

ESSAYS ON METAPHYSICAL EXPLANATION

By

EZRA ALEXANDER RUBENSTEIN

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Jonathan Schaffer and Theodore Sider

And approved by

New Brunswick, New Jersey

January 2023

© 2022 Ezra Rubenstein

ALL RIGHTS RESERVED

ABSTRACT OF THE DISSERTATION

Essays on Metaphysical Explanation

by EZRA ALEXANDER RUBENSTEIN

Dissertation Directors:

Jonathan Schaffer and Theodore Sider

This dissertation comprises four stand-alone essays unified by a concern with *metaphysical explanation*: that is, with explanations which situate non-fundamental phenomena with respect to what is going on more fundamentally.

The first essay, ‘Grounded Shadows, Groundless Ghosts’, discusses the prospects for explaining the four-dimensional manifest image in terms of an underlying reality with many more dimensions. I argue that this proposal does not face any special ‘explanatory gap’ concerning scrutability or arbitrariness.

The second essay, ‘Two Approaches to Metaphysical Explanation’, clarifies and defends the importance of the distinction between a broadly ‘worldly’ approach to metaphysical explanation, involving the ‘generation’ of facts and other entities, and a broadly ‘representational’ approach, involving the ‘reduction’ of truths and their constituent notions.

The third and fourth essays apply this distinction to two cases. ‘Generalism Without Generation’ argues in favour of a representational approach to explaining particular truths about individuals in terms of purely general truths. And ‘Against Grounding Physicalism’ argues in favour of a representational approach to explaining truths about consciousness in terms of broadly physical truths.

Acknowledgements

First and foremost, I would like to thank my two advisors, Jonathan Schaffer and Ted Sider. I have benefitted immensely from their guidance over the course of my time as a PhD student: my development as a philosopher has been shaped hugely by their generous encouragement, by their own work and ideas, and by their supportive yet challenging feedback in many hours of discussion and many pages of notes. I would also like to thank the other two members of my dissertation committee, Karen Bennett and Alex Skiles, who have consistently given me support, astute advice, and extremely insightful comments. All the members of my committee have pushed me to think and write with greater clarity and rigor, without stifling my half-formed ideas.

I am incredibly fortunate to have ended up in a rich philosophical community at Rutgers and in New York, where I have had the opportunity to learn from a plethora of exceptional interlocutors. For exchanges which have helped shape the ideas in this dissertation, special thanks to David Albert, Nathaniel Baron-Schmitt, David Builes, Eddy Chen, Cian Dorr, Augie Faller, Carolina Flores, Christopher Frugé, Zach Kofi, Barry Loewer, Brian McLaughlin, Jill North, Andrew Rubner, Trevor Teitel, Nick Tourville, Isaac Wilhelm, and Dean Zimmerman, as well as the participants in the Rutgers Metaphysics Group between 2016 and 2022, the 2018 Rutgers Third Year Seminar, the 2020 Rutgers Dissertation Seminar, and the 2022 NYU Mind Group. Many more people have made significant contributions to the writing of this dissertation than I have been able to name here: here's to all those who remain unnamed.

For patiently setting me on my philosophical journey when I had more ideas than sense, heartfelt thanks to Mike Inwood, Edward Kanterian, Anil Gomes, Nick Jones, Lee Walters, Martin Pickup, and Martin Davies. My warmest thanks to Lisette Varón Carvajal, Emmet von Stackelberg, Danny Forman, Adam Gibbons, Morgan Moyer, and Joshua Anthony, for helping me to feel at home in Highland Park, and to Nina Sobell, for being my family in New York. My deepest debt is to all of my family, and especially my parents, Anne Bean and Peter Fink, who have been infinite sources of love and inspiration. Finally, for being my constant companion, in philosophy and in life, I will always be incredibly grateful to Verónica Gómez Sánchez. She has left her mark on this dissertation in so many ways: my work would be much less valuable without her probing insights and her affectionate support.

The first chapter of this dissertation has previously been published elsewhere:

Rubenstein, Ezra (2022). Grounded Shadows, Groundless Ghosts. *British Journal for the Philosophy of Science* 73 (3):723-750.

Table of Contents

<i>Abstract of the Dissertation</i>	<i>ii</i>
<i>Acknowledgements</i>	<i>iv</i>
1. Introduction	1
2. Grounded Shadows, Groundless Ghosts	5
2.1. Introduction	5
2.2. Grounding Connections	12
2.3. Grounded Shadows	16
2.4. Scrutability	20
2.4.1 Triviality.....	21
2.4.2 Transparency	25
2.5. Arbitrariness	31
2.5.1 Other-dimensional ghosts.....	33
2.5.2 Displaced ghosts.....	37
2.5.3 Alternative local beable ghosts	43
2.6. Conspiracy Theories	46
3. Two Approaches to Metaphysical Explanation	51
3.1. Introduction	51
3.2. The Two Approaches	53
3.2.1 Facts vs. truths.....	55
3.2.2 Generation vs. reduction	59

3.3. Significance	67
3.3.1 Ontology.....	67
3.3.2 Meta-ontology.....	70
3.3.3 Realism.....	75
3.4. Adjudication	79
3.4.1 Explanatory targets.....	80
3.4.2 Explanatory goodness.....	84
3.5. Redundancy	89
3.5.1 Against pure generation.....	89
3.5.2 Against pure reduction.....	94
3.6. Conclusion	100
4. Generalism Without Generation	103
4.1. Introduction	103
4.2. Two Versions of Generalism	107
4.3. Against Permissive Generalism	111
4.3.1 Systematicity.....	111
4.3.2 Arbitrariness.....	116
4.4. Strict Generalism	120
4.4.1 The proposal.....	122
4.4.2 Progress.....	128
4.5. Remaining Issues	134
4.5.1 Proportionality.....	134
4.5.2 Singular thought.....	137
4.6. Concluding Remarks	139

5. Against Grounding Physicalism.....	141
5.1 Introduction.....	141
5.2 Grounding Physicalism	145
5.3 Against Grounding Physicalism	148
5.3.1 From grounding to covering.....	149
5.3.2 Against physical-phenomenal laws	156
5.3.3 From covering without physical-phenomenal laws to identity	161
5.4 Responses	163
5.4.1 Derivative laws.....	163
5.4.2 Essences	170
5.4.3 Physicalist panpsychism.....	173
5.4.4 Two-factor physicalism.....	176
5.5 Hybrid Physicalism.....	180
5.5.1 Motivation	180
5.5.2 Systematicity	183
5.6 Concluding Remarks: The Explanatory Gap	190
References	193

1. Introduction

A good deal of philosophical inquiry stems from the suspicion that certain phenomena which are central to our conception of reality are nonetheless derivative, superficial, or even illusory. Think of ordinary objects, causation, colour, consciousness, free will... Which ‘deeper’ aspects of the world could underlie these phenomena? To ascend semantically: what could support the truth, or appearance as true, of the claims we make about them?

