semi-circles:




Figure viii. A circle divided in half by a dotted line.

Intuitively, Schaffer claims, the semi-circles are arbitrary abstractions from the circle. We might say it is the circle – the whole – that is fundamental, while the semi-circles are

merely derivative entities, dependent on the circle for their existence. However, examples abound that are considerably less helpful to Schaffer's case. For instance take this example, adapted from David Lewis,¹ of a digitally produced picture. Specifically, this pixel art sprite of a Star Fighter for a smart phone app:



Figure ix. Star Fighter sprite – 148 x155 pixels*

It seems clear that the properties of the Star Fighter – e.g. being sort of pointy – are dependent on the properties of the 22,940 individual pixels that compose the sprite. That is, the visual properties of the image as a whole are ultimately dependent on the intrinsic properties and arrangement of its pixel parts.

There's a raft of cases that lend support – to a greater or lesser degree – to one side or the other of the Priority debate. There's some *prima facie* intuitive force to the vitalist thought that the parts of living organisms are dependent on the organism as a whole: if you're impressed by the neo-vitalist arguments of van Inwagen (1990), you might find the idea that a living thing is prior to its organs, tissues and cells both plausible and attractive. On the other hand, if you build a Lego house out of Lego bricks, this seems much closer to the case of the pixels and the pixel art. Similarly, if you've got a pool of water, or a lump of clay, I think that someone not entirely in the grip of an atomistic picture of the world would find it intuitive that the pool and the lump are prior to any arbitrary division we might make in the water or the clay. Yet, on the other hand, the planks are intuitively prior to the sailing ship, the buildings to the city, the sand grains to the beach, and so forth.

Trading intuitions about everyday cases won't get us very far. However, I agree with Schaffer that consideration of such cases is useful for one thing in particular: it demonstrates that the notions of whole being prior to parts and of parts being prior to whole are both fairly familiar and easy to grasp by reference, or analogy, to everyday cases

¹ Lewis (1986*b*, p.14).

of composition— even if the technical explications of priority are somewhat less transparent. Neither the idea that the parts are prior to the whole, nor that whole is prior to its parts is straight out crazy, or inexcusably esoteric. What’s also notable, though, is what a mixed impression a survey of the intuitive, everyday cases gives: neither the pluralist orthodoxy nor Schaffer’s Priority Monism are supported. Unless in the grip of a theory that gives us a blanket answer, a non-uniform, case-by-case approach to such questions is most natural, with the answer dependent on the details of said cases:

Q: “Is the whole prior to the parts, or the parts prior to the whole?”

A: “Well, it depends on...”

The intuitive cases not only make conceptual space for both uniform answers to the priority question, and for non-uniform answers. There may be good reasons for preferring a uniform answer – we’ll consider this in due course – but there’s nothing on the face of it that’s confused, crazy or esoteric about giving a non-uniform answer. In fact it’s the type of answer most strongly suggested by our ordinary intuitions. The parallel with the Special Composition Question is easy to spot:

(SCQ) Under what conditions do some things compose some further thing?

An ontologist might have certain theoretical motivations for giving an extremal answer – that composition never occurs, or always occurs – but the intuitive, pre-theoretic thought is that it’s a non-trivial matter whether or not composition occurs in a given circumstance. So, a non-uniform answer to the question:

(SPQ) Under what conditions are the parts prior to the whole?

will be of the form:

(LH) Under [non-trivial conditions X] the whole is prior to the parts, while under [non-trivial conditions Y] the parts are prior to the whole.

“LH” here stands for “(Mereological) Local Holism”, to be distinguished from “Mereological Holism”:

(MH) The whole is always prior to its parts.

Which equates to Priority Monism given certain assumptions outlined in Schaffer (2010) – a notable assumption being that composition is unrestricted, so there is one object, “the World” that has everything else as parts. Compare this with the orthodoxy, call it “Mereological Localism”:

(ML): The parts are always prior to the whole.

which leads to atomism, on the assumption of unrestricted decomposition.

The purpose of this chapter is not argue for LH as such. Instead, I hope to persuade you, firstly, that neither ML or MH are particularly promising theories when it comes to accommodating the lessons of fundamental physics into our metaphysics – despite arguments from physics for both ML and for MH being offered in the literature, from Schaffer (2010) and Sider (2007) respectively. Secondly, I tentatively put forward a version of LH which might do somewhat better given the considerations advanced – my intention is to show that making room for locally holistic theories opens up the potential for a more empirically and theoretically satisfying answer to SPQ. My version of LH will not make any concerted attempt to respect our ordinary intuitions about cases— in this respect it won't be better off than either of the uniform alternatives. As I've said, I take our non-uniform intuitions about priority to motivate allowing conceptual space for LH, but beyond this I don't think our ordinary intuitions count for very much when doing metaphysics. How can a view about the priority of part and whole bear in any interesting, non-trivial, way on nihilism though – a metaphysic in which nothing is ever a part or a whole? Well, because, the discussion and defence of LH serves to motivate a nihilist reconstrual of LH that can both (a) handle putative cases of emergence, and (b) do so without making the controversial move of the last chapter of divorcing dependence and fundamentality.

7.2. Dependence and Fundamentality

Quantum entanglement is often put forward as a putative example of ontological emergence: a novel property *emerges* at some point in the mereological hierarchy, such that it fails to supervene on the properties at the smallest level. However, given the assumption that a thing is independent iff it is fundamental – an assumption weakened in the last chapter following Barnes (2012), but reinstated here – it's infelicitous to count anything as a genuine example of ontological emergence. As is well rehearsed,² given orthodox pluralism the only fundamental properties are the ones at the lowest level of the mereological hierarchy, and these determine all properties of the derivative things higher up the hierarchy: so there is simply no space in this schema for things – properties or individuals – to *emerge* higher up the hierarchy, in either the sense of being a genuine addition to reality (since nothing higher up the hierarchy can be fundamental) or in the sense of not being determined by the base (since everything higher up is dependent on

² E.g. Kim (1999).

the base). This is why orthodox pluralism is generally thought to be ill-equipped to deal with genuine emergence: though this is most often stated as the view that, since the orthodox framework must be assumed, emergence is impossible or even unintelligible. Even given monism or LH, where the usual priority relations are reversed (or locally reversed), it is the properties of the smallest things that are taken to be determined by some larger thing, higher up the mereological hierarchy, and the smallest things are not taken to be fundamental. So, it's not as if there's any sense in which any properties emerge "from" the bottom level— all properties reside at the higher level, or are determined by and derivative from such properties. For genuine emergence, it seems we must have a hierarchy of some description *between* fundamental existents as discussed in Chapter 5.

In Chapter 5 I discussed the idea that fundamentality and independence do come apart, with emergent entities being cast a fundamental yet dependent. Emergent entities are "higher-up" in the dependence hierarchy, by being dependent on the microphysical base, yet genuinely novel in the sense that they are fundamental additions to reality. But despite the potential expressive power unlocked by separating out dependence and non-fundamentality – such that they become conceptually distinguished and not co-extensive – I do take there to be motivation to avoid this if possible. This reason is not at all decisive, in my view – I'd happily give up this metaontological position if doing so saw a way to a much more attractive first order theory – but it should none the less be given due consideration. To put the argument simply, just as the notion of ontological dependence provides an explanation for some cases of supervenience, so too can it explain fundamentality. On this account, we can *explain* why an entity is fundamental as its being ontologically independent. As such, we do away with the need for *two* metaontological notions of dependence and fundamentality. Moreover as we'll see, views that *can* keep the two together have been developed that – their proponents claim – can handle the sort of 'emergent' phenomenon that our nihilist in the last chapter split fundamentality and dependence to explain. If the nihilist, unlike the believer in composites, is *forced* to separate fundamentality and dependence in order to do that explanatory work, it might be thought that the nihilist is thereby at a disadvantage. But, by developing and refining the non-nihilist view in this chapter, we'll get to see that there's a simple enough nihilist analogue which *can* handle putative cases of emergence while keeping dependence and fundamentality intertwined.

7.3. Monism vs. Pluralism

According to the orthodoxy of pluralistic atomism, the *fundamental constituents* of the

world are tiny, perhaps point-sized, mereological atoms. The extensionless simples are the building blocks out of which everything else is constructed: The existence and properties of composite objects depend ('supervene' – but with the implication of asymmetric dependence) on the location and intrinsic properties of the simples.³ This model is attractive because it allows much “higher-level” complexity – cats, chemistry, cabinets and the cosmos – to be explained in terms of a handful of properties of the atoms (the “micro-level”).

Why abandon this attractive picture? The principal concern with atomistic orthodoxy is, I think, the empirical possibility⁴ of ontological emergence: that is, “higher-level”, composite systems with novel properties— properties that are “genuinely new” in some sense, and not merely derivative from the properties of their microphysical parts. It's hard (I suspect impossible) to make sense of this without abandoning at least one aspect of atomistic orthodoxy: either that mereological simples are not the only (absolutely) *fundamental* things, or that not everything depends on the microphysical simples. As I've said, Barnes (2012) gives us one way of salvaging some aspects of atomistic pluralism, if we're prepared to abandon the link between dependence and fundamentality. But what if we are not prepared to make such a radical break with orthodox metaontology? It is at this point that Schaffer's Priority Monism begins to look promising. By adopting MH, the supposed difficulty with accommodating “higher-level” composite systems that are independent of, or more fundamental than, their parts is instantly dissolved. On MH ontologies, such as Monism, any given composite system is *always* more fundamental than its parts, including its mereologically simple parts, and these parts are always dependent on them. Given MH plus unrestricted composition it follows that the only absolutely fundamental object is the whole world – the maximal mereological sum – and that everything else is dependent on the World. This means that for any plurality of non-overlapping objects which appear to collectively exhibit emergent properties, there will always be some further thing they compose to which the emergent property can be attributed.

Aside from being somewhat counterintuitive, is there anything unattractive about the monist's “catch all” solution to the problem of emergence? The main issue, as I see it, is

³ Recall from chapter 3 that on my metaontology, the mere non-fundamentality of composites doesn't render them unsuitable members of the domain of a perspicuous 'ontologese' unrestricted quantifier, and thus their non-fundamentality doesn't *guarantee* their unreality.

⁴ By the empirical possibility of P, I mean its epistemically open that we could come to know by empirical means that P being the case is at least physically possible, if not actual.

that a certain amount of explanatory simplicity is lost by treating everything as one interdependent system, rather than as independent subsystems. This, I would venture, is what lies at the heart of the debate between Sider and Schaffer / Cornell.⁵ In a nutshell, monism may be “overkill” when it comes to accommodating putative emergence in our ontology; its successes coming at an unnecessarily high theoretical cost. I should point out from the start that nothing I say will threaten the coherence or adequacy of Priority Monism or MH as a theory— but I do hope to show that there would be some unattractive consequences of endorsing monism, given some epistemically possible ways the world could be; specifically, given certain ways fundamental physics suggests the world to be. However, the road to this conclusion turns out to be a rather twisty-turny one, as we’ll see.

To the issue of simplicity, then. Take Sider’s (2007) “Screen World” consisting of a 4x4 grid of pixels that can either be on or off, e.g.:




Figure x. Three example Screen World states.

Sider points out that there are 65,536 possible ways for the Screen World to be, given that the 16 pixels vary independently with respect to their instantiating 2 possible qualitative states (on or off) = 2^{16} combinations of on/off pixels. In Sider’s terminology, the Screen World has a *state space size* of 2^{16} . One of the arguments he advances against the monist is that the atomistic pluralist is in a better position to account for this fact than the monist. The pluralist says that Screen World is composed of 16 ontologically independent entities, pixels, which can instantiate one of two determinate properties, on-ness or off-ness: This means we can read off the stipulated qualitative possibilities for the screen world straight from the posited pluralist ontology, in the manner detailed above. By contrast, since the monist thinks the only ontologically independent entity is the whole Screen World, and thus the only fundamental properties of the Screen World belong to the Screen World as a whole, she must say that the fact there are exactly 2^{16} ways for the Screen World to be is just a brute fact. Further, the monist cannot account for natural

⁵ Sider (2007), (2008) and Schaffer (2010) / Cornell (2013).

groupings between the members of the statespace – such as the fact that all the Screen World states where exactly 3 pixels are on “go together” – since the monist can’t appeal to the qualitative states of subworld objects.⁶

David Cornell⁷ replies that that the monist can account for the different combinations in an explanatorily deep way – or give just as satisfying an explanation for the size of the statespace as the pluralist, at any rate. No more is needed than there being precisely 65, 536 fundamental qualitative properties that the entire World can instantiate, which can be seen to be structurally and explanatorily on a par with the pluralist’s explanation: The pluralist has an ontology of two fundamental properties and 16 fundamental objects – so 2 to the power of 16 combinations. By contrast, the monist has an ontology of 65, 536 fundamental properties and one fundamental object – so, 65, 536 to the power of 1 combinations. Both views can be seen to use the same combinatorial procedure, and both take the number of fundamental properties and the number of objects to be brute. Further, the monist’s fundamental qualitative properties are *distributional properties*: for each determinate fundamental quality (i.e., ‘on’ or ‘off’) there is a way that that quality can be distributed (a higher-order property) – e.g., only in the top-left corner, or uniformly across the whole screen world, and so forth. As Cornell argues, this enables the monist to adequately account for natural groupings in the statespace, since we can appeal to distributional facts such as the ratio of on-ness to off-ness being instantiated in a given state.

This is all very well. Monism looks like it can provide *adequate* explanation for the Screen World, along the same lines as the pluralist. However, I take a key feature of Sider’s pluralistic version of the Screen World to be its *simplicity*: it manages to generate the entire space of 2^{16} possibilities from just 16 pixels and 2 fundamental qualities. By contrast, the monist requires 2 fundamental qualities and several thousand higher-order distributional properties⁸, in addition to its solitary object – whichever way you cut, this is

⁶ See Sider (2007), However, Sider (2008) offers a response to this argument on behalf of the monist which proceeds by taking the state-space itself to be fundamental. I leave aside this approach in favour of Cornell’s.

⁷ Cornell (2013).

⁸ Here I’m assuming that the properties to be distributed are distinct from the *pattern* of distribution, which I’m taking to be properties of the distributable qualities— e.g. on-ness and off-ness are both distributed in a chequered way. An alternative would be to take each possible distribution of qualities as an irreducible first-order property: such that red-all-over, black-all-over, white-all-over, black-striped, white-striped and red-striped would be 6 irreducibly-distributed qualitative properties, rather than three basic qualities (red, black and white) and two irreducible distributions (all-over and striped). Compare

extravagant theorising, requiring either a grossly expanded ideology, or an ontology containing a vast number of different universals. Each higher-order distribution is (in a higher-order sort of way) qualitatively different from the alternative distributions. The power of specifying a large number of possibilities from a small number of fundamental postulates is greatly diminished for Cornell's Monism. Yes, the structure of the explanation is there, but it is a limiting case of combinatorial explanation in which combinatorial simplicity gains are non-existent. And I believe this combinatorial power is one of the main attractions for being pluralist.

The above advantage of pluralism comes to little, however, if the actual world is importantly disanalogous to Screen World. Sider's physical possibilities for the Screen World are predicated on nakedly pluralistic assumptions: the ontological independence, thus qualitative independence, of the pixels. If the actual world is not even loosely analogous to the screen world in its metaphysical structure (specifically, in its state space), then the monist need not worry about accounting for worlds with Screen World like possibilities. Or, at least, the worst that could be said about monism is that it fails to account - or fails to account very simply or elegantly - for a certain conceptual possibility; that of the world being qualitatively "screeny". (The dialectical situation would be much like that between the nihilists and proponents of mereological gunk: There would be the question of whether conceptual possibility in this case leads to metaphysical possibility, or whether the seeming conceptual possibility could be explained away,⁹ and whether alien metaphysical possibilities need accounting for at all.¹⁰) Provided the *physical* state space of the actual world is not at all analogous to Screen World, however, there's certainly no direct line to be drawn from Sider's thought experiment to the falsity of monism at the actual world.¹¹

This is precisely where the problem of emergence bites: Suppose there are properties of the screen world that cannot be specified in terms of the individual pixels being on or off, or even the total pattern of on/off pixels, but only by considering the screen world as a

Cornell's (MS) discussion of simple vs. complex distributional properties, which includes an assessment of the costs and benefits of each approach.

⁹ E.g. by the method employed in Williams (2006).

¹⁰ See also extensive footnote on gunk in §1.1.

¹¹ Moreover, bolstering her argument from the conceptual possibility of screen worlds might turn out to be a much less pressing concern for the pluralist if it turns out the actual world is not at all screeny, compared to the task of showing that *pluralism can adequately account for the physical state space of the actual world*.

whole— the ineffable qualitative feel of being a Screen World, say. In this case, *ex hypothesi* the pluralist could not generate *all* the combinatorial possibilities from the screen’s smallest elements— straightforward pluralism must be abandoned. So, the question is, how “screen-like” is the space of physical possibilities for the actual world. If there aren’t really any properties of the world that can be specified by considering only the local properties of subworld objects and their arrangements, then pluralism looks straightforwardly doomed. The assumption that we *need* to account for screen-world-like possibilities thus seems like a substantive and possibly question-begging assumption. If the actual world is completely atomistic or screen-like in its space of physical possibilities, then monism is a rather extravagant ontology for accounting for such a world. However, if the world contains a mix of local and holistic properties, or properties for which it’s not clear cut whether they should be considered holistic or local, then the dialectic becomes harder to evaluate.¹²

Worse, due to what is essentially the familiar problem of emergence, using the assumptions we set out with, a ‘mix’ of local and holistic properties is impossible: On the assumptions of ML, fundamental properties are possessed only by the smallest objects in the mereological hierarchy, so the only way a proponent of ML could posit such properties is by denying the existence of anything smaller than the bearers of the holistic property. But this would leave the proponent of ML, i.e. putative pluralism, with the same style of explanation as the proponent of MH (i.e. monism, given unrestricted composition): There is just one thing – the world – instantiating fundamental properties. So, you can’t have a ‘mix’ of lower-level and ‘higher-level’ fundamental properties on the traditional metaontological framework, and by endorsing fundamental properties of the world as a whole ML collapses into monism.¹³

However, maybe all this shouldn’t worry the proponent of ML overly much: As familiar as the problem of emergence itself are attempts to *explain away* apparent instances – or the apparent possibility – of holistic properties. As we saw in §6.2, one way that pluralists have hoped to explain apparently holistic correlations in systems is by positing fundamental relations between the simple parts of that system. This strategy will not work for properties that are supposed to be *ex hypothesi* holistic – such as the qualia

¹² E.g. Sider (2007) argues that charge is one property that should be treated non-holistically even in quantum mechanics. Mass might be another.

¹³ While the monist finds her view in danger of collapsing into pluralism if she finds herself needing to accommodate ‘submergent’ properties (see §6.2.3).

example above - but one might be sceptical that such properties are even possible, let alone likely to exist in the actual world. Much will depend on the specific claims being made on behalf of holistic properties.