These questions are demands for ‘metaphysical explanation’: we seek a theory which situates some target phenomenon with respect to what is going on ‘more fundamentally’. Perhaps ordinary objects are to be explained in terms of the interactions between huge numbers of tiny particles, causation in terms of patterns of counterfactual dependence, colour in terms of the dispositions of surfaces to absorb and reflect light, and so on. This dissertation comprises four essays connected by a concern with metaphysical explanation in this broad sense.

The first essay, ‘Grounded Shadows, Groundless Ghosts’, considers a dispute which arises within the metaphysical interpretation of quantum mechanics. A central piece of mathematical apparatus in quantum mechanics — the ‘wave function’ — appears, if taken at face value, to represent a fundamental field on an extremely high-dimensional space. This motivates a radical project of re-imagining the familiar three-dimensional world of tables and chairs in terms of an underlying high-dimensional reality (e.g Albert 1996, Ney 2015). Several objections to this project allege that there is an in-principle explanatory gap

between high-dimensional reality and the three-dimensional manifest image: rather than being implicit in the physical theory itself, the needed explanatory connections must be tacked on in a non-transparent and arbitrary way (e.g. Maudlin 2010, Allori 2013a). I argue that these concerns are misplaced: even if the world were fundamentally three-dimensional, the same kind of explanatory connections would be needed.

The second essay, ‘Two Approaches to Metaphysical Explanation’, presents a distinction between a broadly ‘worldly’ approach to metaphysical explanation, involving the ‘generation’ of facts and other entities (e.g. Schaffer 2009, Bennett 2017), and a broadly ‘representational’ approach, involving the ‘reduction’ of truths and their constituent notions (e.g. Sider 2011, Dorr 2016). This essay aims to clearly articulate this distinction, to demonstrate its importance, to provide some resources for adjudicating between the approaches, and to argue that both have a role to play in the overall project of explanatory metaphysics.

The third and fourth essays illustrate the impact of this distinction in two cases: the view that individuals are nothing over and above the qualitative goings-on, and the view that conscious experiences are nothing over and above the physical goings-on. In each case, I argue that a representational approach carries a distinctive advantage.

A notable feature of our best physical theories is that they are purely qualitative: permuting the individuals whilst preserving the overall ‘pattern’ of qualitative properties and relations makes no physical difference. Generalism takes this symmetry to be a guide to the fundamental: on this view, there are fundamentally no individuals (e.g. Dasgupta 2009). The explanatory task is to conjure the particular out of the general. In ‘Generalism without

Generation’, I argue that the worldly approach to this task faces problems having to do with systematicity and arbitrariness, which may be avoided by switching to a representational approach.

Finally, the striking success of physics has led many to accept that everything — including even our ‘inner sensations’ — may ultimately be explained in physical terms. However, in light of the general issue of multiple realizability and special issues pertaining to consciousness, the view that conscious experiences may be identified with phenomena characterized in fundamental physical terms has seemed to some to be implausibly extreme (e.g. Fodor 1974, Schaffer 2013, Pautz 2017). This has motivated a generation approach, which promises to vindicate the idea that consciousness is ‘nothing over and above’ physics without the need for such identifications (e.g. Shoemaker 2007, Wilson 2011, Schaffer 2021). In ‘Against Grounding Physicalism’, I argue that this approach (like dualism) requires an unattractive commitment to compact physical-phenomenal laws. But, I suggest, this problem is avoided by employing a hybrid approach which identifies conscious experiences with high-level physical phenomena that may themselves be generated from, but not reducible to, fundamental physics.

One question lies in the background throughout these essays: what criteria should guide us when choosing between metaphysical explanations? This question is highly abstract, but I hope to have shed some light on it by thinking through various ‘first-order’ disputes about explanatory connections (from high-dimensional to low-dimensional, from general to particular, from physical to phenomenal). One general moral of this dissertation is that the question does not have any univocal answer. Rather, an important prior question must first

be addressed: what kind of claim are we making when we offer a metaphysical explanation — in particular, is it a claim about generation relations within reality or about how our representation latches onto reality? There are different games afoot, being played by different rules.

2. Grounded Shadows, Groundless Ghosts¹

2.1. Introduction

Call the space that the fundamental ontology inhabits ‘the fundamental arena’. It is what David Albert (unpublished-a, p.7) describes as ‘the totality of opportunities for things to be one way or another’.² Specifying the fundamental facts is a matter of specifying everything going on in this fundamental arena.

This fundamental arena is standardly conceived as low-dimensional: the universe consists in some fundamental ontology (perhaps particles or fields) inhabiting four-dimensional spacetime (perhaps with some extra string-theoretic dimensions rolled in). Call this standard view ‘low-dimensionalism’.

According to Albert (1996, 2013, 2015, unpublished-b) and Alyssa Ney (2012, 2013, 2015), however, the success of quantum mechanics suggests a radical rejection of low-dimensionalism.³ They propose that the fundamental arena corresponds to what is, for the low-dimensionalist, the universe’s ‘configuration space’ (together with a temporal

¹ Thanks to David Albert, Verónica Gómez, Barry Loewer, Jill North, Jonathan Schaffer, Ted Sider, Trevor Teitel, Isaac Wilhelm, and the participants of the 2018 Rutgers Third Year Seminar for discussion which has greatly benefitted this paper. Thanks also to three anonymous referees for *The British Journal for the Philosophy of Science* for their helpful comments and suggestions.

² My usage differs slightly from Albert’s (unpublished-a): his ‘fundamental arena’ refers to my fundamental arena *at a time* (my fundamental arena corresponds to his ‘Ur-arena’).