The main argument that Schaffer (2010) advances in favour of monism is from quantum mechanics, specifically quantum entanglement: The gist of which is precisely that the world cannot be neatly separated into independent qualitative units. How can we have a world of independent 'pixels' when the physics tells us that the only physically meaningful properties a given particle has in some situations are the properties belonging to an entangled multi-particle system of which it is a part? We'll consider Schaffer's argument in more detail later, and look in particular at the issue of whether it entails that the actual world is sufficiently disanalogous from Screen World to cause problems for the Sider argument. On the face of it, though, the quantum case Schaffer relies on is plausibly one in which a relational strategy could be employed - indeed Schaffer himself raises the possibility of positing 'entanglement relations' as a way of explaining entanglement in a pluralistic ontology,¹⁴ despite outlining some reservations.¹⁵

The line of argument pursued here will be that where the relational strategy is found to be adequate, it still comes along with certain costs, analogous to the costs the proponent of MH faces in accounting for worlds with a screeny state space. The cost is either the qualitative ontological cost of positing the existence of entanglement relations, the ideological cost of providing a nominalistically acceptable surrogate for such relations, or the explanatory cost of positing brute necessary connections between fundamental property instantiations - which amount to *ad hoc* restrictions on the state space. Call this cost a lack of *elegance* since it seems to require a loss of syntactic simplicity for the theory, introducing terms to explicitly rule out certain putative combinatorial possibilities.

Sider's screen world is the paradigm of atomistic orthodoxy: the space of physical possibilities is exactly what we'd expect if the world was made up of a plurality of independent entities. The monist, as we've seen, can reproduce such a world using her own ontological resources. However, such accommodation comes at an unattractive cost. In the next section, I want to explore some other highly simplified state spaces, in the vein of Screen World, and see how well each of the options on the table: monism (more generally MH), pluralism (more generally ML), and local holism (LH) can accommodate them. Later on we'll touch on the rather thorny question of which of these super-

¹⁴ Schaffer (2010).

¹⁵ Covered in §6.1.1, §6.1.2, §6.1.3.

simplicistic models might be usefully analogous to the state space of the actual world (if any). We'll see that the local pluralist doesn't keep the higher ground she took in the Screen World thought experiment when it comes to modelling some other simplified state spaces.

7.4. Strange New State Spaces

The first world I want to look at I'll call Chess World. Chess World only has four possible states—white, black, chequered_w and chequered_b:¹⁶




Figure xi. The four chessboard states.

Is this a world made up of fundamental, independent sub-world objects? On the face of it, this appears implausible. The possible qualitative state for each pixel is either black or white (on/off) as in the screen world. Furthermore, for each pixel there's a possible world where it is one colour and a given neighbour is a different colour, and a possible world where it is one colour and the given neighbour is the same colour. However, despite all this, if we want to treat each pixel as not directly depending on one of the four overall image states then we have to say, instead, that each of the pixels is highly *interdependent*. Once the state of any two adjacent pixels is fixed (of any two pixels that would be opposing colours on the chessboard), the state of every other pixel is fixed. Moreover, the state of one pixel completely determines the colour of half the pixels.

¹⁶ For chequered_w and chequered_b to be considered different states appears to require a haecceitist account of possibility: since one state is just a 90-degree rotation of the other it has to make a difference *which* pixel has *which* quality to make sense of the white-on-the-right distribution as distinct from the black-on-the-right distribution (alternatively, perhaps, there's just primitive rightness and leftness built into our irreducible distributions). Sider (2007) goes on to argue that the monist can't account for haecceitist possibility, but as Cornell (MS) points out, whether doing so is even desirable depends very much on our assessment of issues already at stake between the monist and the pluralist—specifically issues to do with identity and individuality in quantum mechanics. Anyway, whether Sider is right or not, all my toy worlds will make the simplifying assumption that everyone at the debating table can and should distinguish these possibilities—this shouldn't matter to the results, as I could just as well run my arguments using purely qualitative states.

It's hard to get a much more interdependent world than chessboard. This interdependence can't be explained, by the pluralist, by supposing that each pixel is wholly dependent on the state of the board as a whole: for dependence (simpliciter) – we are assuming – goes hand in hand with being a derivative thing, while the pluralist's claim is that the individual pixels are fundamental. The pluralist must instead model the situation in terms of co-dependence or partial dependence, and argue that this weak form of dependence does not imply non-fundamentality. Specifically, each of the pixels enjoys a degree of dependence on all the other pixels, but no one pixel entirely and asymmetrically depends on any one other pixel. One way or another, then, the pluralist must posit a vast number of *necessary connections* between the fundamental entities she posits. This, as I contend above, is a cost.¹⁷

If the actual world is qualitatively analogous to the chessboard world, we're faced with positing, given pluralism, a vast number of necessary connections between a large number of fundamental things plus (say) a couple of fundamental properties (corresponding to black and white). By contrast the monist just needs one fundamental thing, no necessary connections between fundamental existences, and (suppose) just three fundamental distributional properties (corresponding to wholly black, wholly white and chequered). For the price of very modestly increasing the number of fundamental properties, she ends up with a much more simple and elegant theory overall.

The above helps explain and justify the common assumption in the literature that a localized, classical physics lends weight to a pluralistic metaphysics, while the discovery of quantum holism gives (defeasible) evidence in favour of a monist ontology. However, Screen world and Chessboard are both extreme cases – the former of a highly localised world, while the latter is extremely holistic. There are many conceivable intermediate cases which are epistemically open analogues of the actual world. Consider Switch world:

¹⁷ For some useful discussion on the status of the Humean constraint in contemporary ontology see Cameron (2008*b*).




Figure xii. Example switch world states.

Switch world has many more qualitative possibilities than Screen world, but these are still only a subset of those enjoyed by screenworld. For each set of what we might call “paired pixels” only two qualitative states are available: either the left pixel is on and the right pixel is off, or the other way around (i.e. it cannot be the case that both pixels in a horizontal pair are on, or that both are off):




Figure xiii. Paired pixel states

The Screen world possibilities of such “paired pixel” regions being qualitatively homogenous are ruled out. So the state space size is 2 mutually exclusive qualitative states to the power of 8 pixel pairs = 2^8 combinations.

In the first instance, we might usefully compare how the standard pluralist and the monist should model the Switch world. The monist requires 256 distributional properties to account for the number of possible qualitative states of the whole system. This is significantly fewer than the number of distributional properties required to model Screen World, but still a large number: much greater, for instance, than the two qualitative properties employed by the pluralist. On the hand, though, the pluralist committed to 16

fundamental pixels and needs to posit a necessary connection between each pixel and its partner: this is less extravagant than the pluralist's modelling of Chessworld, which required connections between each pixel and every other pixel (120 pairings), but still a large number of necessary connections.

We might take this as a jumping off point for a close-quarters debate over which is the marginally more attractive way to model the possibilities for this world and its more complex analogues. However, the lesson I draw is that neither performs particularly well when faced with Switch world. Certainly, the pluralist model of Switch world is considerably less elegant than its model of Screen world, and the same goes for the monist's account compared to its model of Chessboard. We should hope do better, and we can. We need a form of pluralism that utilizes resources from the monist's ontology: what I'm calling Local Holism (LH). On LH, we allow that intermediately-sized entities can be fundamental (i.e. objects larger than the smallest qualitatively variable units - points or "pixels" - but smaller than the whole system). This requires the adoption of the Monist's fundamental distributional properties (or similar strategy for dealing with fundamental heterogeneity).

Given LH, the obvious candidates for the fundamental furnishing of switch world are the wholes composed of two paired pixels - i.e. the wholes composed by two pixels whose intrinsic properties modally co-vary with one another. Call these fundamental objects "switches". Like the pluralist, we will only need two fundamental properties; but like the monist these will be distributional properties. The properties required are "left-half-off-and-right-half-on" and "left-half-on-and-right-half-off". By allowing each switch to have one of these properties, completely independently of which of these properties is had by any of the other switches. By allowing each fundamental object to vary its properties independently of all the other fundamental objects, this strategy requires the posting of no fundamental necessary connections. So, tallying up, LH can model switch world with as many fundamental properties, half the fundamental objects and none of the necessary connections required by the pluralist. Against the monist, while LH posits a few more fundamental objects, she requires far fewer fundamental properties (while neither require necessary connections between fundamental objects). All thing being equal, then, I think we should prefer LH over either orthodox pluralism or monism if the actual world is suitably analogous to switch world.

However, there are other screen-like worlds we can envisage on which neither monism, orthodox pluralism or LH fair particularly well. Consider Variety world:




Figure xiv. Variety states (X = disallowed)

Variety world, on the face of it, is much more like Screen world than chessboard: Variety world has all the possible states of Screen world, bar two. Variety world can never be wholly white or wholly black. Might it then be fruitful to model Variety as the pluralist models Screen world? This accounts very well for the very large number of possible qualitative states that Variety has: from just two qualitative properties – black and white – and the assumption that each pixel can have either of these properties (more or less!) independently of what properties the other pixels have. However, how to account for the “more or less” caveat? It seems we still need a whole network of necessary connections between the pixels to ensure that if one pixel is off, say, at least one other pixel is on. But as Variety is not very similar to Chessboard, we shouldn’t expect monism to fair much better either. Indeed, the monist requires all of the vast number of fundamental distributional properties required for the monist to account for Screen world, save two (homogeneous black and homogeneous white).

Does LH, which so adeptly accounted for Switch World where the others failed, do any better than its rivals? It doesn’t seem that selecting any group of multi-pixel sub-world objects as fundamental will account for Variety world’s possible states very elegantly. It would require a fair number of distributional properties – a minimum of 4, if it posited switch-sized (2 pixel) fundamental objects. But unlike with its modelling of Switch world, it doesn’t escape the need for necessary connections: each sub-world object must be given the freedom to have any permutation of black and white pixels intrinsically; but this freedom must be restricted such that if e.g. all other fundamental sub-world objects instantiate an all-white distribution property, it instantiates an all-black or heterogeneous distributional property.

So, if the actual world is analogous to Variety, it looks like all the ontologies on the table are (more or less) equally unattractive. But what does all this rather abstract speculation about pixelated worlds amount to? The question is, which of these toy worlds most closely resembles the actual world in its qualitative structure (assuming any resemble it enough to be useful to our purposes)?

7.5. Screen Worlds and the Actual World

Schaffer's case for believing in a world of emergent or holist properties comes primarily from quantum mechanics. Specifically, he thinks the fact that the world forms '*one vast entangled system*' (2010, p.52) is sufficient to establish that the world's physical properties are maximally holistic and incapable of being given a completely local, pluralistic explanation. From this starting assumption, a compelling – though by no means unassailable – case for Priority Monism can be built. If the world is an entangled system then the properties of no part of that system can be specified without knowing the whole state of the system— the plausible ontological upshot being that the properties of any part of the system depend on the properties of the whole system. But is the assumption correct?

Schaffer's argument for the world consisting of a single entangled system rests on the observation that the Schrödinger equation tells us that as time progresses, systems which started out separate will eventually become entangled: "Schrödinger evolution tends to spread entanglements," Schaffer claims (2010, p.52). He concludes that, as such, even if the world did not start out in an entangled state – and here he quotes Roger Penrose¹⁸ (2004, p.591) – "...eventually every particle in the universe must become entangled with every other" (Schaffer 2010, p.52).

Or, rather, depending on the initial conditions at the Big Bang it may be that his assumptions: 'yield a pluralism in letter but not spirit, in which the universe contains many vast isolated bubbles. Everything we ever encounter—everything in our bubble—would still form one entangled system.' (p.52f) That this would lead to 'a pluralism in letter' is, I think, more problematic than Schaffer's gloss suggests, in the context of the current debate: rather than monism and the priority of the whole being a necessary truth, what we in fact have is varying extremes of LH depending on the potentially contingent, certainly unknown, initial conditions of the universe: we have a plurality of things which must be treated separately from one another. If we live in one of these vast entangled bubbles, then, while things are certainly very different from how the Humean would have us picture them, this doesn't seem like Priority Monism either in letter or spirit.

Even if we grant him all of the above, however, including favourable initial conditions, a further big assumption made explicit by Schaffer (2010) is that the Schrödinger equation is the only mechanism that governs the time-evolution of quantum systems. But this is far

¹⁸ Whom, as it happens, is – as I point out in the Appendix to this chapter – far from sympathetic to the conclusion that entanglement is universal.

from trivial or obvious: In fact, I outline in the Appendix to this chapter some reasons for actively doubting this assumption, depending on one's interpretation of QM. I suggest that if collapse interpretations of QM are correct, the world as a whole is most naturally described as a collection of localized facts— not maximally localised to, say, individual spacetime points, but likely still on a microscopic rather than macroscopic scale (certainly on a scale immensely smaller than the vast entanglement bubbles of the previous paragraph— and so much more a pluralism 'in spirit' even by Schaffer's lights). As such, there is considerably less pressure to accept a monistic account of property dependence; for the properties of certain parts of the world can be fully specified without reference to the state of the world as a whole. Given that the properties of parts of the world *can* be specified independently of the whole, there would also a *prima facie* cost to the monist for not doing so: the monist misses out on the combinatorial explanatory advantages associated with pluralism.

If, contra Schaffer's set of controversial assumptions, entanglement is not ubiquitous, then what we have instead of one big entangled cosmos is a plurality of separate, small-to-mid-sized systems, exhibiting holistic behaviour internally but independent from each other. Further to this, I also suggest in the Appendix that even if entanglement is universal, as Schaffer argues, the ubiquitous phenomenon of decoherence – by which large quantum systems come to exhibit approximately classical (and thus quasi-separable) behaviour – means that the argument from universal entanglement to monism from theoretical virtues is still not straightforward: On a large scale the world not only looks localised, rather than holistic, but obeys laws that are classical in the limit, and these are data that the monist must somehow accommodate.

So, with all this in mind how do the various screen worlds shape up when it comes to assessing whether any are a good (albeit highly simplified and approximate!) fit to the actual world? Well, Sider's Screen World is paradigmatically classical: If we're taking seriously the lesson of quantum entanglement, that the possible states of an entangled composite system are not obtained by free re-combination of the properties of its simple parts, then Screen World can be ruled out from the start as a useful analogue of the actual world, even if it turns out that a subset of fundamental properties can be treated combinatorially. This is bad for orthodox pluralism: our world is not sufficiently analogous to the toy world that the pluralist can account for elegantly and economically.¹⁹

¹⁹ Though see Miller (2013) for an argument that endorsing the Bohm Theory interpretation of QM plus a radically Humean reduction of the pilot wave to best-systems style laws across a 4D manifold

Narrowing down the options from the remaining three will depend on what precisely we take quantum theory to be telling us about the qualitative structure of the world. However, I think one of these options – Chessboard – can also be ruled out fairly quickly. From just the evidence of our eyes, let alone physical theory, we can see that the world as a whole instantiates more than a handful of uniform distributional properties. There are many physically possible ways the world can be, and things on an everyday macroscopic scale at least appear to vary independently of what’s going on elsewhere in the universe. The world rather obviously allows for more localized qualitative variance than the sort of wholesale holism caricatured by Chessboard. The physics backs this up: On a large scale our world does look approximately classical – i.e. more like Screen World than Chessboard. So the favourite candidate of monism meets almost as swift a fate as the pluralist’s poster child, Screen World.

Switch World and Variety are more promising candidates: both are able to account for local heterogeneity while taking onboard some degree of quantum holism. Switch World allows for quantum effects at the micro-level, while also accounting for observed local heterogeneity at the macro-level. Each “switch” is somewhat analogous to a pair of entangled particles: Each side of the switch can either be black or white, but one side of the switch being a certain colour is not a matter that can be fixed independently of the colour of the other side of the switch. Similarly, one member of a pair of entangled particles has two possible states – spin up and spin down, say – but each member can only instantiate a given spin-value if its partner instantiates the opposite spin-value. On the other hand, since each “switch”, or pixel pair, can vary its qualitative properties independently of the qualitative properties of the rest of the grid, Switch World allows for a great deal of qualitative variation locally: for instance, most of the grid could be made up of off-to-on switches, with a cluster of on-to-off switches in the middle, or the left half of the grid could be made up of off-on switches with a random mixture on the right of the grid; or the whole grid could be made up of off-to-on switches with a single on-to-off switch in one corner, and so on. By restricting qualitative interdependence to small units – the pixel pairs or switches – it allows for appearance of qualitative independence at the macro-level by genuinely ensuring independence between small sets of pixels.

can, contra for Maudlin (2007), allow one to recover the Humean mosaic picture of reality as localised spacetime points. The joint commitments of Bohmian Theory, a best systems account of laws and reductive account of Bohm’s ontology might not give this view especially wide dialectical appeal; certainly, it make it hard to assess the merits of it succinctly. But perhaps this represents one promising avenue to explore for those who want a less concessive account of how we should modify our ontology in light of the holistic features of QM.

Of course, quantum mechanics allows for entanglement in systems larger than two particles— but these pixels worlds are obviously huge simplifications. However, as we've noted²⁰ it's an open interpretive question whether quantum entanglement is likely to occur in systems much larger than the characteristic scale of quantum mechanics. If there is a cut off-point to the size of such systems in nature, then we might take Switch World to be sufficiently analogous to the actual world for our purposes. It is relatively straightforward and desirable for the proponent of LH to allow that fundamental objects need not be of a fixed sized (i.e. pixels or points vs. the whole world) but rather could come in a number of different sizes (e.g. two, three, four pixel - or particle - systems, etc.). So, one of the two live options here tells strongly in favour of LH, which can account for switch-like worlds with a simplicity and elegance that neither orthodox pluralism or monism come close to matching.²¹

But what if there is no cut-off to quantum effects in nature once we approach the macroscopic scale? In this case, it doesn't seem that LH provides a very natural account of what's going on, since quantum effects will always extend beyond, or between, any sub-world objects that LH takes to be fundamental. The monist might then hope to come to the rescue: positing a single fundamental object encompassing the whole world ensures that the scale of quantum effects will never exceed the scale of our fundamental ontology. However, we've already seen that the sorts of holistic worlds that monism excels at modelling, like chess world, would already to be ruled out both by quantum mechanics and the evidence of our senses: In these sorts of worlds, there is little room for qualitative variation in one sub-region of the world without variation in the rest of the world. By contrast, in the actual world, at scales where quantum effects become negligible, there can be a great deal of localized variation and heterogeneity that appears independent of what is going on elsewhere in the world (as we can both see by looking around us and by examining the theory). We're left with two things to reconcile: On the one hand, the Schrödinger equation without the addition of a collapse postulate tells us that (more or less) everything will be in a giant entangled state, such that the qualitative possibilities for the parts of the system can never be specified without reference to the overall state of the system. On the other hand, sufficiently large sub-systems will look and behave pretty much *as if* they are independent of one another.

²⁰ More details in the Appendix.

²¹ This should hold, to a greater or lesser extent, provided that the fundamental physical quantities that require holistic treatment are in the majority.

Of course, this isn't (on the face of it) a problem for QM itself, which shows mathematically how entanglement leads to the washing out of the interference effects that are prevalent in small isolated systems and give rise to distinctive quantum phenomena. However, it does present a bit of a stumbling block for the ontological debate as we've set it up. For the sorts of worlds which accommodate both data points are most closely to analogous to Variety world (of the four options, anyway). In a variety world you can't know a given colour is permitted for an arbitrarily large system without knowing that at least somewhere in the world there is a pixel of the opposing colour (or a heterogeneously coloured composite system) - so possible qualitative states for any proper part of the world can only be specified with reference to the whole system. Yet, at the same time, subsystems can vary hugely independently of one another, provided that the relatively weak constraints on what the overall state of the system can be are obeyed.