³ See also (Loewer 1996; North 2013; Ismael unpublished).

dimension). This arena has $3N+1$ dimensions, where N is the number of particles.⁴ This allows the wavefunction to be straightforwardly interpreted as a field taking values at the points of this space, evolving through time according to a fundamental dynamical law. Any further fundamental ontology also inhabits this high-dimensional arena. For example, Bohmian mechanics is to be interpreted as positing, in addition to the wavefunction field, a single ‘world-particle’. Call this revisionary view ‘high-dimensionalism’. (It is standardly called ‘wavefunction realism’. However, the name is misleading: opponents may agree that the wavefunction is *real*, whilst denying that its reality requires a high-dimensional fundamental arena – either because it is non-fundamental or because it doesn’t inhabit a high-dimensional arena.)⁵

High-dimensionalism faces an obvious challenge: accounting for the low-dimensional world of our everyday experience. ‘The particularly urgent question’, as Albert (2013: 54) puts it, ‘is where, in this picture, all the tables, and chairs, and buildings, and people are’. As Chen (2019: 6) emphasizes, at stake in this question is not just high-dimensionalist theories’ ability to accommodate our common-sense conception of the world, but their very empirical coherence, given that our evidence for quantum mechanics consists in readings on macroscopic low-dimensional instruments. Meeting this challenge requires an explanation, in terms of the high-dimensionalist’s fundamental ontology, for the ‘manifest image’ (by which I mean the collection of ordinary truths involving apparently low-

⁴ This is merely a heuristic: the number of dimensions is supposed to determine the number of particles, not *vice versa*.

⁵ There is also a hybrid view which envisions both a low-dimensional and a high-dimensional fundamental arena – see Dorr unpublished.

dimensional objects – such as there being a table at such-and-such a location at such-and-such a time, and the pointer pointing a certain direction at the conclusion of the experiment.)

Now, given what low-dimensionalists themselves are capable of offering, it would be unreasonable to demand that high-dimensionalist explanations of the manifest image be explicit, detailed and complete. After all, the application-conditions for ordinary notions like ‘table’ are vague and complex; the best anyone can do is provide a sketchy story in terms of, say, there being some particles arranged ‘table-wise’, whilst justifying the viability of an account along these lines.⁶ This, then, is all we can reasonably require of high-dimensionalism.

Can high-dimensionalism provide some such sketchy account – on a par with the low-dimensionalist’s – and justify its viability? Several critics – notably Tim Maudlin (2007, 2010, 2019) – have claimed that there are in principle barriers to this project succeeding.⁷ In particular, these critics worry that high-dimensionalist theories are unable to exploit the low-dimensionalist’s familiar explanatory scheme, in which the distribution of certain ‘primitive ontology’⁸ – fundamental spatiotemporal entities such as particles or fields – determines ordinary truths involving macroscopic objects like tables. As Allori (2013-a: 69) puts it:

⁶ See Sider 2011: §7.6 on ‘toy’ metaphysical truth-conditions.

⁷ See also Allori 2013-a, 2013-b; Chen 2017; Hawthorne 2010; Monton 2002, 2006.

⁸ For this terminology, see Allori 2013-a.

[T]he concern with [high-dimensionalist] theories is that because the wave function lives on configuration space and not three-dimensional space, the explanatory scheme developed in classical theories in terms of a primitive ontology must be drastically revised. A new explanatory scheme is needed, and nobody has found one yet. Hence, [high-dimensionalist theories] at present are not satisfactory.

As I describe in §2.3, high-dimensionalists (most explicitly, Albert (2015)) have offered an explanatory scheme, involving ‘bridge principles’, or as I will be calling them, ‘connections’, which link the distribution of their high-dimensional fundamental ontology to the manifest image. However, critics have deemed this approach unsatisfactory.

Hawthorne (2010: 149) articulates the worry thus:

If you claim that certain bridge principles are true but you can’t see for the life of you, no matter how much you look, why they are true while certain competing principles are false, then there’s going to be an uncloseable explanatory gap. It is at least a significant cost to a theory if that is the upshot.

Meanwhile, Maudlin (2010: 137) seems to go further:

I think we would do well to abjure all talk of ‘links’ or ‘rules’ at all: a physical theory should posit a physical ontology and a dynamics, and the rest should be a matter of *what is comprehensible in terms of that ontology*. If something is *not* easily comprehensible ... then what is called for is either argument or new physical postulates, not just a *rule* or a *link*.

And Chen (2017: 351) echoes these concerns, concluding that ‘we are right to doubt whether there can be any principled way to close the apparent explanatory gap’.⁹

⁹ Hawthorne (2010: 147) provides a vivid illustration of this type of concern. Imagine someone claims that the world consists fundamentally in real numbers and sets of real numbers, and proceeds to explain the manifest image by connecting objects to numbers and properties to sets in such a way that an object instantiates a property just in case the corresponding number is a member of the corresponding set. Such a theory could be made empirically adequate, but surely it would seem absurd. The challenge for high-dimensionalists lies in saying what is different about their view.

Two main ideas underlie these criticisms. The first is that connections – insofar as they are needed at all – ought to be ‘scrutable’: roughly, the distribution of the fundamental ontology ought to render the manifest image ‘easily comprehensible’, as Maudlin puts it. (I will consider two alternative precisifications of this idea below.) The second is that connections shouldn’t be ‘arbitrary’ – there should be something which privileges the selected principles over alternative candidates. High-dimensionalist explanations of the manifest image are allegedly unsatisfactory since their connections do badly by these two criteria.

High-dimensionalist accounts of the manifest image are likely to strike anyone as radical and bizarre, so these concerns clearly carry a strong intuitive pull. The debate has (understandably) tended to focus on the allegedly problematic features of high-dimensionalist connections. However, it has not carefully examined the connections that low-dimensionalists themselves require. I will argue that there is no metaphysically significant difference between high- and low-dimensionalist explanatory schemes: each involves substantive connections between fundamental and non-fundamental. In particular, the consideration of scrutability and arbitrariness ultimately yields no reason to favour low-dimensionalism. In §2.4, I argue that the low-dimensionalist’s connections are also somewhat inscrutable – and what contrast in scrutability there is merely reflects our own conceptual schemes in a way that makes it an unsuitable basis for metaphysical theorizing. In §2.5, I argue that the arbitrariness of high-dimensionalist connections has been overstated, and any genuine arbitrariness afflicts low-dimensionalist connections too. The

upshot is that the idea that there is some important dividing line between high- and low-dimensionalist explanations of the manifest image appears unmotivated.

Considering the various criticisms of the high-dimensionalist explanatory scheme systematically allows us to diagnose them as stemming from a common source. As I see it, this source is an approach on which metaphysics largely consists in understanding what is already implicit in our best physical theories, rather than being a matter of substantive theorizing in its own right. As I will argue, this approach is importantly misleading. In particular, even in familiar cases, the connections between the theory's fundamental description of the world and non-fundamental ordinary truths are theoretical posits about the structure of reality. As such, they ought to be judged as one key constituent of an overall theory, to be evaluated as a complete package according to the standard norms of scientific theorizing. Explanations of the manifest image live and die by the overall theories they are a part of.

In this vein, I intend to shift the debate away from its current focus on the question of whether high-dimensionalist explanations of the manifest image are satisfactory, and onto the question: is there all-things-considered reason to prefer low-dimensionalist to high-dimensionalist theories?

This paper aims to contribute to, but certainly not comprehensively address, this latter question. I argue that the currently most prominent and influential objection to high-dimensionalism – that it is unable to satisfactorily explain the manifest image – is based on

a faulty conception of such explanations, and ultimately provides no reason to favour low-dimensionalism. However, there are several other important considerations – pertaining both to the explanation of the manifest image and to the other components of theorizing – which are beyond my scope here. For example, Albert (unpublished-b) and Ismael (unpublished) argue that high-dimensionalism provides a compellingly natural and elegant explanation of paradigmatic quantum phenomena such as entanglement, whilst others have argued that high-dimensionalism cannot explain various striking features of the dynamical laws, such as their symmetries.¹⁰ I cannot discuss these arguments here, but I think they should occupy the heart of the future debate.

One other limitation of this paper is worth mentioning at the outset. I am assuming a ‘realist’ rather than ‘instrumentalist’ approach to quantum mechanics, according to which it is a guide not merely to prediction but to the true fundamental theory of the world. It is in this context that the debate between high- and low-dimensionalism takes place.

I begin by describing the general structure of ‘fundamental theories’, and briefly presenting the high-dimensionalist account of the manifest image. I then go on to consider – and reject – the reasons that have been given for deeming this account unsatisfactory.

¹⁰ See, for example, Lewis 2004; Allori 2013-b; Maudlin 2013; Chen 2017; Gao 2017.

2.2. Grounding Connections

Fundamental theories are usefully divided into three parts. First, they posit some fundamental ontology – the ‘fundamentals’ – inhabiting or belonging to a fundamental arena. The fundamental ontology standardly consists in some fundamental objects (such as particles or spacetime points) and some sparse list of fundamental properties and relations that these objects instantiate (such as distances, masses or field-values).

Second, they posit some fundamental laws – the ‘dynamics’ – which systematize or govern the distribution of the fundamental ontology across the fundamental arena. These standardly constrain the evolution of this ontology through time (hence the name), although we needn’t assume they must take this form.

Third, they posit some bridge principles – the ‘connections’ – linking the fundamentals to the ordinary phenomena which the theory aims to explain. These principles take the fundamental ontology (or facts about this ontology) as input, and yield derivative ontology (or facts about this ontology) as output. They may state, for example, that when the particles are distributed in such-and-such a way, there is a table at such-and-such a location.¹¹ Thus, they support metaphysical explanations of the manifest image in terms of the fundamentals.

¹¹ You might think that fundamental theories, being couched in exclusively fundamental terms, shouldn’t mention anything like *tables* (along the lines of Sider’s (2011) ‘Purity’). It needn’t matter here whether the connections between the fundamentals and ordinary truths really belong to the fundamental theories themselves, so long as you acknowledge their essential role in the explanatory task we take fundamental theorizing to be engaged in.

Although rarely an explicit part of theorizing in physics itself, these connections are a crucial component in fundamental theories. Our empirical evidence consists in seemingly derivative, macroscopic phenomena like pointer positions (or, on a stricter conception, our own experiences of those positions). Thus, a fundamental physical theory is unable to explain (or, indeed, even predict) this evidence without some (perhaps implicit) connections between its fundamentals and the non-fundamental. As Maudlin (quoted in Saunders et al 2010: 176) forcefully points out, if the theory merely tells us about the behaviour of its fundamentals, we simply can't have any empirical confirmation for it unless we know how these fundamentals connect up to the macroscopic phenomena that we make evidential contact with. For example, suppose we are told that, fundamentally, there are some particles inside the box, and that the dynamics makes it very likely that these particles will quickly spread out to fill the entire box. We have no way whatsoever to test this theory unless we assume something about what these particles give rise to, non-fundamentally – a gas, say, of the sort that we can detect by its colour or smell.¹²

¹² It might be objected that we can make predictions in quantum mechanics without positing any connections, simply by using the Born rule, which relates the wavefunction of a given piece of experimental apparatus to the probabilities of experimental outcomes obtained using that apparatus. The Born rule is not plausibly regarded as a 'connection' in the sense defined above; at best, if we take the wavefunction in question to be fundamental, the Born rule connects some fundamental ontology to the probabilities of certain non-fundamental truths obtaining. However, in the present realist context, we are considering theories which specify some fundamental ontology (which may or may not include the wavefunction itself) and some dynamics pertaining to that ontology. We cannot extract predictions from such theories using the Born rule whilst staying neutral on the connections between this fundamental ontology and macroscopic phenomena. Consider, for example, the Born rule inference that we should expect (with 100% certainty) to find an alive cat when we open the box. This inference can only be licensed if we assume that the fundamental ontology pertaining to the box actually gives rise to an alive cat. For the assumption that it gives rise to anything else – a *dead* cat, say, or, for that matter, a *giraffe* – would be *incompatible* with this inference. (Thanks to a referee for pushing me to clarify this point.)

For the sake of definiteness, I propose to adopt the ideology of ‘grounding’ as a way of conceptualising these explanatory connections (hence, I will sometimes refer to them as ‘grounding connections’, and I will talk of fundamental ontology/facts grounding derivative ontology/facts.)¹³ I will not be relying on any of the more controversial claims about grounding in what follows, and I leave it to the reader to translate my discussion into their own preferred terms.¹⁴ However, one feature of this ideology deserves emphasis: grounded ontology/facts are non-fundamental but can nonetheless be perfectly *real* (as opposed to fictional or illusory.) For example, it is commonly supposed that chemical ontology, such as hydrogen atoms, is grounded in microphysical ontology, such as electrons orbiting protons. This view does not imply any kind of anti-realism about chemistry.¹⁵

This is worth emphasizing because much of the debate over high-dimensionalism has followed Albert’s (1996: 277) influential early presentation of the view, according to which the impression that we live in a low-dimensional space is ‘somehow flatly illusory’. This has invited several objections. For example, Monton (2006: 784) objects that high-

¹³ This ideology is not standard in the existing debate, but it has become an influential way of conceptualizing the explanatory connections between fundamental and non-fundamental elsewhere. See Fine 2012; Rosen 2010; Schaffer 2017.