But we've already seen that none of the three options on the table - orthodox pluralism, LH or, indeed, monism - model such worlds in a very simple or elegant manner. The monist requires myriad fundamental distributional properties to reproduce what could be produced much more simply via a combinatorial approach using far fewer fundamental properties, given a few caveats about which of the combinatorial properties need to be ruled out. Meanwhile the pluralist, with the combinatorial approach at her disposal, needs to posit lots of necessary connections between fundamental entities to ground said caveat. Where LH fits in to the picture is more complex, but when modelling standard Variety World, requiring both a large number of distributional properties and necessary connections, though fewer in each case.²²

What to make of all this? Well, putting a positive spin on the results - that each contender does more or less equally well - we can at best conclude that the issue of holism in quantum mechanics vs. more classical systems doesn't have the tight connection with the pluralism vs. monism debate than has been assumed. The upshot is a sort of multilateral disarmament: the monist can't use quantum non-locality as ammunition against the orthodox pluralist; while the orthodox pluralist can't use the virtues of combinatorialism against the monist. Similarly, neither combinatorialism nor holism can be used to garner support for or against LH. The ontological debate becomes detached from the physics, even at the level of which ontological theory has the best extra-empirical virtues given the physics, since -as it happens - the world and physical theory don't co-

²² The problem will be worse the greater the number of physical quantities that need to be treated holistically.

operate in helping us choose between our rival hypotheses. This might be welcome or worrying, depending on your perspective.

What is worrying, however, or should be worrying to anyone committed to this approach, is the prospect that we're going about things in entirely the wrong way. By theory choice criteria we were happy to accept at the outset of the investigation, none of our initial theories did very well at all (relative to how much better they do in modelling more amenable worlds tailored to their respective strengths). It's tempting to conclude that the approach of positing a mereological structure and designating a subset of the entities in that structure as independent and fundamental and the rest as dependent, non-fundamental is simply the wrong way of going about trying to model an entangled world, ontologically. The obvious competing approach in the same methodological ballpark is to return to the strategy adopted in the previous two chapters,²³ separating out the notions of fundamentality and independence not just conceptually but extensionally, and positing emergent objects or properties as fundamental yet dependent entities. Having fundamental objects at multiple mereological levels opens up more options for possible ontological theories when it comes to accounting for diverse state spaces – this additional flexibility might be needed to account simply for qualitative interdependence in an entangled world.

To illustrate the potential that's opened up, let's return to our simplified state spaces—specifically, Variety World. All of the dependence-fundamentality linking views considered had trouble accounting for Variety World in an economical fashion. But what if we allow both the 4x4 screen as a whole to be fundamental *and* the individual pixels? My suggestion here is that (non-distributional) fundamental qualitative properties are assigned to each of the pixels, just as the pluralist does, with the World's colour properties being both dependent on, and merely derivative from, the colour properties of the pixels:




Figure xv. Extended world and many simples

²³ Following Barnes (2012.)

The World may derivatively, in virtue of its fundamental parts, have the property of being black on its left edge and white over the rest. However, it also has a less-fine grained fundamental property: Namely, the property of being in a *heterogeneous* state, as opposed to a *homogeneous* one (i.e. black or white all over). The brute fact in this Variety World's ontology is perspicuously described by how we would intuitively characterise the state space restriction in play— that the World as a whole is necessarily heterogeneous. That is, in this model, the fundamental colour property the World instantiates must, of the two theoretically allowed options, be the *heterogeneous* one rather than the *homogeneous* one.

There is two-way dependence between the colour properties of the whole and the parts, though this does not quite amount to having to posit symmetric dependence.²⁴ If you're comfortable with symmetric dependence, we can posit something of a dependence loop: The heterogeneity of the whole depends on the colour properties of its parts, taken collectively, while the colour property of any individual is dependent on – constrained by – the heterogeneity property. The alternative, though, is just to take the heterogeneity property to be both independent and fundamental, simply constraining the colour properties of individual pixels. So, it's more of a *priority of the whole* situation, but with fundamental sub-world objects that are constrained by the whole, but also partially independent of it.²⁵ The derivative (fine-grained) colour properties of the whole will depend on the properties of the pixels, but nothing further depends on what the derivative colour properties of the world are, so the loop is not closed:



Figure xvi. Chessboard and heterogeneity fact.

The necessary connections the pluralist posits between each of the pixels (120 in total) are in this case not brute or unexplained, but derivative from a smaller class (just 16)

²⁴ Though Barnes (Draft) gives us good reason to be comfortable with this too.

²⁵ This idea of partial dependence goes beyond Barnes' own proposals: The thought is that an otherwise independent and fundamental object can be constrained in what properties it can have through a necessary connection of some form being posited between it and some other fundamental object.

necessary connections between each individual and the World as a whole: The heterogeneity property forces at least one pixel of each colour to be instantiated in the grid— but this is ultimately due to the relation between each pixel and the world, rather than because of complex interrelationships between pixels. In a predominantly white world, the last black pixel couldn't turn white, since this would necessitate the impossible state of the World entity being fundamentally homogenous. The advantage over the monist is that the World only needs to be fundamentally in one state, rather than have a possible fundamental state for each allowed lower-level distribution. The advantage over the pluralist is the reduced number of irreducible necessary connections.

So, this is just one suggestion for how allowing absolute fundamentality on multiple mereological 'levels', with dependence relations of varying directions between them, may help with the modelling of recalcitrant state spaces. More radical departures, such as the adoption of a fundamentally structuralist or relational ontology, might be warranted if less minimal revisions such as decoupling do not ultimately prove satisfactory.²⁶

The most interesting immediate result as I see it, however, is that if entanglement is not as ubiquitous as Schaffer claims (for any of the suggested reasons), LH comes out as the clear winner from the options we've been considering. Given an interpretation of QM that includes a collapse postulate, adopting LH seems to be the best way of accommodating the lessons of quantum entanglement into your traditional metaphysical theory, as opposed to sticking with classical pluralism or embracing maximal holism in the form of Priority Monism. But the dialectical situation is not symmetric in this regard: If entanglement is ubiquitous, LH is not obviously disadvantaged relative to its nearest rivals. The present results suggest that orthodox pluralism and monism fail to elegantly account for universal quantum entanglement, whichever class of interpretations wins out. So, absent some independent reason to completely rule out collapse-based interpretations, this tells in favour of LH: the significant benefits it has over its nearest rivals are hostage to fortune, but even in the worst case scenario we can't do much worse by endorsing LH. So, overall what I think we have here is a reasonably compelling - if highly defeasible and hostage to empirical fortune - argument in favour of LH. Given everything that's been said, we would at least have some principled reason for preferring LH over its near rivals, priority monism and orthodox pluralism.

²⁶ See, e.g., Ladyman & Ross (2007) or, for a more moderate take, French (2014).

7.6. Merging (an intuitive gloss)

So, I've advocated a view on which "medium-sized" objects are fundamental: They're medium-sized in the sense that they're neither at the bottom of the mereological hierarchy on the top.²⁷ The considerations I've offered from quantum mechanics suggest these medium-sized objects would in fact be very, very small indeed and perhaps not even spatially continuous (*one half here, another the other half all the way over there...*), reminiscent of the universalists scattered composites (*your left foot and the red spot of Jupiter...*) but mereologically simple.

This view has, I think, both nihilist and non-nihilist formulations, which almost exactly parallel the choice between existence monism²⁸ and priority monism.²⁹ On LH, monism is in effect true for a *localised* region – and, as I've cashed it out, the proponent of local holism will be employing the same ideology as the monist to explain how there could be a heterogeneous extended object with properties not derived from its parts.

My reason for preferring a nihilist formulation over non-nihilist formulations is an entirely defeasible one: Namely, that nihilism is *simpler* than belief in composite objects, and all else equal we should prefer simpler theories. Now, there are different ways of cashing out and justifying such appeals to Ockham's razor, of varying degrees of plausibility: I say a bit more about this and try to show that parsimony is genuinely applicable to the composition debate in §8. For now, though, I just want to try and outline what a nihilistic version of local holism might look like.

So, how should we picture a world which has all the features I've argued for, i.e. a world that:

²⁷ Where the theoretical mereological hierarchy concerned is one constructed from our everyday talk and common sense, informed by fundamental physics and if necessary stipulatively extended to accommodate additional ontological commitments of the view. In the final analysis this might amount to effectively cutting and splicing together several incompatible description of reality, but at this point it's just a heuristic for introducing the idea of a medium-sized object.

²⁸ As set out in Cornell (MS); Horgan and Potrc (2000); Schaffer (2007). Some (e.g. Ross Cameron – in conversation) might be uncomfortable with my term use of the 'existence monism' here, given my metaontological commitments (i.e. I say that Moorean objects like tables and chairs *exist* in the ordinary sense, whatever one's ontology). However, by 'existence' here I mean existence in the ontologically privileged sense— while I reserve the label 'priority monism' for monists with a more inflationary view of the fundamental (i.e. that such things exist derivatively – whatever that means – but are still in the domain of the unrestricted quantifier of ontologese.)

²⁹ As argued for in Schaffer (2010).

- Has no proper parthood: nothing ever composes anything larger or has smaller things as parts.
- Contains systems of contingent size ranging, from unextended single particles to extended yet mereologically simple “multi-particle” systems. (With the upper-bound for system size governed by a stochastic collapse postulate, given the argument in the Appendix).
- “Multi-particle” systems having properties “over and above” what the physics tells us are the properties of their single particle subsystems.

The best way to conceive of such a world, I suggest, is to do away with the familiar notions of composition and emergence in favour of a single equally intuitive notion—merging. I take merging to be a close cousin of composition, in the sense that it’s another way in which some (small) things can come together to create some further (bigger) thing. However, unlike in an ordinary cases of composition – call this ordinary sense ‘construction’ – we can say loosely that the merged individuals “lose their individual identity”.

Given the way I’ve talked about it just now, composition is an *event* – the occurrence of some things becoming appropriately arranged such that they bring about the existence of some further thing. This is how the concept of composition is often introduced in the literature, including in van Inwagen (1990). However, in practice metaphysicians tend to be less concerned with composition as an event, and more with the relation of composition, or parthood: The composition relation is the relation that holds, at any given time, between the whole and the parts (taken collectively), while ‘parthood’ is the (transitive, anti-symmetric) relation that holds between each part taken individually and the whole. One of the reasons for this, I suspect, is that even among those philosophers who think that some things compose other things, and that some objects have parts, a substantial number don’t believe that there are any conditions under which some things could come to compose some further thing. The mereological universalist thinks that when you have some things they always compose some further thing: so, according to the popular universalist position, there are ubiquitous instances of the composition/parthood relation but no events of some things that previously failed to compose some further thing now composing some further thing. However, considering an event of composition is often used as an easy, intuitive way to grasp the relation. I’ll be taking a cue from this strategy when contrasting composition and merging – except that in the case of merging I’ll be arguing that there is no metaphysically interesting relation behind the event, and thus that merging is compatible both with the letter and the spirit of mereological nihilism.

Composition, then, is the familiar process of rearranging things, like bricks, so that you get something new, like a house (*see Figure 1*). Other putative examples of composition:

- Nailing some planks together to make a ship
- Reattaching a handle to a hammerhead
- Beavers using wood to build a dam
- Particles of sand accumulating over time to form a sandbar
- *Constructing* a house (re-arranging some bricks, etc.):

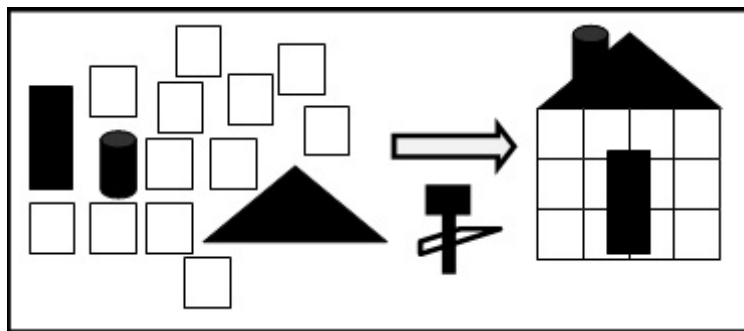


Figure xvii. Building a house

Merging, on the other hand, is a process by which some things come together to create some new thing *but don't survive the process*, or are at least *radically changed*.³⁰ Again, some putative examples of what I have in mind might help:

- Two light beams are combined and become a single light beam.
- Two companies are merged or amalgamated into a single company.
- Some rain drops collide, becoming a single rain drop.
- Two clouds merge and become a single cloud.
- Two computer networks are connected and become a single network.

³⁰ See Bennett (*draft*).

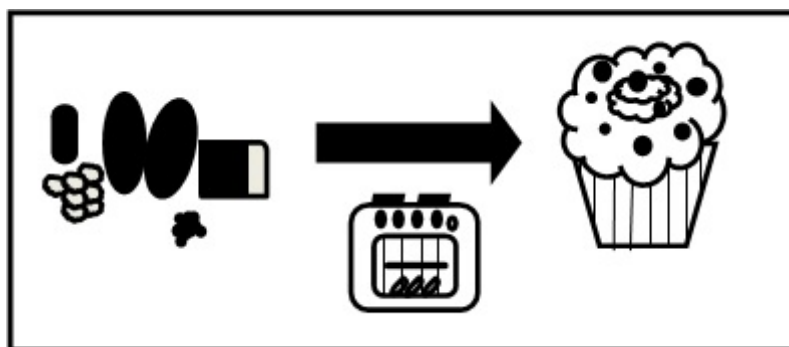


Figure xviii. Making a cake

The above are intuitive examples of composition and merging, respectively— but it’s up for grabs whether any of these are *genuine* cases of the phenomena I’m trying to get a fix on (especially in the case of merging, as we’ll see). Considered as events, the key similarity between the two sorts of case is that you start out with some things (eggs and flour, brick, water droplets, planks of wood, etc.), do something to them - rearrange them, squash them, let them run into each other, cook them, whatever - and end up with some new thing (a house, a cake, a ship, a bigger water droplet...). What’s more, there is a certain kind of intimate relationship between the things you started out with and the new thing you ended up with— the new thing will inherit the location, mass and many of the physical properties of the original things, for example. The old things and the new thing are made out of the same matter, the same “stuff”. Whither the distinction? I think there are two broad distinctions worth focusing on. For now, both will be described in fairly sketchy terms: we’ll get to a more rigorous analysis in due course.

The most obvious difference is what happens to the original things after the event. In paradigm cases of composition, the parts after composition are largely left unchanged from how they were before composition occurred - and, moreover, we can diachronically identify the original parts with some things post-composition. On the other hand, in cases of merging, the properties of the parts are often radically different once merging has occurred. In particular, the original things tend to be less *separable* after merging: That is, for one thing, they tend to be less physically separable: It’s easy to pull apart and recover the original Lego bricks when you’ve made a Lego house. It’s very much harder to recover the original eggs or grains of sugar after baking a cake. Further, for want of a better term, the ‘original parts’ are not discernible post-merging - we cannot diachronically identify the entities that are the original parts with any entities post-

merging.³¹

When you start out with some eggs and flour that you then mix together and bake into a cake, it's not only mechanically difficult to recover the original eggs and flour, but unclear whether the eggs and flour still exist: it's seems right to say the cake contains egg and flour, but odd to say that the individual eggs or flour granules are still present – that the egg or flour are now *parts* of the cake.

The big difference I want to draw attention to is that the properties of the resultant merged entity are usually very different from the properties one would expect from the whole in a case of construction. Specifically, the properties of a merged entity tend to be *novel*, in a particular sense, in a way that the properties of wholes produced by construction tend not to be. When we simply put things next to each other or stack them on top of each other— paradigm cases of construction – we expect the properties of the whole to be just the properties of the parts, taken collectively. An event of construction is just changing the locations of the parts, rather than any of their other properties, and the properties of the resultant whole *supervenes* on the properties of the parts after the event. The properties of the whole are just the properties of the parts *in a certain arrangement*. Intuitively, though, merged entities have properties over and above those of the original parts.

Are there any events that fit our intuitive conception of merging in nature? Everyday instances of putative merging can be explained – or, more aptly, explained way – by an atomistic account of what is going on: All that's really happening is that the simple parts of the two composites you started out with are being rearranged: the resulting composite and its non-simple parts will have novel intrinsic properties only to the extent that their simple parts now stand in different spatial relations from how they did before. This accounts for the apparent novelty of the new composite, and also the changes in the original objects and their parts, including their potentially ceasing to exist depending on what you think the persistence conditions for composite objects are. That no matter is created or destroyed is explained by their being the same simple parts you started out with, just re-arranged. However, I think our common sense conception of merging is prior to this atomistic explaining away of the phenomenon, and that we should be open to finding events in nature that more perspicuously correspond to our intuitive idea of merging.

The intuitive gloss I've given of merging so far is compatible with both a nihilist and

³¹ Or, perhaps, we cannot *determinately* identify them.

pluralist view - for all I've said, merged entities may be composite. The distinctive claim of merging was that there is a difficulty in diachronically identifying the 'original parts' which merge with any of the entities post-merging. How best to explain that? I suggest that the best explanation of this is afforded to us by nihilism: when some simples merge, the reason that you cannot synchronically identify those simples with anything post-merging is that the merged entity is not a composite, and - given the plausible assumption that they can't all be identified with the non-composite merged entity - there is thus nothing available to identify them with. The event of merging is, on the nihilist account, the coming to be of an extended simple, brought about by the merging of a plurality of simples.³²

By contrast, the most intuitive universalist gloss on merging - and the one which also seems like a reasonably apt for the entanglement case - is that it involves a 'switching' of priority relations: cases where the whole becomes prior to the parts. The distinction between nihilist and non-nihilist interpretations of merging chimes with our intuitive ambivalence concerning what happens to the original parts in everyday merging cases: the non-nihilist view of merging corresponds to the intuitive feeling that the eggs and flour granules in some sense survive when we bake the cake but are now very hard to separate out from the cake as a whole. Alternatively, there's intuitive pressure towards the view that the cake and flour granules are indeed destroyed, though there's still "the same amount of stuff" in the world.

But, why care about merging at all? Well, I think that it gives us an intuitive way to think

³² Similar considerations lead Bennett (*draft*) to conclude, in her analysis of similar cases, that the metaphysically interesting relation here is *diachronic*, holding between the things at t , immediately before the new object comes to be, and the new thing at t , immediately after it has come to be. To the extent I feel under any pressure to posit a real external relation in this case, however, it is to provide a standard causal explanation of how the things at t *caused* the thing at t to exist. So, either I'll already be committed to this extra relation (of diachronic causation) to ground causal activity more generally, or I'll have a reductive account of causation that does not require me to posit a genuine diachronic relation here or anywhere. Since Bennett's own view is that ontological dependence and causation are of a kind, she could perhaps argue that taking the first horn is essentially to commit oneself to ontological dependence, leaving one ideologically in no better position than the local holist with an inflationary ontology. However, to have much impact on my argument, it would have to be shown there's both good reason to be non-reductive about causation *and* that the two relations are unified by more than just analogy.

about LH *nihilistically*. If you accept my tentative proposal for a form of LH based in quantum mechanics, then one can diagnose it nihilistically as the occurrence of many instances in the world of both merging and unmerging. But even if you don't take entanglement to be a reversible process, if you buy Schaffer's argument that quantum entanglement entails holism, there's good reason to believe in merging: As Schaffer (2010, p.52) admits, it's an epistemically open and perhaps nomically contingent matter that everything started in an entangled state at the big bang: if not, then the world contains at best vast "bubbles" of entanglement (i.e. potentially as large as everything in our light-cone) which we can expect to remain ever separated in an expanding universe due to the fact that causal interaction is limited by the speed of light. However, we should still take it to be the case that *were these bubbles to collide*, they will become entangled with each other and form a single physical system: we can imagine this happening, say, if they were in a spacetime with a weird global geometry, that would send them converging on the same point even though they're hurtling apart. But it's not the mechanism that matters: If God were to pick them up and smash them together, the laws tell us they would form a single entangled system.

Now, we could go with Schaffer's very strong assumption that monism is a *metaphysical necessity*. While this may come with theoretical benefits of its own, it entails that even non-entangled, nomically independent systems, are all really part of a single indivisible whole. I've argued here that this is a theoretical cost: the most simple solution is, by and large, to treat nomically independent systems as metaphysically independent objects (while going along with Schaffer in treating nomically dependent systems either as metaphysically dependent parts of a greater whole, or nihilistically by diagnosing the entangled system as an extended simple). Plus, on the face of it, it makes sense to suppose that the mechanism that lead us to concede holism in the first place should be the mechanism by which it is wrought: that is, that quantum entanglement is a nomically necessary and sufficient condition for ontological holism (pluralistically or nihilistically construed). This route to seeing ontological holism as restricted gives us at once a satisfactory answer to what we might call the 'Special Merging Question' (SMQ): When do objects merge? *Answer*: When they become entangled with one other (i.e., when they have physically interacted with each other).

But let us set immensely contentious arguments from quantum mechanics aside for a moment: The general metaphysical lesson, I propose, is that once we've made conceptual room for the possibility of merging, as I've described it, is that every putative case of emergence is a putative case of *merging*. This being so, we have the theoretical option of dispensing with emergence and posting merging in these cases, while treating the datum

that *the whole is something over above its parts* with complete metaphysical seriousness. Let's suppose, loosely taking inspiration from van Inwagen and Merricks, you take *living things* or *human persons* to be in some sense more than the sum of their parts, perhaps even having novel causal powers. One can now account for this by saying that while things like atoms, molecules and cells are in general independent entities, once they become part of a life or a conscious being, they surrender that independence to the resultant unity - they merge. In so far as merging is best understood, as I claimed, nihilistically as the coming to be of an extended simple, this allows the nihilist to circumvent the need for the positing of dependent objects with novel or fundamental properties, and thereby allows the nihilist to retain the orthodox framework that ties dependence to non-fundamentality.

7.7. Local Holism and Ideological Economy

I've argued in this chapter that Local Holism offers a simple and elegant theory that accommodates examples of emergent properties from quantum mechanics, and shown how a nihilistic Local Holism can be described by appeal to the notion of 'merging' (to be explained, I claimed, as the coming to be of an extended simple). However, I can see a potential objection that any putative simplicity gains are washed out by a bloated ideology: On the one hand, the "pointy" or *punctual* nihilist who believes only in un-extended simples will point to the need of both LH and Monism to posit fundamental distributional properties. Parsons (2004) has a persuasive independent argument for belief in fundamental distributional properties, but it relies on the premise that mereological gunk is possible— a premise likely to be rejected by any nihilist who thinks answers to SCQ are necessary, and rendered ineffectual by modal deflationists in the nihilist camp, such as Sider and Cameron. On the other hand, Cornell (MS.) argues that monism is ideologically parsimonious because it eliminates relations from one's ontology: at the very least, ontologists like the punctual nihilist and the believer in LH who believe in more than one thing are going to need spatial relations. While punctual nihilist and the monist can argue over which is the greater cost, it seems the proponent of LH must pay *both* spatial relations and distributional properties.

As I've argued in §5 and §6, the punctual nihilist is going to have to modify their ideology of property instantiation *in some way* to account for emergent properties: one option is to account for putatively emergent properties by positing plurally instantiated monadic properties in their place. Conversely, positing entanglement *relations* between simples can provide an ontological basis for constraints on what property distributions are

possible, consistent with quantum entanglement. However, a trouble for this strategy arises when another important feature of quantum entanglement was described above:³³ That at the given moment at which the component simples of a system are entangled with respect to some determinable F, they may not have any F properties at all, while the system itself does have F properties.³⁴ As such, unless the posited entanglement relations are going to - somehow - entirely ground the monadic qualitative properties of the system, the punctual nihilist is going to have to resort to something like collective instantiation. Either way, a substantive revision to the orthodox picture of fundamental qualitative properties being instantiated one-one by individual objects is required. The nihilist who believes in extended simples can go down the alternative route of positing extended simples with fundamental distributional properties.

In §5.4, I argued that the relaxing of the orthodox picture required to accommodate the plural instantiation strategy is not really an ideological cost at all - or not a significant one. I attempt to bolster this claim further, defending it against a potential objection drawing on the work of Cowling (2013), in §8.2. I also don't see any reason to suppose that relaxing the orthodox picture to allow for fundamental distributional properties is any *more* of an ideological cost than allowing for plural instantiation.

Sure, if we start out with, say, Red, Blue and Green, and then add fundamental distributional properties for each possible distribution in a 5x5 screen world, then this will require a proliferation of properties - arguably a vice by the measure of ideological or qualitative ontological parsimony. But this was all accounted for in our previous discussion. What's at issue *here* is whether allowing your theory to admit primitive distributional qualitative properties as a *kind* is a significant ideological cost. Here I find the issues harder to evaluate, but I also see little putative reason for considering there to be a significant cost. On the traditional view, for an entity to exhibit some given quality distribution is for it to either have or not have that quality at each of its parts. But why should the traditional restrictions on quality distributions be considered significantly more ideologically parsimonious than the more generalized model that emerges once we admit that heterogeneous distributions are possible?

I feel similarly about the monist's claim to economy via getting rid of fundamental relations. We're not talking here about getting rid of individual relations, but the *category*

³³ Recall §6.4.

³⁴ Or perhaps, as I prefer to see it, the simples have no *determinate* F properties, while the system has determinate F properties.

of relations. However, I'm sceptical that there's any economy to be had, once you have the notion of a property/relation, by placing restrictions on its adicity – say, having only three placed relations and lower rather than four placed relations, five placed relations, etc. Even in the seemingly special case of restriction to a one placed property/relation, I'm not sure what there is to be gained. Nonetheless, I think there's more intuitive force to the monist's claim to gain something important by eliminating relations than the punctual nihilist's rejection of distributional properties: Yes, the restriction means that our picture of reality is much simplified once we've done away independent things standing in relations to one another in favour of things (or rather *thing*) that just have qualities and that's that. The goal of explaining away all apparent relational phenomena using monadic properties is, I think, an attractive and powerful one. But I'm inclined to see any fulfilment of this ambition as representing an explanatory or conceptual victory for those who are distrustful of fundamental relations, rather than ideological economy as such. To the extent that the pluralist's commitment to relations as a general category – rather than to individual types of relations – *is* seen as a cost, it should be weighed up with all the other simplicity considerations on either side. However, I remain doubtful that this particular consideration will conclusively tip the scales.

7.8. Sider-Lewis Objection and Merging

A final objection I'll consider: That a merging model of LH succumbs to a version of the Sider-Lewis argument, usually directed against restricted composition. According to the Sider-Lewis argument³⁵ composition must be unrestricted or never happen at all, since any restricted set of criteria for when composition occurs would make it a vague matter when composition occurs, but this would lead to vagueness in the unrestricted existential quantifier (since it will be vague whether a given composite object is in the domain of the unrestricted existential quantifier), but logical terms such as the existential quantifier cannot be vague.³⁶ Analogously, you might worry that any criteria for when merging occurs will be vague, but this will lead to vagueness in the unrestricted quantifier since it will be indeterminate whether we are still quantifying over a plurality of things or over a single new merged thing.

Sider's (2001) construction of the argument has three premises:

P1. [If composition is restricted], then there must be a pair of cases connected by a

³⁵ Formulated in Lewis (1986*b*), later refined in Sider (2001).

³⁶ See Sider (2001, §4.9).

continuous series such that in one, composition occurs, but in the other, composition does not occur. (p.123)

P2. In no continuous series is there a sharp cut-off in whether composition occurs. (p.124)

P3. In any case of composition, either composition [determinately] occurs, or composition [determinately] does not occur. (p.125)

By P1 + P3 there must a sharp (determinate) cut-off in a continuous series when composition occurs, if composition is restricted, contrary to P2 (see p.125). I shall give some reasons to doubt each in turn for the specific case of merging proposed. Against the first premise, we should immediately note that once we move to a quantum setting it can't be assumed that it will be possible to construct a *continuous* series of cases connecting a case of non-composition (or pre-merging) with a case of composition (or merging). This is for the simple reason that not all physical variables are treated as continuous in QM— it allows for both for continuous and discrete quantities. So, we might still create a continuous series of particles getting closer spatially, yet there will be examples of things that vary in nature by fundamentally discrete jumps: For instance, electrons must occupy one of a discrete number of electron shells around an atom due to the exclusion principle. As such, if the physical quantities that our criterion for composition/merging are parasitic on are discrete rather than continuous, there will be no continuous series of extremely similar cases to be constructed: So there could be principled reason to choose between adjacent cases - to make one a case of composition/merging and the other not - since they may vary in significant and important ways.

But suppose our criterion does rely on continuous variables, and so P1 stands— there are still a number of reason one might have for doubting P2. I'm inclined to reject P2 because I can think of several ways of interpreting and justifying it, but none of them make remotely plausible the blanket rejection of a non-vague cut-off in composition/merging, even though they may serve to make *some* such cut-offs implausible. In particular, I'm inclined to think that sharp cut-offs are objectionably arbitrary if - but only if - motivated by an anthropocentric need to recover our ordinary conceptual scheme:

Say we want to make it the case that our answer to SCQ rules in familiar objects like my desk, and my shoe, but rules out weird objects like my-desk-plus-my-shoe. Consequently, we find some criterion of bonding or fastening, or what have you, that rules in the parts of my desk composing a desk and the parts of my shoes composing a shoe, while ruling out that the desk and the shoe compose some further thing (a desk-shoe)

despite my resting my foot on the desk as I type. The snag is, even if (contra van Inwagen's worries) we found a criterion for fastening, or what have you, that fitted our intuitive scheme, that whatever macro-features are relevant to us deciding that the *xx* compose or don't compose are going to supervene on some (set of) continuous microphysical variable(s). As far as everything on the micro-level is concerned, there's just no point in the series that gives one a principled reason to say *here* composition occurs, but wouldn't occur if you varied the conditions infinitesimally. Here it looks like coarse grained concerns about what macroscopic objects there should be are driving you to posit distinctions in the micro-realm where there is no principled difference: Thus, Sider's worry is that the restrictivist is committing herself to a macro-realm that, while still respecting micro-supervenience, is somehow 'autonomous' (p.124) from the microphysical facts in an objectionably 'arbitrary' (p.124) way.

But my proposal is that things merge if and only if they form an entangled system. Things differ from the desk-shoe case here in a few important ways: Firstly, macro-micro supervenience is already *ex hypothesi* violated, since we're assuming that the properties of entangled systems don't supervene on the properties of their (putative) parts, so the 'floating free' boat has already sailed in a much more egregious fashion: At some point, it just becomes the case that systems cease to obey micro-supervenience— if we're already prepared to give up on micro-supervenience generally, is it any great cost to further stipulate that the conditions under which this failure happens don't track any sharp distinction at the micro-level?³⁷ Isn't it, rather, just a furtherance of the theme that happenings at the micro-level do not determine the state of the world as a whole?

Secondly, my condition is not driven by anthropocentric macro-concerns. The concerns are neither anthropocentric *nor* primarily concerned with the macro-level. They are not anthropocentric because the condition for when things merge is taken directly from fundamental physics: things merge, metaphysically, if and only if they form an entangled system according to the physics. This is not an arbitrary or gerrymandered criterion driven by an *ad hoc* desire to recover our ordinary conceptual scheme, or any other preconceived idea about what mid-sized objects should exist. What's more: at the risk of stating the obvious, merging is not composition. On this view, there is no macro/micro distinction *as such...* Yes, pointy systems merge into extended systems...

³⁷ This line is closely analogous to that taken in Merricks (2005), which Barnes (2007) argues is dialectally effective against the proponent of the Sider-Lewis argument in normal conditions – but in a context in which micro-macro supervenience has been given up by *both sides* from the outset, I think the outlook is more promising.

extended systems into bigger systems. But, really, there is only one ontological ‘level’ as it were. What we have is a criterion, from physics, telling us when things stop behaving independently from one another, and thus form a new object. It’s an empirically motivated principle concerning the behaviour of things at the bottom (only) level of reality, and so is not subject to the concern that there are arbitrary or under-motivated *inter*-level principles.

Thirdly, Sider himself concedes that the sharp cut-off premise could be resisted if a restriction were placed on composition to *terminal cases* in a series, such as that composite objects can only occupy ‘regions in which any two points are connectable by some continuous path confined to the region’ (p.124). Call this restriction *absolute contact*. Sider objects to absolute contact by pointing out that, at least under a classical conception of matter, all macroscopic physical objects are spatially discontinuous (they’re made up of dispersed arrangements of particles). So, unlike intuitive contact views,³⁸ absolute contact would not allow for the existence of tables or even molecules,³⁹ let alone more obviously discontinuous objects such as galaxies. But this should only be a problem if our concern is to recover our everyday intuitions concerning what composite objects exist— and as I’ve argued extensively this should *not* be our concern: our criterion is motivated by fundamental physics. Sider further says that ‘a classical world should not turn out devoid of macro-objects’ (p.124), but this only holds good if a nihilist strategy fails to account for ordinary classical objects, and I’ve argued there is every reason to think that such an account is adequate. Indeed, the main focus of Sider’s argument is on showing that there can be no sensible *intuitive* restrictions on composition, which makes it hard to evaluate how seriously we should take its conclusion when proposing a theoretically motivated restriction that’s non-intuitive but independently motivated. In any case, if our criterion is analogous to contact in being a terminal case – and we might think that systems become entangled, very roughly, when they come into causal (rather than spatial) contact with each other – then we have another potential way of resting P2.

Finally, if none of the above has moved you, perhaps the most straightforward way of rejecting the argument is by throwing out P3. Sider rejects that there could be semantic vagueness in the quantifier to motivate P3, but does not take the possibility of worldly vagueness in whether composition occurs seriously. The Barnes-Williams (2011)

³⁸ i.e. views on which things compose under the conditions which we would ordinarily describe them as touching.

³⁹ To pronounce on the fate of objects smaller than this would require a consistent classical description of the atom.

framework for indeterminacy gives us a straightforward way of modelling vagueness in the world— while Barnes (2010) shows us specifically how we can respond to the Sider-Lewis argument by employing metaphysical indeterminacy with respect to the quantifier.⁴⁰

To conclude, then - having seen off the Sider-Lewis and ideological economy objections - the merging model (a subspecies of the extended simples account of nihilistic emergence defended in §6.1) provides a coherent, attractive account of nihilistic LH. Nihilist versions of LH were themselves found to be preferable to non-nihilist versions on defeasible grounds of ideological economy. LH itself, a form of pluralistic nihilism which takes into account localized instances of micro-macro supervenience, was, given certain assumptions about quantum theory, found to provide a more elegant account of micro-macro correlations in the world, compared to the traditional pluralism of ML or the Priority Monism of MH. We also saw that, even granted all Schaffer's assumptions concerning QM, there's reason to doubt monism provides a particularly elegant model of an entangled-yet-decoherent world. All this took place under the assumption that a 'priority' metaontology is attractive, for explaining ontological dependence via fundamentality: If one is prepared to break the link between dependence and fundamentality, as discussed in §5.2, a wider space of potential options opens up.

⁴⁰ Cornell (MS.) objects that even if we employ ontic vagueness to smooth-over sharp cut-offs, we can iterate the problem by demanding that there be no sharp cut-off between the determinate cases and the vague cases: What you make of this argument will depend pretty much entirely on what you make of the coherence of higher-order vagueness worries in general. It's true that the Barnes-Williams model can't accommodate higher-order metaphysical indeterminacy in its canonical form; but nor would they be likely to countenance such higher-order worries. If you do think the no cut-off worry has teeth when iterated in this way to higher-order concerns, you'll need a model of metaphysical indeterminacy that accommodates higher-order unsettledness in the world. As I've said, I'm inclined to doubt the bite of P2 even at the first-order, for our purposes: I'm even less inclined to accept that there can't be a cut-off between determinate and indeterminate cases.

8. Nihilism and Parsimony

So, I've left what I take to be some of the most tendentious remarks of this thesis until last. Not a prudent course for those concerned with lasting impressions, perhaps, but there is a smidgen of method in my madness: Discussing extra-empirical theoretical choice virtues is hard. Applying such discussions to highly abstract debates in metaphysics is an order of magnitude harder. Harder even than common or garden metaontology—and I think that's really saying something. And I'm not harping on about this just to be whiny, or to try and garner sympathy— if you've waded through all the pages above to come this far, Dear Reader, that boat has surely sailed. No, rather: You will recall that my approach to the metaontological issues in *Part I* was to try and gain traction on them through the process of doing first order ontology— through trying to defend the attractiveness of nihilism as an answer to SCQ. I think this was helpful for keeping us grounded— that is, for example, for helping keep track of the important issues and reminding us what was at stake; for getting to grips with key concepts or debates; for allowing us to carve up dialectical space into more manageable chunks, to facilitate modest progress on piecemeal problems.

I want to take more or less the same approach here, when it comes to discussing parsimony, the theoretical virtue (or close family of theoretical virtues) which I have relied on most in my defence of nihilism. This time, however, rather than attacking the issues head on - as I did with the metaontology debate - I just want to abstract some of my key assumptions that I've made in previous chapters relating to theory choice and see if I can't say anything vaguely persuasive or original in their defence. What I'm not offering is even close to a complete account of how or why parsimony matters, either generally or in metaphysics, but I do want to draw some connections between different debates in this thesis, and relevant discussion elsewhere - which may go some way to convincing you that if simplicity matters at all as a theoretical virtue, it matters in the composition debate.

The least controversial way in which nihilism is claimed to be more parsimonious than its rivals is the claim that nihilism is *ideologically* or *qualitatively ontologically* more parsimonious than its rivals. That desirability of this is pretty widely held, in that several

philosophers have followed David Lewis¹ in endorsing this form of parsimony in the process of specifically disavowing *quantitative* parsimony. I think it's reasonably clear – as much as anything is in the murky territory of theoretical virtues – that nihilism does well by the lights of such a principle. I'll defend these claims further in §8.3. However, defending such a principle in terms a sceptic about parsimony would accept is notoriously difficult and I have little if anything new to add on this front. Nor do I have a non-question begging response to the theorist who is more generally sceptical about whether the sorts of extra-empirical virtues employed by metaphysicians track truth.²

Perhaps the most uncontroversial formulation of the razor, however, is that one should not posit entities that are *strictly redundant* by the lights of the theory itself. That is, not just that a different theory could explain the target phenomenon using fewer entities but that some of the entities posited do no explanatory work even within the first theory. While relatively uncontroversial that falling foul of such an *anti-redundancy principle* is a bad-making feature of a theory, it is hard to put such a principle to much use, since most theorists have internalised such a rule to the extent that they're careful to avoid positing entirely redundant entities in any case.