¹⁴ For example, I will not be assuming that there is a single unified relation of ‘big-G’ *Grounding* (Wilson 2014), that grounds metaphysically necessitate the grounded, that grounding is transitive, *etc.*

¹⁵ Some metaphysicians have argued that only the fundamental is real. For example, they think that, strictly speaking, there are no tables. As I discuss below (§2.6), the truth of ordinary platitudes often turns on subtle meta-semantic questions. However, these issues are orthogonal to the debate between low- and high-dimensionalism: both require that their fundamentals underlie, in some sense, the manifest image. Whatever role the low-dimensionalist thinks that these fundamentals play in accounting for tables (or talk of tables), the high-dimensionalist can take the same role to be played by their fundamentals.

dimensionalism is ‘even more radical than the brain-in-the-vat scenario’, Maudlin (2007: 3166) objects that ‘we cannot appeal to mere fictions’ to explain our empirical evidence, and Allori (2013-b: §7.6) objects that high-dimensionality – since it makes essential appeal to experience – requires a solution to the mind-body problem. Understanding high-dimensionality as the claim that spacetime and its occupants are grounded in high-dimensional fundamentals avoids such objections; high-dimensionality is compatible with the manifest image being just as real as it is for the low-dimensionalist.¹⁶

The nature of the required grounding connections is central to the debate between high- and low-dimensionality. The key question is: by what criteria ought we evaluate this part of theorizing, and do high-dimensionality connections do worse than low-dimensionality connections by these criteria?

Positing grounding connections certainly isn’t a free-for-all. For one thing, it is subject to evaluation by ordinary theoretical virtues – most obviously: empirical adequacy, and whatever makes for overall explanatory goodness. Since high-dimensionality versions of quantum theories replicate the predictions of their low-dimensionality counterparts, empirical adequacy will not distinguish them. As for explanatory goodness, I will not attempt any comprehensive evaluation. Rather, I will be addressing the two main reasons that emerge from existing criticisms of high-dimensionality for thinking that high-

¹⁶ On the aptness of grounding in this context, see Ney 2013: 180 and North 2013: 198. Albert (unpublished-a, unpublished-b) seems to have moved away from the anti-realist conception of low-dimensional space in more recent work.

dimensionalist connections do not provide satisfactory explanations of the manifest image. But first, I should briefly present these connections.

2.3. Grounded Shadows

According to high-dimensionalism, the ordinary world is constituted by ‘shadows’ of the fundamental ontology, projected onto a low-dimensional space. The latter is a derivative space distinct but constructed from – or, as I will say, grounded in – the fundamental arena. Thus, high-dimensionalist connections should be understood as describing the construction of a ‘new’ space out of the fundamental arena, rather than hooking up the fundamental arena to some ‘pre-existing’ space.¹⁷

The key to this construction is that each point in the fundamental arena is posited to ground an N-tuple of spacetime points – mimicking the mapping between points of $3N$ -dimensional configuration-space and the N particle locations in their corresponding configurations. Given some coordinatization of spacetime (x, y, z, t) , we can represent this correspondence perspicuously by coordinatizing the fundamental arena in such a way that the point which maps to the N-tuple of spacetime points $\langle x_1, x_2, x_3, t \rangle, \dots, \langle x_{3i-2}, x_{3i-1}, x_{3i}, t \rangle, \dots, \langle x_{3N-2}, x_{3N-1}, x_{3N}, t \rangle$ is coordinatized as $\langle x_1, \dots, x_{3N}, t \rangle$.¹⁸

¹⁷ This becomes important in the discussion of arbitrariness below (especially §2.5.2).

¹⁸ I am glossing over complications regarding the relativistic extension of quantum theories here, in the spirit of assuming that they won’t affect the core metaphysical issue that I am concerned with. Whether these complications ultimately impact the debate between high- and low-dimensionalism will have to be left for future discussion.

The resulting ‘configuration-space mapping’ from points of the fundamental arena to N -tuples of spacetime points can be used to define a converse mapping, G , from spacetime points to corresponding regions of the fundamental arena. G maps the spacetime point $\langle a, b, c, t \rangle$ to the region containing all points of the fundamental arena of the form $(\dots x_{3i-2} = a, x_{3i-1} = b, x_{3i} = c, \dots, t)$. Intuitively, G maps each space-time point p to all those points of the fundamental arena whose corresponding configurations involve a particle occupying p . The resulting region of the fundamental arena is a fusion of N $3N-3$ -dimensional hyperplanes; the i -th of these hyperplanes corresponds (heuristically) to all the configurations in which the i -th particle occupies p .

The various high-dimensionalist proposals exploit this mapping to define projections of the fundamental ontology onto the newly constructed derivative space. Albert (2015: ch.6) describes four such projections: one for Bohmian mechanics and three for ‘GRW’ (named for Ghirardi, Rimini & Weber (1986)).¹⁹ My interest is in the general structure shared by these proposals rather than their particular details. Nonetheless, to have a concrete target in the arguments that follow, it will help to introduce one of them.

In the case of GRW, unlike Bohmian mechanics, the wavefunction is not supplemented by any additional fundamental ontology. As Bell (1987: 204) puts it: ‘It is in the wavefunction that we must find an image of the physical world, and in particular of the arrangement of things in ordinary three-dimensional space.’

¹⁹ See also Bell 1987: ch.22. Although the same issues arise for Many Worlds interpretations (see Saunders et al 2010: Part II), I don’t consider the extension directly here as it involves extraneous difficulties – for example, in making sense of probabilities (see Albert 2015: ch.8)).

One natural approach is to posit a non-fundamental ‘mass-density’ field on spacetime. This field can be imagined as being formed by superimposed particle configurations, with the ‘density’ of each configuration in the resulting spatiotemporal image being determined by the wavefunction’s amplitude at the corresponding point of the fundamental arena. This layering of particle configurations amounts to a pointwise ‘projection’ of what is going on in the fundamental arena onto spacetime. More formally, the mass-density at any given spacetime point p is determined by integrating the wavefunction’s squared-amplitude across the N $3N-3$ -dimensional hyperplanes that constitute the region $G(p)$.

This describes a relatively simple and mathematically precise rule for going from the wavefunction in the fundamental arena to mass-density in derivative spacetime. The requisite ordinary objects are then formed out of the various clumpings of this mass-density, arranged in a suitably intricate manner. Tables, for instance, are table-shaped clumps of high mass-density.


Albert’s other proposals all share the same two-step structure as the mass-density proposal.