However, I will also defend the more controversial claim that nihilism's *quantitatively* ontological parsimony over its rivals is *also* an advantage. That is, I claim that we should prefer nihilism because it posits fewer *individuals* than its rivals, as distinct from the above claim that it posits fewer kinds of things (where its assumed for sake of argument that all the entities in question play a non-redundant explanatory role). While I'm inclined to believe that quantitative ontological parsimony is a virtue in itself, many philosophers have explicitly reject this, and I don't have much to offer in the way of fresh argument to the contrary. Instead, I'll be trying to show in §8.3 that the attractiveness of quantitative parsimony can be explained in terms of other explanatory virtues.

8.1. Kinds, Primitives, Theories and Cosmologies

The Divers-Melia debate has significant implications for the current discussion. Melia (1992) criticises Lewis' (1986*b*) Genuine Modal Realism (GMR) for being uneconomical with respect to qualitative ontological parsimony (qual-OP), since it contains many kinds of things that actualist ontologies would say don't exist. Divers (2002, p.155) replies by

¹ Lewis (1973, p.87).

² See, e.g., Bueno & Shalkowski (2014). Though for a forceful defence of the view that theoretical virtues in metaphysics track truth see Paul (2012*b*).

suggesting that:

GMR proports to be a comprehensive theory of ontology, and in that light it is plausible that the relevant measure of the quantitative economy of GMR ought to be the number of ultimate, *sui generis*, ontological kinds that it takes to be instantiated. [...] In matters of fundamental ontology, individuals form one ultimate (instantiated) kind and the relevant question is how many other (non-overlapping) ultimate kinds of entity are postulated.

So, according to Divers, the relevant kinds that count for measuring qual-OP are *categories*: individuals, sets, universals, states of affairs, and so forth (see p.155). How many entities you posit *within* such kinds has no bearing on the question of economy in fundamental metaphysics. If this exhausted the requirements of parsimony, this would have consequences in the context of the composition debate: Is being a composite thing to be a fundamentally different *kind* of thing from a simple thing? Plausibly not, you might think: a composite thing is just another concrete individual that happens to have smaller concrete individuals as parts.

However, I think there's a widely accepted, if not often articulated, line of thought on which the Diver's criterion is a little *too* general, even if we're sympathetic to the line of thought which says that the metaphysician shouldn't concern herself with *all* kind distinctions that one could make between individuals. The middle ground is that kind distinctions that are marked by sparse properties also count against qual-OP: having regular horses and horned-horses in your ontology might be no sin against qual-OP, but - whether we're marking the having of a sparse property nominalistically with sets or plantonistically with universals - adding a new sparse property to your ontology, such as charge or spin, perhaps, is a sin against qual-OP.

Why do these kinds of things matter for parsimony? To see why, consider the following. In fundamental physics there is the Standard Model, which is concerned with the kinds of particles that exist and the universally applicable dynamics that govern their interactions, but does not tell us anything directly about the present state of the world or its history. By contrast, physical cosmology tells us what the early universe was like, how it has evolved over time, and very generally what things we're likely to find in the universe now and where. Drawing an analogy, a lot of metaphysical theorising seems to be a mix of these sorts of two broad enterprises. For instance, in Lewis' case he gives a reductive analysis of possibility in terms of concrete possible worlds - which I take to be akin to the physicist giving a unifying analysis of, say, the dynamic forces of electromagnetism and the weak force in terms of the electro-weak force. Lewis tells us that there is no primitive possibility in the world; the physicist that there is no fundamental electromagnetism.

Meanwhile, however, Lewis also gives us a theory of the extent of possibility by

pronouncing on what possible worlds are out there:³ This is as if – if you will – he’s giving a metaphysical cosmology: a theory of what particular things there are in the world, what they’re like, how they’re arranged, etc. Our general theory will have already sorted (we hope) how many fundamental *sui generis* kinds of things there are. The task here is to lay down some doctrine⁴ positing instances of these things – and general patterns in and relationships between these instances – in the service of explaining relevant features of the world we see around us.

Now consider the sorts of kinds that are playing the role of ultimately explaining the target phenomenon. Each theoretical kind fills a given theoretical role in our theory: for general metaphysical theorising these roles take kinds such as individuals or universals or states-of-affairs, but when giving a metaphysical cosmology the theoretical roles to be filled require kinds that are much more fine-grained. In either case, I believe that parsimony concerns arise from a general concern that we should avoid complicating our explanations by multiplying *explanatory* kinds beyond necessity. And a lot of the time when doing ontology, such as in the composition debate or the priority debate, we find ourselves giving both general and cosmological theories— so we should be careful both about the multiplying of coarse grained kinds and fine-grained kinds. So, if the nihilist is right and parthood is not fundamental then even our cosmologies won’t count ‘composite’ as a theoretical kind, but if the universalist is right then the world furnishes the universalist with the tools to make a distinction between composites and individuals that she then goes on to employ in her cosmological explanations (where the nihilist refrains).

In §8.3 I’ll try to argue that when in the business of giving something along the lines of a metaphysical cosmology, even quantitative parsimony should matter: that is, how many individuals you posit *within a kind*, rather than just the number of kinds one postulates, loosely construed.

8.2. Ideological Parsimony

Sider’s (2013) master argument against belief in composite objects is that admitting the ideological primitive of *parthood* complicates your theory’s fundamental ideology. That belief in composition requires that you complicate your fundamental ideology with belief in a parthood primitive, or that it requires a *qualitative* complication of your fundamental ontology through belief in a parthood relation, is the most straightforward parsimony

³ See Cameron (2012, pp.10-11).

⁴ See Sider (2011, §9.2).

argument for nihilism. We should, I think, treat the qualitative ontological cost of belief in a fundamental relation of parthood as exactly commensurate with the ideological cost of taking parthood as a bit of primitive ideology. Notice that this does not require taking any stand on whether composite objects are of a fundamentally different kind to simple objects in the sense that they have some monadic property of being composite or simple: it's simply observed that parthood requires the posting of a new perfectly natural relation or a complication of our fundamental ideology. Since our best theories should aim to avoid such additional commitments, all else being equal we should prefer nihilism on this basis.

However, Cowling (2013) argues that nihilism offers no gain to ideological parsimony because a plausible principle of ideological parsimony seeks to minimise *kinds* of ideological primitives rather than individual ideological primitives, and he takes parthood to be an instance of a kind of ideological primitive that the nihilist remains committed to. Specifically, Cowling thinks that parthood and identity fall under the same ideological kind, so there is no ideological saving to be had by ditching parthood but keep identity—and, he says, any plausible theory needs primitive identity.⁵

This is similar style of argument to one I push elsewhere in the thesis.⁶ However, while I find Cowling's initial set-up and motivation of the quantitative/qualitative distinction for ideology plausible, I think the resulting principle is too conservative to get him the claim he wants about mereology— that is, unless he's happy to endorse a stronger version of Composition as Identity (CAI)⁷ than he appears to commit to in the paper. In summary, I think that the argument is more plausible the closer one takes the family resemblance between composition and identity to be: For super strong CAI on which composition and identity are the very same primitive, you don't even need the qualitative/quantitative distinction for the basic argument to go through: you can't gain anything by eliminating parthood from your theory and keeping identity, since to eliminate parthood *is* to eliminate identity. Now weaken CAI to the claim that parthood and identity are very intimately related primitives: there's some plausibility to the particular argument that

⁵ Contrast with the puzzling argument of Cowling (*forthcoming*) which says that provided we're happy to go further than Sider (A-Parthood) and eliminate *improper parthood* from our ideology as well as *proper parthood*, and hold that while there are particular concrete things they are neither simple nor composite, we have the makings of a plausible and ideologically parsimonious theory.

⁶ Both defending the ideological innocence of plural instantiation §5.4 and in §7.7 discussing the ideological commitments of my version of LH.

⁷ See, e.g. Baxter (1988a, 1988b); Wallace (*draft*, 2011a, 2011b).

Cowling runs, but this requires substantive commitment to some form of CAI. Taken as the claim that the argument for nihilism from ideological parsimony doesn't go through even if interesting versions of CAI are false and parthood shares at best a loose family resemblance with identity, I find the argument to have almost no plausibility at all.

Cowling initially motivates his case by appealing to the obvious innocence of equivalent or interdefinable bits of ideology: admitting a diamond operator into one's language, is not an *extra* cost over and above a primitive box operator. So *just counting* ideology in your theory isn't the way to go when weighing up ideological costs. Cowling's gloss on this: possibility is of the same ideological kind - the modal kind - as necessity, so once you have primitive modality of any sort all the rest comes for free. However, everyone's going to agree that ideological parsimony isn't simply a matter of counting ideology,⁸ but Cowling needs more to establish that more loosely related⁹ ideology such as parthood and identity should be counted together. Enter his appeal to the theoretical work that a distinctly *qualitative* principle of ideological parsimony could do for us were we to accept it.

Cowling seeks to wield his principle of qualitative ideological parsimony not so much as an Ockhamite razor, cutting expressively extravagant theories down to size, but as a Moorean cudgel, knocking on the head arguments for 'counterintuitive metaphysical theories'. I have worries both about the general strategy and the particular examples given. The general strategy is suspect since the principle is justified by its ability to undercut arguments for 'counterintuitive' theories, while the particular parameters of the principle - how similar two ideological notions need to be to be counted as of *the same kind* - seems to be set *ad hoc* at whatever threshold is needed to cause problems for the target theories. Furthermore, it is assumed that merely being 'counterintuitive' is a bad-making feature of a metaphysical theory: My suspicion is that a principle employed to bolster our otherwise baseless Moorean prejudices is being justified simply by its effectiveness in doing just that. Leaving aside methodological concerns, I don't think any of the three cases Cowling uses to support the principle work in the way Cowling needs them to. Nihilism is the case at issue between us, so I'll leave discussion of that til last. In the first two cases, Growing Block Theory and Genuine Modal Realism, Cowling is just wrong in his assessment that ideological parsimony - or certainly *qualitative* ontological parsimony - has any role to

⁸ E.g. Goodman (1951), Sider (2011).

⁹ On the *interesting* reading of his argument on which Cowling isn't simply assuming or begging the question in favour of CAI.

play in elucidating the suspicions about these theories that he tries to point to.

With respect to the first, Cowling says that the Growing Block theory of time is widely agreed to be strictly worse than either Presentism or Eternalism. His idea is that this is because the Growing Block takes on both the quantitative ontological bloat of Eternalism (past and future ontology) and the ideological bloat of Presentism (primitive WAS and WILL sentence operators). However, it seems the Growing Block theory only gets encumbered with half the extra qualitative ontological commitments of Eternalism (past ontology) and half the quantitative ideological commitment of Presentism (the WILL operator). Hence, the cudgel needs to be wielded to pummel the putative ideological costs into the desired shape: If the WAS and WILL operators are decreed to fall under the same ideological *kind* then Growing Block theorist gains nothing by doing away with the WAS operator, and so ends up with effectively the same ideological commitments of the Presentist *plus* half the extra ontological commitments of the Eternalist. A strictly less parsimonious theory.

Here is the argument in Cowling's (2013) words (p.3901-2):

I believe the proper explanation of the inadequacy of the growing block view [of time] is available only given [qual-IP]. To begin, let us suppose that presentism and eternalism are equivalent with respect to [overall parsimony] yet differ markedly with respect to [ideological parsimony] and [ontological parsimony]. While eternalism requires ontological commitment to both past and future entities, presentism takes on the ideological cost of primitive WAS and WILL operators. Given a quantitative conception of ideological parsimony, the growing block view is therefore equivalent in cost to either presentism or eternalism, since it accepts what we might think of as half the cost of each of presentism and eternalism. So understood, the case against the growing block view must turn on relatively small concerns about the cost of foregoing a purely general account of the past and future. But, if these concerns are indeed this small, the standard assessment of the growing block view [as obviously inferior to rival reviews] is puzzling. [...] In contrast to the quantitative conception of ideological parsimony, the qualitative conception of parsimony guarantees that the growing block view is more costly than either presentism or eternalism. On the one hand, since the WAS and WILL operators of the presentist are of the same ideological kind, the growing block theorist's commitment to a primitive WILL operator is at least equal to the total ideological cost of presentism. On the other hand, the growing block theorist also takes on the ontological cost of quantifying over past entities and is thereby guaranteed to bear a greater theoretical cost than either presentism. And, while the precise ontological cost of quantifying over past entities depends on one's preferred view of ontological parsimony, once coupled with the ideological expense, the growing block view is also assured to be more costly than eternalism. For this reason, [qual-IP] provides a plausible and intuitive

diagnosis of why the growing block view is reasonably dismissed as a alternative to presentism and eternalism.

However, it's absurd to complain in the first place that the eternalist is committed to a different *kind* of ontology, past and future ontology, when on the eternalist's theory the only difference between existing now, in the past or in the future is *indexical*. Only by accepting the presentist's ideology of a primitive present does this claim have any plausibility in the first place. So, if there's a dialectical stand-off between the Eternalist and the Presentist, it's not a straightforward case of a trade-off between qualitative ontological parsimony and ideological parsimony. Indeed, it's more reminiscent of the Genuine Modal Realist's trade off of *quantitative* ontological parsimony for ideological parsimony, except it seems much harder to make the Melia (1992) move that the eternalist is committed to alien natural kinds: we tend to assume more or less the same natural kinds existed in the past as exist today and will exist in the future.

You could try to save Cowling's argument with an explicitly *quantitative* form of the ontological parsimony principle, but *a)* even though I'm inclined to think quantitative ontological parsimony is a virtue myself, it's much more controversial than its qualitative counterpart, and so a shaky basis from which to legislate on other theory choice principles, *b)* Cowling makes the point that belief in a principle of ideological parsimony is needed to balance out a tendency towards always going with the most ontologically austere theory over an ideologically austere one if we accept a principle of ontological parsimony (and *vice versa*). But, it seems methodologically haphazard to posit a theory of ideological parsimony that makes quantitative economy entirely subservient to *qualitative* economy, while saying that this can be traded off on an equivalent basis with *quantitative* ontological parsimony.

Of course, some growing block theorists¹⁰ have been lead to metaphysically privilege objects at the edge of the block from those further back in the block – saying that only present things have causal powers or consciousness, for instance. So, perhaps we have the start of a competing explanation of the badness of growing block in terms of parsimony: however one might go about justifying the ideological commitments of presentism, growing block comes with these ideological commitments *plus* qualitative ontological commitments— that is, commitment to both past ontology and metaphysically privileged present ontology. This argument would, all else equal, make presentism the leading contender to eternalism, rather than putting both alternatives on an equal footing. But this

¹⁰ Such as Forrest (2004).

argument does not hinge on how we cash out the nature or severity of presentism's ideological commitments.

With respect to the debate on genuine modal realism, Cowling argues that a key motivation for GMR – that it gains ideological parsimony through giving a reductionist account of modal notions – is undercut once we move to a qualitative conception of ideological parsimony. This is because, he claims, the Lewisian reductionist account of actuality is doomed: actuality can't be defined just as everything in the pluriverse that's spatiotemporally related to us, because island universes are possible, which contain a plurality of spatiotemporally disconnected bubbles. If an indexical spatiotemporal account of actuality is out, then the only plausible alternative is having primitive actuality in one's ideology. Under a quantitative conception of ideology, the modal realist might be able to argue that while they're still committed to one modal primitive, actuality, they've still economized by giving a reductive account of necessity and possibility. But once the ideology in question is put under the cosh of Cowling's qualitative principle, it's clear that as actuality, necessity and possibility are all modal notions, you don't save anything by getting rid of the others if you keep one.

Again, in Cowling's words (pp.3904-5):

According to Lewis, actuality is an indexical rather than absolute matter, so worlds and their occupants are actual or non-actual only relative to other worlds or individuals. And, while this account handily resolves an epistemic worry regarding how we might know ourselves to be actual rather than merely possible, it precludes a satisfactory view about the plenitude of possible worlds. In particular, there is ample reason to believe that there could have been a world of island universes (i.e., wholly disconnected spacetimes); however, given Lewisianmodal realism, worlds are isolated or unified by their spatiotemporal relations. Island universes, which are not spatiotemporally unified, are therefore ruled to be impossible. Assuming, for present purposes, that island universes are indeed possible, the Lewisian view of actuality must therefore be rejected. In its place, the leading alternative is absolutism about actuality, which takes actuality to be an irreducible property, which could be instantiated by a plurality of disconnected spacetimes. (In addition, the standard analysis of modality that proceeds using singular quantifiers must be amended to appeal to irreducibly plural quantifiers over co-actual worlds.) Absolute actuality is therefore required to accommodate the possibility of island universes within the modal realist framework. Since a serviceable version of modal realism must view actuality as an irreducible, unanalyzable property, it also comes at an ideological cost. Furthermore, it directly undermines the Ambitious Thesis, since modal realism offers no advantage in ideological parsimony over any view that already accepts primitive actuality.

However, no self-respecting reductionist project should be happy with a dialectical

move that forces her to accept primitive actuality. That's because even the biggest sceptic concerning Cowling's more wide-ranging principle of qualitative ideological parsimony is likely to take it as obvious that actuality, possibility and necessity are a closed circle of concepts which should be taken together if at all. If one is attracted to GMR on the basis of ideological simplicity, it's rather obvious that accepting primitive actuality is going to undercut these motivations (regardless of whether someone attracted to GMR for independent reasons might in principle be happy with the move). Hence, the parsimony conscious proponent of GMR is simply going to get off the boat at a different step in Cowling's rather rushed dialectic. Either they'll reject the possibility of island universes or they'll give a different *reductionist* account of actuality, presumably in terms of some other non-modal relation. To the first option: The putative possibility of island universes just isn't a decisive objection to GMR, absent extensive further argument. To the second: The project of finding an alternative reductionist account of actuality can't be sweepingly pre-judged. So there's no sense in which the most 'plausible' versions of GMR are committed to the ideological bloat Cowling alleges. And yet, there will still be a lingering thought among many objectors that, contra Lewis, GMR is unparsimonious— if this is to be made sense of at all, it clearly can't be in terms of qualitative ideological parsimony.

So, we come to an assessment of nihilism's ideology without any prior support for the theoretical utility of the qualitative ideological parsimony distinction as Cowling would have it applied, beyond rather tame observations of the sort that no one wants to pay extra to employ a box operator in their theory if they're already committed to a diamond operator, and that debates concerning whether to *either* take the box operator *or* the diamond operator as primitive have the smell of non-substantivity. But the box and the diamond operators are two sides of the same coin: they are interdefinable. You're going to need more than that to argue that one primitive notion that's similar yet independent of another might turn out to be so similar that the former is no ideological commitment over and above the latter.

However, despite Cowling's inability to spell out or defend such, I do think there's a plausible and interesting way, besides box and diamond type cases, in which an ideological notion can be no (significant) cost above and above another, closely related, ideological notion. This is where we start with established ideological primitive and then augment our theory with a new notion that does basically the same work as the old primitive but, say, involving objects from different ontological categories or involving a different number of objects or, more generally, allows that primitive to apply in different situations or in a structurally different way. You may well then be in a situation to posit a new, more general, primitive that subsumes the original two - one that, from your new

dialectical position can be cast as removing restrictions on the old primitive – or it may be that the old predicate can be defined in terms of the new more powerful or applicable primitive, such that it takes on the role of the general, subsuming primitive. In any case, this is obviously importantly different from the case of two *interdefinable* primitives and I take it we're on more controversial ground. This is because, despite the hope of being able to get rid of one primitive by defining it terms of the other being present in both cases, adding the new primitive increases the overall expressiveness of the theory, even if it ultimately subsumes, or is subsumed by, a different primitive of the theory.

Perhaps the best example to discuss, especially in the given context, is when you start out with a primitive that works in the manner of a one-one function and then add a primitive that does the same theoretical work, except in the manner of one-many function. Take the example of plural instantiation employed in §5. Here we start out with instantiation, which is standardly construed as holding one-one between an object and a property, and introduce plural instantiation, which is taken to relate a plurality of objects to a property in the same way that ordinary instantiation relates a single object to a property. To pick out the things that are to instantiate a given plural property, we similarly need to allow the existential quantifier to quantify not just over individual objects but also over pluralities: on the face of it, we complicate our ideology by adding a plural quantifier, but note that the old singular quantifier can be defined in terms of the new plural quantifier. Now, even on a *quantitative* conception of ideological parsimony, we get the result that there's no cost here, because of the interdefinability of the old primitive in terms of the new one— which suggests that the qualitative/quantitative distinction isn't quite getting to the heart of what's going on here. It's too quick to conclude that because it doesn't fall foul of a flatfooted quantitative principle of ideological parsimony – and therefore, for sure, Cowling's qualitative principle – that there's no cost here. On the face of it, there *is* a complication in ideology because less expressive quantificational vocabulary is being replaced by more expressive quantificational vocabulary. However, this sort of increase in expressive power that just allows quantification to be done over more individuals, rather than introducing distinctively non-quantificational machinery, is I think *analogous* to a quantitative cost in the ontological case: No cost at all for those on Cowling's side of the qualitative/quantitative debate. In my view, still a cost – in both the ideological and ontological cases – but not of the same order. But nothing much here should hang on this difference.

A similar case: Is it a complication of our ontology to allow for variably polyadic relations? For instance, take the debate between Karen Bennett (2011) and L. A. Paul (2012a), on the one hand, who argue broadly that composition or 'building relations'

should be treated as a unified single relation across widely different theoretical roles and categories, despite holding between different numbers and kinds of relata in each case. On the other hand, Kris McDaniel (2010) argues that the composition relation should be considered a fundamentally different relation, with a different adicity, when employed for certain theoretical purposes compared to others: specifically, with spatial parthood holding between regions and sub-regions, while concrete parthood holds between a concrete whole, a part and a spatial region. The two sorts of views, while not straightforwardly incompatible, suggest two radically different pictures of the ideology of composition: One on which we have multiple primitives, each dealing with a fixed number of entities and entities in specific ontological categories; another on which we just have the one primitive that can deal with varying number of entities across all relevant ontological categories. One thing which would look very bad to say is that we should simply dismiss the McDaniel picture out of hand, since going down this route will commit us to more than one parthood primitive. Here the quantitative/qualitative distinction could be put to work in explaining the intuition that whatever substantive issues may (or may not) be at stake between the Bennett-Paul picture and the McDaniel picture, that one does not have the advantage over the other on grounds of parsimony. Furthermore, though, by accepting that the posited composition relations in McDaniel belong to the same close-knit family of notions, we see that McDaniel doesn't have to pay a significant ideological cost for belief in both spatial composition *and* concrete composition compared to someone who *just* wanted to believe spatial composition.

With this background in place, what to make of Cowling's claim that parthood is sufficiently similar to identity to be no cost over and above identity? Can we make sense of the idea that composition is just like identity, only differing in the number or types of entities it concerns? Well, yes, easily enough, in that we can take the composition to be the very same relation as identity, except holding between an individual and an (identical) plurality, rather than between an individual and an (identical) individual. But this is composition as identity in a strong form: the highly committal and controversial thesis that the whole is the very same thing as its parts. Sure, if we are prepared to assent to *that* then there's a strong claim that composition is no cost over and above identity: let's not quibble overmuch about the extra expressive power given to the identity relation here. Yet Cowling seems to be presenting an argument that is neutral on strong CAI, and arguing from mere superficial similarities between parthood and identity that parthood is no cost over and above identity (2013, p.3906):

Like identity, it contributes nothing to the non-structural, qualitative character of the world, and, like identity, facts about its general nature seem to be a non-contingent matter. Furthermore, regardless of whether one endorses nihilism, classical extensional

mereology demands certain conceptual ties between these relations. Most notably, the uniqueness of composition precludes distinct entities being composed of the very same objects.

While I'm perfectly happy to accept that parthood, along with identity, deserves the status of a 'broadly logical relation', and would even go further and assent to the claim that parthood is an identity-*like* relation in a meaningful sense, the gestured-to similarities are much looser than those present in the strong CAI case and, in my opinion, just not enough to support the claim that parthood is no cost over and above identity. Identity is a relation that never holds between distinct entities, while composition is a relation between some entities and some *further*, distinct entity – however closely qualitatively related the distinct entity may be to the other entities. To reject this fundamental difference between composition and identity is to endorse strong CAI. There's no sense here in which a well understood identity relation is simply being given enhanced expressive power to take more or different kinds of entities into its fold: rather, we have two entirely different, incompatible ideological primitives that are noted to be similar in interesting ways, but which must ultimately be treated separately. As such, the qualitative ideological parsimony principle as Cowling presents it does not provide a 'novel interpretation' (p.3906) of CAI in any substantive sense: either it should be read as a plea for full-fledged CAI along the usual lines, or as merely drawing attention to certain similarities between composition and identity which have no doubt contributed to the motivation between weak forms of CAI in the past – either way, there is no new argument for nihilism.

8.3. Quantitative Parsimony

Further to the discussion of the previous section, I think some support can be provided for an anti-quantity principle of parsimony, in addition to the less controversial anti-redundancy principle, which the previous section relies on.¹¹ What's more, the principle at work will be one that targets not just a proliferation of different *kinds* of things but the proliferation of individuals: That is, a principle of quantitative ontological parsimony. While I think that minimising commitment to individuals is a theoretical virtue in its own right, many, following David Lewis, have been sceptical of this and producing convincing arguments against the sceptic is notoriously difficult. As such, here I have the more modest ambition of trying to show that quantitative parsimony is closely correlated with other explanatory virtues. My central contention will be that increasing the number of

¹¹ See E.C. Barnes 2000 for a discussion of the distinction and the relative plausibility of the different principles.

objects that a theory posits tends to lead to more necessary connections between objects, and such necessary connections are an explanatory vice.

More objects leads to more combinatorially generated *prima facie* possible states. As such, the total resources the theory provides – both ontological and ideological – would seem to allow for compatibility with a greater range of conceivable scenarios. But being compatible with a greater range of (epistemically) possible data isn't necessarily a good thing: You want your theory to be compatible with all the data that in fact gets observed (both up to this point and in the future) – i.e. you want empirical adequacy, and you probably want it above all else. But if your theory is compatible with possibilities that continue to not be observed indefinitely, things begin to look embarrassing for the theory.

If you're sympathetic to a broadly Popperian line, you may think theories compatible with a greater range of possibilities should be avoided, all else equal, since such theories are harder to disprove. The main worry I want to focus on here, though, is that if you start out being compatible with options that don't ever get observed, then you end up having to add brute, unexplanatory – potentially *ad hoc* – additions to the theory to explicitly rule out these possibilities in the theory. In the case of ontologically abundant theories, this comes in the form of unexplained necessary connections between entities.

Take Daniel Nolan's (1997) case of quantitative parsimony in which the hypothesis that one neutrino is admitted in beta decay competes against the two neutrino hypothesis, the seventeen-million neutrino hypothesis, and so forth. What makes the one neutrino hypothesis the better hypothesis? It's not that extra neutrinos would be redundant, since according to the seventeen-million neutrino hypothesis, each neutrino posited accounts for a seventeen-millionth of the missing spin. Yet, there still seems to be something right about preferring the one neutrino hypothesis – which was born out empirically.

Baker (2003) puts forward an analysis of the Nolan case that he thinks shows a principle of quantitative parsimony is 'demonstrably rational, for a certain characterizable class of case' (p.245). Baker suggests that more quantitatively parsimonious hypotheses tend to have more explanatory power. In the neutrino case, it is part of the background set up that no fractional spin values smaller than $\frac{1}{2}$ have been observed to be lost from an atom during beta decay— it's always $\frac{1}{2}$ spin that is missing, never $\frac{1}{4}$ or $\frac{1}{3}$ or $\frac{1}{17,000,000}$. While all the competing neutrino hypotheses are compatible with this data, it's not clear that all of them are supported to an equal extent. The one neutrino hypothesis only allows for spin to be emitted in units of $\frac{1}{2}$, assuming there is only one kind of neutrino with a fixed spin value, so there is no need to further explain why no fractional spin losses are ever observed once the one neutrino hypothesis is accepted. By

contrast, the two neutrino hypothesis either needs to find some explanation for the absence of $\frac{1}{4}$ spin losses: perhaps make it a brute fact that, or posit some deeper explanatory mechanism by which, neutrinos are only ever emitted in pairs. The seventeen million neutrino hypothesis needs to explain the absence of *a lot* of missing spin values: 1/17,000,000, 2/17,000,000, 3/17,000,000 and so on. Now, a similar explanation could be given as in the two neutrino case, but that nature co-ordinates things such that neutrinos always come in swarms of exactly 17,000,000 is, I think, a lot harder to swallow than the pairs hypothesis without some deeper story: Certainly, in any given case of beta-decay, we have to believe in many more *de re* connections between the particles involved. Perhaps such apparent connections could be explained away by a deeper story, but that's just to draw attention to the weightier explanatory burden the view faces. Baker's conclusion is that 'Quantitative parsimony tends to bring with it greater explanatory power. Less quantitatively parsimonious hypotheses can only match this power by adding auxiliary claims which decrease their syntactic simplicity.' (p.258) My focus here, however, is not on syntactic simplicity as such but specifically on the link between a liberal ontology and unexplanatory necessary connections— which I take to also be applicable to metaphysics.

In metaphysics, these connections will tend to hold between two classes of entities with two different *theoretical roles*, usually where it's not the case that the entities in the second class belong to a different *ontological category* than those in the first. Hence, it applies where qualitative ontological parsimony would not if you take Divers' line in the Divers-Melia dispute. Further, it also applies even where an infinite number of entities are already posited by a theory, since it advises against adding new entities into a theory in the context of a particular explanation, regardless of the theories over all commitments.¹² Ontological elegance is primarily a feature of metaphysical cosmologies, rather than our general theories, but having an economical metaphysical cosmology matters when defending one is (part of) the task in hand.

Consider the debate between the monist and the pluralist above. As I cashed it out, there were several things being employed as theory choice criteria in that debate:

- Ideological parsimony / qualitative ontological parsimony - positing fewer fundamental properties.
- Quantitative ontological parsimony - positing fewer concrete individuals
- Ontological elegance - positing fewer necessary connections

¹² Avoiding the worry to the contrary in Nolan (1997).

Where the primary theoretical role of concrete individuals is as bearers of qualitative fundamental properties, it's possible that 2 above may be explained in terms of (and perhaps fully or partially reduced to) 3. While 2 and 3 are clearly not equivalent, 2 may work pragmatically as a guide to theory choice if achieving 2 is more often happens to result in achieving 3.

Think about the pluralist trying to account for Chessboard. She has to posit a veritable web of necessary connections between her pixels to make the statespace come out right—to make sure neighbours are off if a pixel is on and vice versa. Why? It's all down to having so many bearers of fundamental properties: By positing things with the primary theoretical role of being the bearers of fundamental qualitative properties, when such an x can't instantiate some F in circumstance C the theory needs to say so: that is, it needs to do more than just posit the entities, the properties and a free recombination principle, but has to make case-by-case determinations about what are the admissible instantiation patterns. The more entities that are posited, the more combinatorial possibilities are opened up, so the greater the chance that *ad hoc* dependencies will be needed to shut them down again.¹³

Perhaps we're being unwarrantedly Humean in our assumptions here.¹⁴ Why assume free recombination is the default, especially as the examples from quantum mechanics motivating such of this discussion appears to undercut any initial motivation for being Humean in the first place - namely, the assumption that we can safely take *as a starting point* that the physical world can be accounted for as merely a plurality of modally independent individuals (see Barnes *draft*)? Well, I think if anything dropping the independence assumption - at least as the default - makes things a lot worse for any sort of pluralist. If we make the default complete dependence, then even to account for highly interdependent cases like chessboard, the pluralist still has to spell out exactly what modal variation is allowed— it's not enough just to say that the colour property of every pixel depends on any other; that alone doesn't nail down the statespace to Chessboard. But worse, consider going back to accounting for Switch World. Now the pluralist has to decree that each and every pixel is free from dependence constraints, on a dizzying case by case basis. Essentially, if Humean recombability isn't something we can help ourselves to for free - and rule exceptions to when needed - then positing pluralities is an

¹³ In the limiting case of one object, there could only be anything analogous to such connections if there were certain fundamental properties in the theory that are necessarily uninstantiated.

¹⁴ See Wilson (2013).

even less elegant exercise than we might have thought. Therefore, I don't think challenging the recombability assumption really helps dispel the worry that positing more property bearers will tend to lead to less elegant theorising.

Applied to the composition debate, the analogous issue that believer faces is an unexplained connection between the existence and properties of the parts, on the one hand, and the existence and properties of the whole on the other. While the deflationary theorist concerning the non-fundamental does not believe there is really any such connection - because there aren't really any such things as non-fundamental entities - the inflationist does need to contend with this issue. Merely pointing out that wholes are dependent on their parts for their existence and properties is just to restate the fact that these dependencies need to be posited in the first place. Such connections are only in need of further explanation because the inflationist starts out explaining with more than what could have been explained with less.

9. Some Simple Restatements Arranged Conclusion-wise

PART I:

In **Chapter 1** we examined four positions one could take concerning the relationship between our ordinary existential claims and revisionary ontology. We found several reasons to doubt the Moorean case against revisionary ontology, finding the error-theoretic paraphrase strategy offered by van Inwagen to be a more principled way of reconciling conflicting revisionary ontological and ordinary existential claims: Instead of the revisionary ontologist always losing out, revisionary claims can be sustained if a *systematic paraphrase* can be found for the conflicting ordinary claim. A challenge to this view is presented by the semanticist, who claims that the metasemantic principle of charity ensures that our ‘Moorean’ existential claims (in English) are always true (just not for Moorean reasons): Ontology can thus only proceed without conflict if we assume the revisionist is speaking a different language from the ordinary folk— but resolution is gained at the cost of substantivity, since there’s no privileged language that’s *better* for evaluating existential claims. The Siderean responds with a competing metasemantic argument on which charity loses out to reference magnetism, and so the argument that revisionary ontology can’t take place in ordinary language never gets off the ground.

The stage now set, **I then put forward my own take on the revisionary ontology debate:** I first noted that it would be attractive to have an account of revisionary ontology independent of the reference magnetism story. I then pointed to the central problem for error-theory: its failure to properly separate out the theoretical roles of *factuality* on the one hand, and metaphysical *perspicuity* on the other. I proceeded to set forth a positive account which had the desired features: I then situated this answer within the current literature by noting it could be glossed as one way of implementing a certain general strategy for revisionary metaphysics in the absence of error-theory: The Ontologese Strategy (OS).

Chapter 2 defended OS from the charge that it is unintelligible or esoteric. I gave reasons to suppose that the concept of existence in Ontologese is exactly equivalent to a more determinate *conception* of the ordinary language quantifier: starting from a shared (ordinary) concept of existence we can come to grasp the more determinate meaning of

the Ontologese quantifier simply by understanding the more weighty theoretical role it plays. After revisiting the issue of Mooreanism specifically from the perspective of OS, we are in a position to declare the Ontologese Strategy (OS) the most plausible account of revisionary metaphysics: It clearly separates out factuality from perspicuity in a way which allows us to secure the factuality of Moorean claims while not allowing our ordinary existential claims to artificially constrain our ontological theorising.

In **Chapter 3** after tidying away some metaontological loose ends, I gave an account of the status of relatively natural or partially joint carving existential discourse, such as Chemistry. Or rather, I declined to, but said why that's OK. Instead, I showed that we can still have substantive ontological discussions about such discourse, in a certain important sense: by employing a hypothetically privileged quantifier. I then examined the status of social ontology, *qua* a putative example of substantive but highly unnatural discourse: I argued that we could have substantive debates here too, since this discourse is treated on a par with moderately natural discourse on my account, but stronger (e.g. constructivist) construal of such discourse should not be dismissed out of hand.

In **Chapter 4** we learn that when we ask the Special Composition Question (SCQ) in Ontologese, nihilism yields the best answer according to generally accepted theory choice criteria: it is adequate, gives a powerful unified answer to material object puzzles, does not posit causally redundant entities, is ideologically parsimonious and (arguably) ontologically parsimonious in at least one respect. Hence, the charge that the composition debate is epistemically intractable is dismissed and (*pro tanto*) victory is declared for nihilism.

The **conclusion** of PART 1: OS is the best strategy for conducting revisionary ontology and, thus, for asking SCQ. There is strong (defeasible) reason to adopt nihilism as the answer to that question.

PART II:

Chapter 5 considered what I took to be the most pressing objection to the conclusion of *PART 1*. Namely, that nihilist can't account for genuinely *ontologically emergent properties* (e.g. from quantum mechanics) since there are no suitable bearers in the nihilist's ontology to instantiate such properties. I argued that while there are no individually suitable bearers, collectively simples are adequate bearers for such properties - I then defend the coherence and innocence of plural quantification.

In direct continuation, **Chapter 6** showed that the Emergent Simples strategy, Entanglement Relations strategy and the Fundamental Indeterminacy strategy (therein detailed) were also adequate and parsimonious, giving the nihilist an attractive range of

options for dealing with putative ontological emergence. I conclude that the nihilist should pick whichever of the four strategies best fits with her overall theory.

Chapter 7 conceded that given the lessons of quantum mechanics, there's *pro tanto* reason to suspect that monism (the view there is just one simple concrete object) is a more attractive theory than traditional pluralistic nihilism. But on close examination, there are attractive versions of pluralistic nihilism which appear to be worthy rivals to monism: Further work will be required to settle this question. (So watch this space....)

Finally, in **Chapter 8** I make some remarks in defence of my use of parsimony as a theoretical virtue in the composition debate. For instance, I argue that *qualitative ontological parsimony* might only apply to *sui generis* kinds when giving a general theory of the categories of being, but parsimony with respect to kinds of concrete individuals is reasonably held to be a virtue when one is in the business of giving a *metaphysical cosmology* given that the goal of this enterprise is ostensibly to tell us what kinds of individuals there are and where, and how they are related. I also argue, on similar lines, that when giving a metaphysical cosmology one should seek to minimise relations of dependence between individuals, giving an indirect argument for *quantitative ontological parsimony*. Further to this, I defend my own criterion of ideological parsimony against a rival position by showing that my criterion delivers more principled results.

10. Appendix to Chapter 7: Collapse and Decoherence.

This appendix sketches a way one might attempt to argue for a world of many small and separate entangled systems, rather than treating the world as one big entangled system as in Schaffer (2010). As such, it can be read as a supplement to §7.5 Screen Worlds and the Actual World. If there is a grain of truth in what follows, it would lend urgency to the argument of §7, since there would be a positive reason to doubt Schaffer's assumptions concerning the quantum facts we must take for granted for his argument to get off the ground: as things stand in that chapter, however, I am content to merely point out that his assessment of the facts remains an open interpretive question. Here, though, I want to offer some background on my reasons for doubting some of the quantum mechanical background Schaffer asserts as assumptions of his argument for Priority Monism, and point the reader to passages and problems in the literature on QM that might be built on in the construction of a positive case for Local Holism in metaphysics. My central concern here is to show that objective collapse models of QM are a live option and suggest that on such models then, contra Schaffer, entanglement may be highly localised rather than universal or near-universal.