Firstly, they posit a precise and systematic connection between their fundamental high-dimensional ontology and some derivative spatiotemporal ontology capable of underlying ordinary objects, such as the mass-density field – what Bell (1987: ch.7) terms ‘local beables’. Each of these connections exploits the mapping G described above, extended in the natural way to take in arbitrary regions, so that what is going on vis-à-vis the local


beables at spacetime region R is metaphysically determined by what is going on at the region of the fundamental arena $G(R)$.

Secondly, they posit some further connections between the local beables and ordinary objects, replicating the connections posited by their counterpart low-dimensionalist theory.

For example, suppose this latter theory posits a connection of the form:



The high-dimensionalist simply replicates this connection (the second step), whilst supplementing it (the first step) with a further connection of the form:



Here, ‘table-wise*’ is a place-holder standing for however it is that the high-dimensional ontology has to be arranged at $G(R)$ to give rise to local beables arranged table-wise at R , given the posited connection between the high-dimensional ontology and the local beables. Since this latter connection is made perfectly precise on each proposal, ‘table-wise*’ will be no sketchier than ‘table-wise’ is already. Any precisification of what ‘table-wise’

amounts to in terms of the low-dimensionalist's preferred local beables will automatically induce a corresponding precisification of 'table-wise*' in terms of the high-dimensionalist's own fundamentals.

As a mathematical matter, the posited connections in all these proposals will, by design, replicate the predictions of their low-dimensionalist counterparts. Whenever the low-dimensionalist says that the local beables ground a pointer pointing up, say, the high-dimensionalist agrees, merely adding that those local beables are in turn grounded by some fundamental high-dimensional ontology.

But, as already mentioned, empirical adequacy alone seems insufficient to fulfil high-dimensionalism's obligation to explain the manifest image; the mathematical operations being described by the high-dimensionalist must represent satisfactory grounding connections. Two main reasons for thinking that they do not have been suggested by existing criticisms: firstly, the resulting connections seem 'inscrutable', and secondly, they seem 'arbitrary'. I will examine these concerns in turn, arguing that neither provides good reason to prefer low-dimensionalism.

2.4. Scrutability

Several criticisms of high-dimensionalist connections concern their apparent lack of 'scrutability', in contrast to their low-dimensionalist rivals. According to this idea, the emergence of a pointer pointing up from some particle arrangement is somehow

straightforward, innocent and unmysterious, unlike its emergence from the undulations of a wavefunction in some distinct high-dimensional space, say. For example, Allori (2013: 65) alleges that ‘once the primitive ontology and its temporal evolution are given, everything else follows’, and Maudlin (2007: 3160) writes that positing fundamental local beables ‘makes the connection between the theoretical picture and the world as we pre-theoretically take it to be transparent’.

It is worth distinguishing two ways of fleshing out what this apparent asymmetry in scrutability amounts to – I label them ‘triviality’ and ‘transparency’. In each case, I argue, the extent to which low-dimensionalist connections themselves are scrutable has been overstated, and what genuine asymmetry there is turns out to be merely conceptual, in the sense that it reflects features of the concepts with which connections are represented. Thus, scrutability ultimately provides no reason to prefer low-dimensionalist connections.

2.4.1 Triviality

Maudlin (2010) argues that the emergence of pointer positions from Bohmian particle configurations is easy to understand. He (2010: 123) observes that a ‘perspicuous’ visual representation of different Bohmian configurations makes it obvious which corresponds to which pointer position ‘without further interpretive machinery being added’, concluding that:

There is a certain obvious sense in which a world described by [low-dimensionalist theories] comprehensibly corresponds to the world as we experience it. If the notion of the emergence of a low-dimensional spacetime with localized objects from a high-dimensional reality can be made equally

comprehensible, then [high-dimensionalist theories] will have passed one hurdle. (Maudlin 2010: 142)

Meanwhile, Allori (2013-a, 2013-b) repeatedly emphasizes the apparent smoothness with which low-dimensionalist theories explain the manifest image. According to Allori, every macroscopic property ‘just “arises” from’ (2013-b: 20) or ‘can be appropriately “read off” from’ (2013-b: 15) the history of the primitive ontology, allowing us to ‘directly compare’ (2013-a: 66) its macroscopic behaviour to the manifest image. In Maudlin’s (2007: 3167) metaphor, the manifest image is to be revealed merely by ‘squinting’ at low-dimensionalist fundamentals.

To be sure, the concern being gestured at here remains somewhat vague. Nonetheless, enough has been said, I think, to capture a distinctive criterion that acceptable grounding connections are being alleged to satisfy – let’s label this criterion ‘triviality’. We can gloss triviality thus: a grounding connection is trivial if the grounds render the grounded easily comprehensible or intuitively visualizable, with no ‘metaphysical codebreaking’ required. Through trivial connections, the emergence of the manifest image from the fundamentals is natural and obvious – indeed, the explicit specification of such connections is unnecessary. (As we will see, triviality comes apart from the more precise criterion of transparency.)

It seems clear that the proposed connections between high-dimensional fundamental ontology and the manifest image are not trivial in the relevant sense. The wavefunction just isn’t the sort of thing which bears squinting at; we are unable to visually represent such

high-dimensional ontology at all, let alone in a way which allows the manifest image to seamlessly emerge. And the alleged emergence of familiar three-dimensional reality from such ontology is certainly not intuitive or easy to comprehend.

Two questions remain: are the low-dimensionalist connections themselves trivial in the relevant sense? And if so, does this give us good reason to prefer them?

As for the first question, I doubt that low-dimensionalist grounding connections are really as trivial as these critics suggest. Even in Maudlin's simple case of pointer positions arising from Bohmian configurations, there must actually *be* a pointer – at least, something that is disposed to appear pointer-like. We need a solid, stable object which moves rigidly, is visible, and so on.²⁰ In order to account for such dispositions, it is not enough merely to have a pointer-shaped collection of particles; one needs these particles to be disposed to behave in the right kinds of ways. In this case, the particle dispositions are underwritten by the dynamics, including – crucially – the evolution of the associated wavefunction itself.²¹

A full account of the existence of a pointer, then, requires a low-dimensionalist account of the wavefunction – an account which is hardly likely to be trivial. And even supposing we have been given such an account, surely the fact that complex dynamical properties of the particles are required – including, presumably, those underlying the chemical bonds which

²⁰ As Maudlin (2010: 124) acknowledges, merely having the configurations is insufficient: 'we need at least some schematic reason to think that we could *see* these configurations'.

²¹ Thanks to a referee for pointing this out.

make the pointer into a stable substance – should shatter any illusion that the emergence of the pointer is somehow straightforward or intuitively obvious.