Made explicit by Schaffer (2010) is an underlying assumption that the Schrödinger equation is the only mechanism that governs the time-evolution of quantum systems. But the bare theory of quantum mechanics doesn't just contain the Schrödinger equation, which provides the means to calculate the deterministic evolution of a system's quantum state, but also the project postulate, which tells you how a quantum state will collapse indeterministically into an eigenstate of an observable upon measurement of that system: call this process *state reduction*.¹ As Schaffer admits, 'It is controversial whether temporal evolution is always via the Schrödinger dynamics (*unitarity*), or whether there is a further dynamics of wave-function collapse.' (2010, p.52)

The indeterministic collapse of the system into an eigenstate is clearly a violation of the

¹ See Penrose (2004, p.528).

deterministic evolution described by the Schrödinger equation.² Since the Schrödinger equation gives no prediction of when such a collapse will occur, the two components of the theory are independent of one another. Given the existence of this independent and conflicting bit of theoretical machinery in the theory, it's too quick to conclude from the outset that Schrödinger evolution is the only relevant factor in determining any aspect of a quantum systems' behaviour. On the face of it, this is obviously not the case:³ The Schrödinger equation plus appropriate application of the measurement principle give us the quantum mechanical story of how the world changes over time. The dynamics of quantum systems are governed by both Schrödinger evolution and state reduction.

Hence, the mere fact that the Schrödinger evolution tends to spread entanglement is not by itself reason to conclude that in reality entanglement will inevitably increase over time. The additional assumption is required that either it is *never* appropriate to postulate collapse when determining the evolution of a quantum system, or that the way in which application of the postulate affects the dynamics is not relevant to the question of global entanglement. Regarding the latter part of this assumption, I think there is good reason to think that collapse *is* relevant to the dynamics of entanglement— in such a way that is again at odds with the evolution described by the Schrödinger equation.

Indeed, after noting that Schrödinger evolution tends to spread entanglements and immediately following the excerpt Schaffer quotes, Penrose (2004, p.592) says:

Yet, we seem to get along pretty well, in everyday life, without even noticing these entanglements. Why is this? If we are to get no help from quantum theory's U process [Schrödinger evolution], then we must turn to its other essential ingredient: the R process [state reduction]. In fact, we have already seen something of the way that this might help in our considerations of EPR effects. Recall that I envisaged performing a measurement on an EPR pair, the other member of which was approaching my colleague on the planet Titan. If I make my measurement first, then upon my performing this measurement, this very act would cut my colleague's particle free of its entanglement with mine, and from then on (until it became measured by my colleague) it would possess a state vector of its own, unencumbered by any further responsibility to its partner, no matter what I might subsequently do to it. Thus, it seems, it is measurements that slash through these entanglements.

While the Schrödinger equation spreads entanglement, plausibly state reduction destroys entanglement: A pair of particles entangled with respect to some property (e.g. z-

² See Penrose (2004, p.530).

³ Sneed (1965, p.22)

axis spin) is necessarily in a supposition with respect to that property (i.e. neither has a definite spin state), while collapse forces the quantum particle system into an eigenstate, destroying the superposition (stochastically assigning a definite z-axis spin value to each particle in this example). Elsewhere, Penrose sets out the view thus (1998, p.1931):

From time to time, whenever it is considered that a 'measurement' has taken place, unitary evolution is abandoned, and the state vector of the quantum system is taken instead to 'jump' into an eigenstate of some Hermitian operator, determined by the apparatus that is doing the measurement. It is this latter process that, in effect, 'cuts' the entanglements that the system under consideration previously had with the outside world.

The key question is when, or if, state reduction occurs. Several major attempts have been made to rid quantum mechanics of the stochastic vagaries of state reduction through reinterpretation of the formalism: Many Worlds, the Modal Interpretation, Bohm Theory and Consistent Histories⁴ for instance. Each boasts an interpretation of QM which rely only on an unaltered Schrödinger equation without employing collapse of the wave function – relegating the projection postulate to a non-fundamental part of the theory, with at most pragmatic value.

Yet, if there is no state reduction, how do we explain the fact that in everyday life we do not (so it would certainly seem) observe superpositions and entanglement, but rather well localised objects with well-defined physical properties? Or, in the language of many worlds, how come we don't constantly observe these other worlds interfering with events in our own? Without the help state reduction, it appears our hopes must rest on the phenomenon known as *decoherence*.⁵ Decoherence is the process by which, as quantum systems become larger – i.e. as they causally interact with other independent systems and become entangled with them – they exhibit a closer and closer approximation to classical behaviour.⁶ Coherent states produce distinctly quantum behaviour, such as the correlations exhibited by an entangled pair in EPR experiments: but when the complex amplitudes of two systems are out of phase and they become entangled there is a loss of

⁴ For an examination of each of these interpretations in light of the phenomenon of decoherence see Schlosshauer (2004) – esp. p.22, p.26, p.31, p.33 respectively. Compare Schlosshauer (2007), esp., §8.2, §8.3, §8.5.

⁵ The decoherence programme is sometimes presented as a solution to the measurement problem, and thus a complete interpretation of quantum mechanics in itself. The plausibility of this claim is at best a controversial matter, however. See Tanona (2013); compare Pessoa (1998).

⁶ For an extended accessible discussion of decoherence see Schlosshauer (2007).

coherence, with distinctly quantum correlations becoming much less probable. So, as an initially coherent system becomes entangled with other systems of arbitrary phase through interaction, it will tend to become less coherent, and its state space will begin to look more classical. As Kallfelz (2009, p.4001) explains:

Taking advantage of a wholly quantum mechanical description of system, measurement apparatus, and environment (*S-M-E*), decoherence-based approaches quantitatively show how in vastly short time scales “off-diagonal” elements of the *S-M-E* density matrix ρ_{S-M-E} become negligible, when “traced” over the “pointer basis” (or the basis describing the states of the measurement apparatus coupled with the environment). This is a technical way of characterizing the Projection Postulate via a statistical “washingout” mechanism: The vastly greater number of degrees of freedom of the environment + measuring apparatus render it nearly impossible to macroscopically observe a system in a superposition state -- in practically all cases when the system is a vast ensemble of fermions.

Crucially, decoherence predicts genuinely classical behaviour *at the limit*.⁷ Even a vast entangled system – such as the entire world – will not have a completely classical state space. It just becomes increasingly *unlikely* that quantum effects will be observed on a macroscopic scale. What’s more, although the state space of the system will be approximately classical, this does not mean the system can begin to be treated as made of independent approximately classical sub-systems: For, it’s precisely an increase in the distinctively quantum phenomenon of entanglement which, ironically, causes the system to behave more classically. A more decoherent state is generally a more entangled one: thus less apt to be treated as being composed of independent sub-systems. The system as a whole encodes ever more information than its classical counterpart would, in order to account for the possibility of its ever less likely deviations from classical behaviour.

If decoherence is the sole reason for the appearance of classicality in the world, then Schaffer’s starting assumption can happily be considered correct, even if it is particularly non-straightforward what ontological conclusion we should draw from the *entanglement-with-decoherence* picture (an issue explored in §7.5). However, if Penrose is right that the state reduction process is fundamental and mitigates against the entanglement propagated by Schrödinger evolution, then there is an alternative starting point to Schaffer’s that seems worthy of consideration. The details would depend on how one incorporated collapse into the theory: One need not have recourse to pervasive instrumentalism or antirealism of the Copenhagen interpretation, on which a description of reality is always

⁷ See Tanona (2013).

relativized to a measurement,⁸ or to the inescapable mind-dependence of Consciousness Causes Collapse.⁹ Indeed, objective collapse interpretations such as CSL (Continuous Spontaneous Localisation) provide principled accounts of how state reduction might interact with Schrödinger evolution that are entirely realist and mind-independent. Adler (2003) writes:

Heuristically, the idea here is that quantum mechanics may be modified by a low level universal noise, akin to Brownian motion, possibly arising from new physics at the Planck scale, which in certain situations causes reduction of the state vector. This approach, which has been developed in great detail, reproduces the observed fact of discrete outcomes governed by Born rule probabilities, and predicts the maintenance of coherence where that is observed (including in systems with large numbers of particles, such as recent superconductive tunneling and molecular diffraction experiments), while predicting state vector reduction when the apparatus parameters are those characterizing measurements.

The picture is one of many small-scale state reductions or mind-independent ‘measurements’ destroying entanglement for the quantity so measured on each occasion. The larger a system becomes the more likely it will be to undergo collapse: State reduction will tend to occur on quite a small scale while still leaving room for some large scale quantum effects in keeping with empirical observation. And ‘...after the hypothetical spontaneous collapse the combined system is no longer described by the entangled state[...] but by one of the product states...’ (Gisin 1991, p.445) Thus, although there will still exist far reaching entanglement on this picture – for this is not a way to salvage the view that nature consist of entirely independent particles – we might expect that reality can in principle be carved up into a plurality of separate systems.¹⁰

This provides a competing or complimentary explanation of why we don’t observe superpositions in everyday life, and why we can treat the world as if it were made up of local, independent systems – because it *is* made up of fairly local, independent systems,

⁸ See, e.g., Barad (1996).

⁹ See, e.g. Kallfelz (2009: 400).

¹⁰ Schlosshauer (2007, §8.4.3.) discusses the potential problem that ‘all collapse models lead to an only approximate reduction of the wave function’ (p.351) As such, there’s a sense in which collapse models are in the same boat as decoherence approaches. However, since collapse models ‘at least approximately select[...]’ (p.351) a single outcome, one only needs to allow for the ‘weakening of the eigenvalue-eigenstate link to make this state of affairs correspond to the (objective) existence of a determinate physical property’ (p.351), whereas on the decoherence approach, even after ‘complete decoherence’ (p.351) one would still be left with ‘an improper mixture’ (p.351) of states.

with no entangled system ever gets very large or spreads its entanglement very far before collapsing.¹¹ Here decoherence and state reduction will be working together to produce the appearance of classicality on the large scale, yet state reduction will be working against entanglement, the distinctively quantum mechanism by which decoherence operates: The result being a world which is not just superficially classical, but has the fundamentally classical feature of separable independent systems— even if these systems themselves exhibit distinctively quantum behaviour. As Penrose (1998, p.1932) puts it, discussing his version of objective collapse:

...state-vector reduction is considered to take place objectively and spontaneously, in contradiction with the strict unitary evolution of the quantum state. Accordingly, when a measurement is made on a system, its entanglements with the outside world are indeed objectively cut, so that the aforementioned problem of the persistence of complicated entanglements of a system with the rest of the universe is thereby removed.

Given that one is adding new dynamics, relative to interpretations that only take the Schrödinger equation as fundamental, the deviation from the predictions of the Schrödinger equation are in principle empirically testable,¹² but are at present empirically adequate. Such conflict in predictions should not be a surprise, since we're here in search of a theory that provides a more clear 'cut-off' between the quantum and classical domains than the phenomenon of decoherence alone provides for. Adler (2003) continues:

The stochastic approach may ultimately be falsified by experiment, but it constitutes a viable phenomenological solution to the measurement problem. Decoherence, in the absence of a detailed theory showing that it leads to stochastic outcomes with the correct properties, has yet to achieve this status.

It is massively beyond the scope of this work to attempt to adjudicate between the two interpretations. My intention here has been to present a clear alternative to the view on entanglement and holism in quantum mechanics that Schaffer sets out which I think is also worthy of serious consideration, the existence of which may alter our approach to the

¹¹ As Schlosshauer (2007, §8.4.1) notes 'The Preferred Basis Problem' of *which* observable is spontaneously 'measured' by 'postulating a mechanism that leads to spatial localisation.' (p.349). He notes that the physical implications of decoherence make this a natural choice.

¹² See Ghose (1998) for a discussion of how collapse theories and the decoherence approach might in principle be empirically differentiated. On the other hand, Schlosshauer (2004) suggests the existence of decoherence may make disconfirmation of collapse theories impossible in practice. Meanwhile, Pearle (1997) and Diosi (2014) provide reason for thinking the decoherence approach is inadequate or incomplete for omitting collapse.

Chapter 7 discussion of ontology. Another important thing we've learnt is that even if there is no collapse, and the world is just one big superposition, the world-system will not by and large behave in the very holistic way exemplified by a smaller entangled systems due to decoherence. Indeed, the world will behave very much as if it is made up of lots of independent, local subsystems— just *not quite*. We have the puzzling situation that by increasing entanglement we reduce the holistic tendencies of a system, yet without ever making it possible to entirely separate out independent subsystems.

Though I'm not trying to argue for collapse interpretations, I just want to say something about a common objection to such interpretations - namely, that collapse theories are less simple or otherwise objectionable because they involve 'new physics'. Dickson (1998, pp.58-61) gives a detailed and convincing reply to the new physics objection. His main argument is that what constitutes "the physics", or quantum mechanics proper, is not clearly defined: If the projection postulate is considered part of the core theory then no collapse theories deviate from this by ditching it. So, advocates of such interpretations must mean that they add nothing to the Schrödinger equation: But it's not clear why we should insist on agreement with the Schrödinger equation beyond the current level of experimental success— something that all interpretations share. What we take to be "the physics" beyond the current experimental evidence seems to be largely the interpretational issue at stake when we decide to interpret the bare theory. Blind allegiance to the Schrödinger equation seems to give unwarranted weight to social and historical accident. Dickson (p.60) concludes:

When "adding no new physics" is rendered in such a way as to make it clearly a virtue, it becomes a virtue shared by many interpretations. When it is rendered in such a way that only few interpretations do it, it is no longer clearly a virtue. And worse, because "quantum mechanics" is already ill defined, it is not clear what we are "not adding" in the first place.

The above is only a brief sketch of a defence that would do nothing to persuade someone convinced of the soundness of the new physics objection: My purpose is just to point in the direction of arguments in the literature that might be successful in heading off the objection, in support of my contention that both collapse and non-collapse views should be considered as live possibilities— even if decoherence approaches are more in vogue.¹³

¹³ As for how this choice interfaces with wider theory: While earlier objective collapse models struggled to incorporate Special Relativity, this has been rectified in more recent formulations: now

To sum up, then, if a collapse interpretation of QM is correct the world as a whole is most naturally described as a collection of localized (though not maximally local) facts. For a physically holistic system rarely gets to a macroscopic size without collapsing. As such, there is much less pressure to accept a monist account of property dependence; for the properties of certain parts of the world can be fully specified without reference to the state of the world as a whole. Given that the properties of parts of the world *can* be specified independently of the whole, there would also a *prima facie* cost to the monist for not doing so: the monist misses out on the combinatorial explanatory advantages associated with pluralism. One important result that this analysis throws up, however, is that even if entanglement is universal, as Schaffer argues, quantum decoherence means that the argument from universal entanglement to monism from theoretical virtues is still not straightforward.

considerations from General Relativity might actually constitute a strong positive reason to adopt certain collapse theories. See, e.g., Okon & Sudarsky (2014). Also, Penrose (2004).

Bibliography

- Adler, Stephen L. (2003). 'Why Decoherence has not Solved the Measurement Problem: A Response to P.W. Anderson', *Studies in History and Philosophy of Modern Physics* 34 135-142.
- Armstrong, D. M. (1997). *A World of States of Affairs*. UK: Cambridge
- Azzouni, Jody (2004). *Deflating Existential Consequence: A case for nominalism*. NY: Oxford
- Baxter, D. L. M., (1988a). "Identity in the Loose and Popular Sense", *Mind* 97: 576-582.
- Baxter, D. L. M., (1988b). "Many-One Identity", *Philosophical Papers* 17: 193-216.
- Barad, Karen (1996). 'Meeting the universe halfway: Realism and Social Constructivism without Contradiction' *L.H. Nelson and J. Nelson (eds.), Feminism, Science and the Philosophy of Science*, pp.161-194. Kluwer
- Barnes, Elizabeth (2007). 'Vagueness and Arbitrariness: Merricks on Composition', *Mind* 116 (1)
- Barnes, Elizabeth (2009). 'Indeterminacy, Identity and Counterparts.' *Synthese* 168 (1)
- Barnes, Elizabeth (2010). 'Arguments Against Metaphysical Indeterminacy', *Philosophy Compass* (5): 953-64.
- Barnes, Elizabeth (2012). 'Emergence and Fundamentality', *Philosophical Studies*.
- Barnes, Elizabeth (*forthcoming-a*). 'Going Beyond the Fundamental: Feminism in Contemporary Metaphysics.' Proceedings of the Aristotelian Society.
- Barnes, Elizabeth (*forthcoming-b*). 'Fundamental Indeterminacy'. *Analytic Philosophy*.
- Barnes, Elizabeth (*draft*). 'Symmetric Dependence'.
- Barnes, E. & Cameron, R. P. (2009). 'The open future: bivalence, determinism and ontology.' *Philosophical Studies* 146: 291-309
- Barnes, E. & Cameron, R. P. (2011). 'Back to the Open Future', *Philosophical Perspectives* 25: 1-26.
- Barnes, E. & Williams, J. R. G. (2012). 'A Theory of Metaphysical Indeterminacy.' *Oxford Studies in Metaphysics* 6
- Bennett, Karen (2009). 'Composition, colocation, and metaontology.' In David John Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.