And this is before one considers more complex macro-phenomena such as temperature and colour; understanding such phenomena in microphysical terms is surely a huge intellectual achievement, far from a matter of smooth ‘reading off’ or ‘direct comparison’. One needn’t be an anti-reductionist to doubt that the task of reducing the rich macroscopic features of the manifest image to fundamental ontology will be trivial in the relevant sense.

The triviality criterion is vague enough, however, that some contrast in the vicinity between high- and low-dimensionalist connections might be salvaged. The question is then whether this contrast really gives us any reason to think low-dimensionalist connections more likely to be true. And it seems implausible that it should; any discrepancy in this case surely concerns our grasp of the concepts involved, rather than the connections themselves. Local beables are conceptualized through a direct extension of ordinary concepts like *TABLE*; the concept *PARTICLE*, for instance, is introduced to us as referring to tiny bits of tables and chairs.²² This makes low-dimensionalist connections between local beables and ordinary objects familiar and built into our understanding from the outset. Our grasp of high-dimensional concepts, on the other hand, is more indirect – since we are unable to point to or visualize high-dimensional ontology, we are introduced to it via abstract mathematical formalisms. It ought to be unsurprising that descriptions in terms of concepts acquired in this way fail to trivialize the emergence of the manifest image.

²² I denote concepts in *CAPS*.

Using such discrepancies in our access to the relevant concepts as a basis for fundamental theorizing seems unduly anthropocentric.²³ This theorizing may just be conceptually difficult – even intractable – for beings like us. If anything, supposing that the fundamental ontology should turn out to be adequately described by concepts as naïve and directly graspable as *PARTICLE* seems suspiciously impatient. It's possible that the world we inhabit is kind enough to bottom out in 'primitive' ontology like that, but this sort of optimism may just begin to look like blind faith in the face of empirical evidence.

Methodologically speaking, it might be reasonable to try to get by with trivial connections insofar as it proves possible. Perhaps there is even some theoretical benefit in doing so, insofar as triviality correlates with simplicity (whatever that amounts to). But it is hard to see how triviality could be a genuine virtue in its own right: among two equally simple connections, the one which is harder to comprehend or less visualizable or less familiar may thereby be less comforting or convenient for limited cognizers like us – but is it really thereby less plausible?

2.4.2 Transparency

Triviality seems unattractive as an articulation of the intuitive demand that grounding connections be scrutable. But there is an alternative interpretation: transparency. The

²³ This point is illustrated by the 'phenomenal concepts strategy' for defending physicalism in the philosophy of mind (Loar 1990; Hill 1997), according to which the conceivability of zombies (physically identical but non-conscious duplicates of people) reflects certain features of 'phenomenal concepts' – and hence is an inappropriate basis for the metaphysical claim that consciousness is non-physical. One may disagree that this is the right explanation of the conceivability intuitions in this case, but the point remains that *if* a consideration turns out to merely reflect our concepts, then that considerably lessens its weight in metaphysical arguments.

demand that grounding connections be ‘transparent’ has clear precedent in explanatory gap arguments for the non-physicality of phenomenal consciousness (Levine 1983; Chalmers 1996, 2012). These arguments suggest a parallel worry about high-dimensionalism: ordinary ontology could not be grounded in high-dimensional ontology because such connections are ‘opaque’ (non-transparent).

In fact, following Schaffer (2017: 4), ‘three not-obviously-equivalent’ ways of unpacking opacity are familiar from the philosophy of mind literature, involving conceivability, logical possibility, and epistemic a priori. According to the first, a connection between some grounding state and some grounded state is opaque just in case it is conceivable that the former obtains without the latter obtaining. For example, connections between physical states and consciousness are allegedly opaque in this sense since one can conceive of or coherently imagine a zombie-world: a physical duplicate of our world entirely lacking in consciousness (Chalmers 1996). According to the second, a connection is opaque just in case the grounding state’s obtaining does not logically entail the grounded state’s obtaining. And according to the third, a connection is opaque just in case it is impossible to know that the grounded state obtains purely on the basis of a priori reasoning from the knowledge that the grounding state obtains.²⁴

Transparency comes apart from triviality, in both directions. A connection can be transparent without being trivial: for example, the proposition that certain mathematical axioms obtain without certain difficult-to-prove theorems obtaining is not conceivable,

²⁴ See Chalmers 2012 on ‘a priori scrutability’.

logically possible, or a priori open, but the connection is far from easily comprehensible or intuitively visualizable.

Conversely, some mereological connections fail to be transparent, despite seeming trivial. Consider the claim that the whole is grounded in its parts: a triangle, for instance, is grounded in the three lines that compose it. This connection certainly seems to satisfy triviality: indeed, once one visualizes the three lines one thereby also visualizes the triangle. Nonetheless, the claim that there is no triangle – and more generally, no composite objects at all – is conceivable, logically possible, and a priori open. This is evidenced by the consistency of mereological nihilism, according to which there are only mereologically simple atoms. So the connection between parts and wholes is opaque on any of the three precisifications.

Schaffer (2017) uses this observation as the basis for an interesting defence of physicalism against explanatory gap worries. He argues that transparency cannot be a constraint on grounding connections, since there are ‘gaps’ all over the place – even in familiar and apparently unproblematic cases like mereology. These gaps are bridged by substantive, opaque connections.

As Schaffer (2017: 5) notes in passing, this point has application to the debate in quantum metaphysics too. For example, the low-dimensionalist might wish to ground the fact that some pointer is pointing up at time t in a certain particle configuration at t . But such a connection would be opaque since it is conceivable/logically possible/a priori open that

this particle configuration fails to yield the existence of any pointer (as nihilism claims), let alone one that is pointing up. Perhaps low-dimensionalist connections like these are in some salvageable sense more ‘trivial’ than their high-dimensionalist counterparts, but any asymmetry in transparency is only apparent. If we are open to such connections (as the low-dimensionalist ought to be) then we have already bought into opacity.

Now, I suspect the low-dimensionalist may be tempted to complain that, *pace* its prominent defenders (such as Dorr & Rosen 2002, Sider 2013), the nihilist view that particles fail to compose to yield pointers is not just implausible but incoherent. After all, the low-dimensionalist might say, all we *mean* by ‘the pointer is pointing up’ is nothing more than that the particles are arranged pointer-pointing-up-wise. So, given that the particles are arranged in this way, it *is* inconceivable that there is no pointer pointing up. On the other hand, they continue, it is perfectly conceivable that the wavefunction be arranged any way you like, and it fail to be true that the pointer is pointing up – indeed, the pointer may, for all that has been said, be pointing *down*. So there is, they allege, a genuine asymmetry in transparency after all.