- Bennett, Karen (2011). 'Construction area (no hard hat required)', *Philosophical Studies* 154:79-104
- Bennett, Karen (*Draft*). *Making Things Up*. OUP [Preliminary draft— limited circulation.]
- Boolos, George (1984). 'To be is to be a value of a variable (or to be some values of some variables)', *Journal of Philosophy* 81 (8):430-449.
- Brouwer, Thomas (MS.). 'The metaphysical commitments of logic'. URL = <http://theses.whiterose.ac.uk/5042/>
- Bueno, O & Shalkowski, SA (2014). 'Modalism and Theoretical Virtues: Toward an Epistemology of Modality.' *Philosophical Studies*
- Burke, Michael (1994a). 'Preserving the Principle of One Object to a Place: A Novel Account of the Relations Among Objects, Sorts, Sortals, and Persistence Conditions', *Philosophy and Phenomenological Research*, 54 : 591-624.
- Burke, Michael (1994b). 'Dion and Theon: An Essentialist Solution to an Ancient Puzzle,' *Journal of Philosophy* 91: 129-39.
- Cameron, Ross P. (2007). 'The contingency of composition', *Phil Studies* 136: 99-121
- Cameron, Ross P. (2008a). 'Truthmakers and Ontological Commitment: or, how to deal with complex objects and mathematical ontology without getting into trouble.' *Philosophical Studies* 140, pp.1-18.
- Cameron, Ross P. (2008b). 'Truthmakers and Necessary Connections.' *Synthese* 161: 27-45.
- Cameron, Ross P. (2010). 'How to have a radically minimal ontology', *Philosophical Studies* 151: 249-64
- Cameron, Ross P. (2011) 'Truthmaking for Presentists' *Oxford Studies in Metaphysics* Vol.6.
- Cameron, Ross P. (2012). 'Why Lewis's analysis of modality succeeds in its reductive ambitions.' *Philosophers' Imprint*, vol. 12, no. 8, pp. 1-21.
- Carnap, Rudolf. (1965). 'Empiricism, Semantics and Ontology', In *Meaning and Necessity*, pp.205-21. USA: University of Chicago Press.
- Caves, Richard L. J. (*Draft*). 'The Nihilist's Infinite Descent.'
- Chalmers, David J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. OUP.
- Chalmers, David J. (2009). 'Ontological Anti-Realism', in Chalmers, D. J., Manley, D., & Wasserman, R. (eds.), *Metametaphysics: new essays on the foundations of ontology*, pp.77-129. New York: OUP

- Chalmers, Manley & Wasserman (eds.). (2009). *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- Cornell, David (2013). 'Monism and Statespace: a reply to Sider'. *Analysis* 73 : 230 - 236.
- Cornell, David (MS.). *Why There Is Only One Thing: A Defence of Ontological Monism*. URL = <<http://etheses.whiterose.ac.uk/4939/>>
- Cowling, Sam (2013). 'Ideological Parsimony', *Synthese* 190: 889-908.
- Cowling, Sam (*forthcoming*). 'No Simples, No Gunk, No Nothing', *Pacific Philosophical Quarterly*
- Diosi, Lajos (2014). 'Gravity-Related Wave Function Collapse.' *Foundations in Physics* 44(5): 483-491.
- Daly, Chris & Liggins, David (2010). 'In defence of error theory', *Philosophical Studies* 149.2, 209-230.
- Daly, Chris & Liggins, David (*Draft*). 'Hirsch's Charity Argument Against Revisionary Ontology.'
- Derby, George (2010). 'Quantum Mechanics and Metaphysical Indeterminacy.' *Australasian Journal of Philosophy* 88: 227-245.
- Dickson, W. Michael (1998). *Quantum Chance and Non-Locality*. UK: Cambridge University Press
- Divers, John (2002). *Possible Worlds*. New York: Routledge
- Dorr, Cian (2005). 'What We Disagree About When We Disagree About Ontology', in Kalderon, M. (ed.), *Fictionalism in Metaphysics*, pp.234-86. Oxford: Oxford University Press.
- Eklund, Matti (2009). 'Carnap and Ontological *Pluralism*', in Chalmers, D. J., Manley, D., & Wasserman, R. (eds.), *Metametaphysics: new essays on the foundations of ontology*, pp.130-56. New York: OUP
- Evans, Gareth (1978). 'Can There Be Vague Objects?' *Analysis* 38 : 208.
- Fine, Kit (2001). 'The Question of Realism'. *Philosophers' Imprint* 1 (2):1-30.
- Fine, Kit (2009). 'The Question of Ontology.' In David John Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press. 157-177.
- Forrest, Peter (2004). 'The Real but Dead past: a reply to Braddon-Mitchell', *Analysis* 64 (4): 358-362.
- French, Steven (2004). *The Structure of the World: Metaphysics and Representation*. UK: OUP
- Gallios, Andre (1998). *Occasions of Identity*. Oxford: Clarendon Press

- Geach, P. T. (1980). *Reference and Generality*. 3rd edn. Ithaca: Cornell University Press.
- Ghose, Partha (1998). 'Spontaneous Collapses vs. Decoherence', *PINSA-A* 64: 193-210.
- Goodman, Nelson (1951). *The Structure of Appearance*. Cambridge, MA: Harvard University Press.
- Gisin, N. (1991). 'Can quantum entangled states collapse spontaneously?' *Physics Letters A* 155: 445 - 449.
- Haslanger, Sally. (2000). 'Gender, Race: (What) Are They? (What) Do We Want Them To Be?' *Noûs* 34 (1): 31-55.
- Haslanger, Sally (2007). 'Social knowledge, social structure and ideological critique,' *Philosophical Inquiries* 17 (1): 70-91.
- Haslanger, Sally. (2011). 'Ideology, Generics, and Common Ground' in Charlotte Witt (ed.), *Feminist Metaphysics*, pp.179-207. New York: Springer.
- Haslanger, Sally. (2012). *Resisting Reality*. Oxford: OUP.
- Hawley, Katherine (2001). *How Things Persist*. NY: OUP
- Hirsch, Eli (2002). 'Quantifier Variance and Realism', *Philosophical Issues* 12
- Hirsch, Eli (2005). 'Physical-Object Ontology, Verbal Disputes, and Common Sense', *Philosophy and Phenomenological Research* 70, pp.67-97.
- Hirsch, Eli (2008). 'Language, Ontology and Structure', *Noûs* 42, pp.509-28.
- Hirsch, Eli (2009). 'Ontology and Alternative Languages' in Chalmers, D. J., Manley, D. & Wasserman, R. (eds.) *Metametaphysics*, pp.231-59. NY: OUP.
- Hofweber, Thomas (2009). 'Ambitious, Yet Modest, Metaphysics', pp.260-89. In Chalmers, D. J. Manley, D. & Wasserman, R. (eds.) *Metametaphysics*. NY: OUP.
- Horgan, T. & Potrc, M. (2000). 'Bobjectivism and Indirect Correspondence'. *Facta Philosophica* 2:249-270
- Horgan, T. & Potrc, M. (2008). *Austere Realism*. MIT
- Hudson, Hud (2003). 'Alexander's dictum and Merricks' dictum.' *Topoi* 22 : 173 - 82.
- Jackson, Frank (1994). 'Armchair Metaphysics', in O'Leary-Hawthorne, J. & Michael, M. (eds.), *Philosophy in Mind: the place of philosophy in the study of mind*, pp.23-42. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Jenkins, Katharine. (MS.). 'Amelioration and Inclusion'.
- Jones, Nicholas (2012). "Multiple constitution", *Oxford Studies in Metaphysics*

- Kallfelz, William M. (2009). 'A Response to G.B. Bagci's "Ghirardi-Rimini-Weber Collapse Theory and Whiteheadian Process Philosophy"' *Process Studies* 38.2: 399-416.
- Kim, Jaegwon (1993). 'The Nonreductivist's Troubles with Metaphysical Causation', in John Heil and Alfred Mele (eds.), *Mental Causation*. Oxford: Oxford University Press 189 - 210.
- Kim, Jaegwon (1989). 'The Myth of Nonreductive Materialism', *APA Proceedings* 63: 31-47.
- Kim, Jaegwon (1999): 'Making Sense of Emergence'. *Philosophical Studies* 95 (1), 3-36
- Korman, Daniel (*forthcoming*). 'The Language of the Ontology Room'. *Noûs*.
- Koslicki, Katharin (2008). *The Structure of Objects*. OUP
- Ladyman, James & Ross, Don (2007). *Every Thing Must Go: Metaphysics Naturalized*. UK: OUP
- Lewis, David (1973). *Counterfactuals*. Oxford:Blackwell.
- Lewis, David (1983a), "Extrinsic Properties", *Philosophical Studies*, 44: 197-200.
- Lewis, David (1983b). 'New work for a theory of universals', *The Australasian Journal of Philosophy* 61: 343-77.
- Lewis, David (1986a). *Philosophical Papers*, Vol. ii (Oxford: Oxford University Press).
- Lewis, David (1986b). *On the Plurality of Worlds*. Oxford Blackwell
- Lewis, David (1991). *Parts of Classes*. Oxford: Blackwell.
- Lewis, David (1993). 'Many, but Almost One' in *Ontology, Causality and Mind: Essays in Honour of D M Armstrong*, Bacon, John (ed) New York: Cambridge University Press.
- Linnebo, Øystein (2004). 'Plural Quantification', Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Spring 2013 edition)*, URL = <http://plato.stanford.edu/archives/spr2013/entries/plural-quant/>
- Markosian, Ned (1998). 'Simples'. *Australian Journal of Philosophy* 76: 213-26.
- Markosian, N. (2004). 'Simples, stuff and simple people.' *The Monist* 87(3): 405-421.
- Maudlin, Tim (2007). 'Why be Humean?' in *The Metaphysics within Physics*, pp.50-77.
- McDaniel, Kris (2007). 'Extended Simples'. *Philosophical Studies* 133:131-141
- McDaniel, Kris (2010). 'A Return to the Analogy of Being'. *Philosophy and Phenomenological Research* LXXXIX No.3: 688 - 717.
- McGrath, Matthew (2005). 'No objects, no problem?' *Australasian Journal of Philosophy* 83: 475 - 86.

- McKenzie, Kerry (MS.). *Physics without fundamentality*. URL = <http://theses.whiterose.ac.uk/4052/>
- Melia, Joseph (1992). 'A Note on Lewis's Ontology', *Analysis* 52 : 191-2.
- Melia, Joseph (2000). 'Weaseling Away the Indispensability Argument', *Mind* 109 (435): 455-479.
- Melia, Joseph (2005). 'Truthmaking without Truthmakers', in Helen Beebe and Julian Dodd (eds.), *Truthmakers* OUP.
- Merricks, Trenton (2001). *Objects and Persons*. New York: Oxford
- Merricks, Trenton (2005). 'Vagueness and Composition'. *Mind* 114: 615-37.
- Miller, Elizabeth (2013). 'Quantum Entanglement, Bohmian Mechanics and Humean Supervenience.' *Australasian Journal of Philosophy*
- Miller, Kristie (2006). 'The physics of extended simples', *Analysis* 66(3) 222-226
- Miller, Kristie (2009). 'Defending Contingentism in Metaphysics.' *Dialectica*
- Moore, G. E. (1925) 'A defence of common sense' in J. H. Muirhead (ed.), *Contemporary British Philosophy* (2nd series), Allen and Unwin, London: pp.193-223. Reprinted in Moore, G.E. (1959) *Philosophical Papers*, pp.32-59. London: George, Allen & Unwin.
- Morton, Adam (1975) 'Complex Individuals and Multigrade Relations.' *Noûs* 9 (3): 309-18.
- Mundy, Brent (1989) 'Elementary Categorical Logic, Predicates of Variable Degree, and a Theory of Quantity.' *Journal of Philosophical Logic* 18: 115 - 140.
- Nolan, D. (1997). 'Quantitative Parsimony', *The British Journal for the Philosophy of Science* 48: 329-343.
- Nolan, D. (2006). 'Vagueness, Multiplicity and Parts.' *Noûs* 40:4 : 716-737.
- Okon, Elias & Sudarsky, Daniel (2014). 'Benefits of Objective Collapse Models for Cosmology and Quantum Gravity', *Foundations in Physics* 44: 114-143.
- Paul, L. A. (2012a). 'Building the world from its fundamental constituents', *Philosophical Studies* 158: 221-256
- Paul, L. A. (2012b). 'Metaphysics as modeling: the handmaiden's tale', *Philosophical Studies* 160: 1-29.
- Parsons, Terence. (2000). *Indeterminate Identity*. NY: Oxford
- Parsons, Josh (2004). 'Distributional Properties' in *Lewisian Themes*. Jackson, F. and Priest, G. (eds.). Oxford: Oxford University Press, pp.173-180.

- Pearle, Phillip (1997). 'True Collapse and False Collapse.' In Da Hsuan Feng and Bei Lok Hu (eds.) *Quantum Classical Correspondence: Proceedings of the 4th Drexel Symposium on Quantum Nonintegrability, Philadelphia, PA, USA, September 8-11, 1994*. pp. 51-68. Cambridge, MA: International Press
- Penrose, Roger (2004). *The Road to Reality*. UK: Random House
- Penrose, Roger (1998). 'Quantum computation, entanglement and state reduction.' *Philosophical Transactions of the Royal Society A* 356 : 1927 - 1939.
- Pessoa, Osvaldo (1998). 'Can the decoherence approach help to solve the measurement problem?' *Synthese* 113: 323-346.
- Price, Huw (2009). 'Metaphysics after Carnap: The Ghost who Walks?' In Chalmers, D. J., Manley, D. & Wasserman, R. (eds.) *Metametaphysics*. New York: OUP.
- Putnam, H. (1973). "Meaning and Reference", *Journal of Philosophy* 70, 699-711.
- Putnam, H. (1981). 'Brains in a Vat'. In *Reason, Truth and History*, pp.1-21. UK: Cambridge University Press
- Quine, Willard van Orman (1948). 'On What There Is.' *Review of Metaphysics* 2 (5):21–36.
- Quine, Willard van Orman (1960a). *Word & Object*. NY & London: Wiley.
- Quine, Willard van Orman (1960b). "Variables Explained Away." *Proceedings of the American Philosophical Society* 104: 343–347.
- Quine, Willard van Orman (1966). 'On Carnap's View of Ontology', in Quine (1996) *The Ways of Paradox*. New York: Random Houses
- Quine, Willard van Orman (1971). "Algebraic Logic and Predicate Functors." In Richard Rudner and Israel Scheffler (eds.), *Logic and Art: Essays in Honor of Nelson Goodman*, 214–238. New York: Bobbs-Merrill.
- Rayo, Agustin (MS.). *Plural Predication*. URL = <https://dspace.mit.edu/bitstream/handle/1721.1/8763/48123914.pdf?sequence=1>
- Schaffer, Jonathan (2003). 'Is There a Fundamental Level?' *Noûs* 37: 498-517.
- Schaffer, Jonathan (2007) 'From nihilism to monism', *Australasian Journal of Philosophy*, 85(2): 175–191.
- Schaffer, Jonathan (2009). 'On What Grounds What', in Chalmers, D. J., Manley, D. & Wasserman, R. (eds.) *Metametaphysics*, pp.347-383. New York: OUP.
- Schaffer, Jonathan (2010). 'Monism: The Priority of the Whole.' *Philosophical Review* 119 (1):31-76.
- Schlosshauer, Maximilian (2007). *Decoherence and the quantum-to-classical transition*. New York: Springer

- Schlosshauer, Maximilian (2004). "Decoherence, the measurement problem, and interpretations of quantum mechanics," *Rev. Mod. Phys.* 76: 1267–1305.
- Sider, Theodore. (1993). 'Van Inwagen and the Possibility of Gunk.' *Analysis* 53 : 285 – 9.
- Sider, Theodore (2001a). *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
- Sider, Theodore. (2001b) 'Criteria of Personal Identity and the limits of conceptual analysis', *Philosophical Perspectives* 15, pp.189-209.
- Sider, T. (2007). 'Against Monism', *Analysis*, 67, 1, pp.1-7
- Sider, Theodore (2008). 'Monism and Statespace Structure'. In Robin Le Poidevin, ed., *Being: Developments in Contemporary Metaphysics*. Cambridge University Press, pp. 129–150
- Sider, T. (2009). 'Ontological Realism', in Chalmers, D. J., Manley, D., & Wasserman, R. (eds.), *Metametaphysics: new essays on the foundations of ontology*, pp.130-56. New York: OUP
- Sider, Theodore (2011). *Writing the Book of the World*. Oxford University Press.
- Sider, Theodore (2013). 'Against Parthood'. Karen Bennett and Dean W. Zimmerman, eds., *Oxford Studies in Metaphysics*, volume 8, Oxford: OUP
- Simons, P. 2004. Extended simples: a third way between atoms and gunk. *The Monist*. 87(3): 371–385.
- Skow, Bradford (2010). 'Deep metaphysical indeterminacy.' *Philosophical Quarterly* 60: 851-8.
- Sneed, Joseph D. (1965). 'Von Neumann's argument for the projection postulate.' *Philosophy of Science* 33: 22-39.
- Tanona, Scott (2013). 'Decoherence and the Copenhagen cut', *Synthese* 190: 3625-3649.
- Teller, P. (1995). *An Interpretive Introduction to Quantum Field Theory*. Princeton
- Thomasson, Amie L. (2007). *Ordinary Objects*. NY: OUP
- Thomasson, Amie L. (2009). 'Answerable and Unanswerable Questions', in Chalmers, D. J., Manley, D. & Wasserman, R. (eds.) *Metametaphysics*, pp.231-59. NY: OUP.
- Thomson, J.J.T. (1983). 'Parthood and Identity Across Time'. *Journal of Philosophy*, 80: 201–20.
- Thomson, J.J.T. (1998). 'The Statue and the Clay,' *Noûs*, 32: 148–173.
- Turner, Jason (2010). 'Ontological Pluralism', *The Journal of Philosophy* 107.1: 5-34
- Turner, Jason (2011). 'Ontological Nihilism', in Zimmerman, D. & Bennett, K. (eds.) *Oxford Studies in Metaphysics* 6. Oxford: OUP

- Turner, Jason (2013). '(Metasemantically) Securing Free Will.' *Australasian Journal of Philosophy*, 91:2, 295-310.
- Turner, Jason (MS.) *Ontology, Quantification and Fundamentality*. URL = <<https://rucore.libraries.rutgers.edu/rutgers-lib/24072/>>
- Unger, P. (1979). 'There are no Ordinary Things'. *Synthese*, 41: 117-154.
- Unger, P. (1980). 'The Problem of the Many.' *Midwest Studies in Philosophy* 5 (1):411-468 (1980)
- Van Inwagen, Peter (1990). *Material Beings*. Cornell University Press.
- Wallace, Meg (2011a.). 'Composition as Identity: Part 1', *Philosophy Compass* vol. 6(11): 804-816
- Wallace, Meg (2011b). 'Composition as Identity: Part 2', *Philosophy Compass* vol. 6(11): 817-827
- Wallace, Meg (*draft*) 'The Many Advantages of Composition as Identity'.
- Wasserman, Ryan (2013). 'Material Constitution', *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), <<http://plato.stanford.edu/archives/sum2014/entries/material-constitution/>>.
- Weatherson, Brian (2009). 'The Problem of the Many', *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), <<http://plato.stanford.edu/archives/win2009/entries/problem-of-many/>>.
- Williams, J. R. G. (2006). 'Illusions of Gunk'. *Philosophical Perspectives* 20, pp.493-513
- Williams, J. R. G. (2012). 'Requirements on reality.' in *Grounding and Explanation* Correia and Schneider (eds.) Cambridge University Press
- Wilson, Jessica (2013). 'Three Dogmas of Metaphysical Methodology'. *Philosophical Methodology: The Armchair or the Laboratory?*, Matthew Haug (ed.). Oxford: Routledge, pp.145-165.
- Yi, Byeong-uk (1999). "Is Two a Property?", *Journal of Philosophy* 96 :163-90
- Yi, Byeong-uk (2002). *Understanding the Many*. Routledge.
- Yi, Byeong-uk (2005). 'The Logic and Meaning of Plurals. Part I', *Journal of Phil Logic* 34 : 459-506
- Yi, Byeong-uk (2006). 'The Logic and Meaning of Plurals. Part II', *Journal of Phil Logic* 35 : 239-88
- Zeh, H. D. (2003). "There Is No 'First' Quantization." *Physics Letters A*309: 329-34.

Abbreviations

Believer	=	Believer in the existence of composite objects
BUC	=	Bump under the Carpet objection to nihilism
CAI	=	Composition as Identity
CSL	=	Continuous Spontaneous Localisation model
ED	=	Epistemic Dismissivism
GMR	=	Genuine Modal Realism
LH	=	Local Holism (Pluralism + Priority of the Whole)
MH	=	Mereological Holism (Monism + Priority of the Whole)
ML	=	Mereological Localism (Pluralism + Priority of the Parts)
Nihilism	=	Mereological Nihilism
Nihilist	=	Proponent of Mereological Nihilism
OS	=	Ontologese Strategy
QFT	=	Quantum Field Theory
QM	=	Quantum Mechanics
Qual-IP	=	Qualitative Ideological Parsimony
Qual-OP	=	Qualitative Ontological Parsimony
Quant-IP	=	Quantitative Ideological Parsimony
Quant-OP	=	Quantitative Ontological Parsimony
SAQ	=	Special Arrangement Question
SCQ	=	The Special Composition Question
SFQ	=	The Special 'F' Question
SMQ	=	The Special Merging Question
SPQ	=	The Special Merging Question