The suggestion that the existence of the pointer might follow analytically from the arrangement of the particles amounts to a controversial deflationism about ontology – albeit one which has its defenders.²⁵ But without entering into this debate, it is instructive to see how the natural ways of interpreting this outlook fail to secure any metaphysically significant asymmetry between low- and high-dimensionalism.

²⁵ For example, Thomasson 2007. See Sider 2011: §§9.7-9.10 for discussion.

Presumably, any such analytic connection would have to follow from our concept *PARTICLE* or our concept *POINTER*. On the latter interpretation, the low-dimensionalist is claiming that *POINTER* is constitutively tied to particles and not wavefunctions. I doubt that our ordinary concepts come so finely opinionated about metaphysics. But besides, if they do, they run the risk of being defective – like *PHLOGISTON* or *ETHER*. The low-dimensionalist cannot have it both ways: either we can start with risky, metaphysically prejudiced concepts, in which case we ought to inquire into whether they succeed in picking anything out; or we can safely assume that there are pointers, and inquire into the nature of the reality that underwrites this assumption.

Alternatively, the low-dimensionalist is alleging a discrepancy between the concepts *PARTICLE* and *WAVEFUNCTION* themselves. Perhaps the former comes with a fixed connection to ordinary ontology like pointers. The latter concept clearly doesn't: as the measurement problem dramatically highlights, it is a matter of debate among *WAVEFUNCTION*-experts whether, and how, certain arrangements of the wavefunction give rise to a pointer pointing up. Granting the controversial claim about the ordinary concept *PARTICLE*, the high-dimensionalist is free to introduce a new concept – call it '*WAVEFUNCTION+*' – for their preferred fundamental ontology, which builds in similar connections (perhaps indirectly, in terms of its connection to some local beables). For example, on the mass-density proposal, it would be analytic to *WAVEFUNCTION+* that the fundamental ontology it picks out gives rise to a mass-density field (and hence to pointers) in the manner described above (§2.3). Again, the question then becomes whether these theoretical posits – wielded by low- and high-dimensionalist alike – actually refer.

In sum, alleging a conceptual discrepancy is merely bulge-shifting. Either, we can agree on connection-neutral concepts (introducing if necessary *PARTICLE*- and *POINTER*-, with any connections explicitly deleted) and investigate whether the posited connections really hold between their referents. Or, we can stipulate concepts which come with these connections preloaded, and investigate whether they hit a worldly target. Ultimately, these are just different glosses on the same metaphysical enquiry. We shouldn't let contingent differences in our actual concepts blind us to the fact that grounding connections, whether built into our concepts or not, are substantive theoretical posits about the world's structure.

This supplements the main point that opaque connections are the norm with an explanation (in terms of the concepts involved) for any apparent asymmetry in transparency. As such, it can be thought of as a 'low-dimensional concepts strategy', analogous to the physicalist's 'phenomenal concepts strategy' (see n.23). The situations in which certain wavefunction arrangements fail to give rise to pointers pointing up may be conceivable despite being metaphysically impossible, due to the incommensurability of high- and low-dimensional concepts (just as, allegedly, zombie-worlds are conceivable despite being metaphysically impossible due to the incommensurability of physical and phenomenal concepts).²⁶

²⁶ To clarify: I am not committed to such a strategy succeeding in the case of phenomenal consciousness. Perhaps there is a *special* kind of metaphysical rather than merely conceptual incommensurability in this case (due, for example, to phenomenal consciousness lacking any essential functional role). But it is hard to see what could be so special in the case of high-dimensionalist connections. (Of course, to fill out the low-dimensional concepts strategy, more would need to be said about how exactly the incommensurability arises in this case.)

2.5. Arbitrariness

So far, I have considered one prominent respect in which high-dimensionalist explanations of the manifest image have been deemed unsatisfactory – the idea that they are insufficiently scrutable – and I have argued that it ultimately provides no reason to prefer low-dimensionalist connections. However, there is another important idea which seems to underlie the criticisms of high-dimensionalist grounding connections: the idea that they are objectionably arbitrary. In this vein, Maudlin (2007: 3166) characterizes these connections as a ‘choice of one out of an infinitude’ of possible alternatives, and Allori (2013-b: 19) claims that ‘there is no deep justification for the additional rules the [high-dimensionalist theories] need. In fact, the answer to the question “Why these rules?” is nothing but “Because they work.”’²⁷

A cluster of important objections along these lines centre around the apparent existence of several kinds of ‘ghost’: alternative constructions from the fundamental ontology that are alleged to have the same credentials to be material objects as Albert’s shadows.²⁸ The question is: what privileges the shadows over the ghosts? Why aren’t the ghosts also material objects? Or if they are, how come we can’t see them or interact with them – how

²⁷ For worries about arbitrariness, see also Gao 2017; Hawthorne 2010; Lewis 2004; Maudlin 2010, 2019; Monton 2002, 2006.

²⁸ I borrow this useful term from Albert (2015: 154), who coins it to describe the low mass-density correlates of high mass-density material objects whose existence forms the basis of one of Maudlin’s (2010: 135) objections. As Albert (2015: 151-2) points out, Maudlin’s key claim that ‘*the density per se does not affect the structural or functional properties of the object*’ is false, given the GRW dynamics; in fact, unlike high-density shadows, low-density ghosts don’t behave anything like material objects.

come, more generally, they don't seem to have the same significance for us that the shadows (allegedly) do?

It is worth getting clear, before proceeding, on what this question amounts to. The credentials in question are, in a broad sense, exclusively *dynamical*: it is a matter of behaving, or being disposed to behave, in certain characteristic ways. Material objects move continuously through the space they inhabit, are relatively stable, interact with each other when they are close enough, do not tend to split or pass through each other, and so on. Grant that high-dimensionalist connections succeed in recovering a dynamical structure that enacts material objects in this sense – the problem is not that it cannot be done, but rather that this kind of formal adequacy is all too cheap! Indeed, the objection goes, there are many other connections which would work just as well.

The challenge, then, is to steer between the horns of arbitrariness and overpopulation. On the one hand, the high-dimensionalist wants to avoid populating our world with many more material objects than we expected. Such overpopulation would seem to involve objectionable redundancy: positing far more table and chair-like objects than we need to explain our ordinary experiences of tables and chairs. On the other hand, these objections run, the high-dimensionalist must avoid arbitrary stipulation; it shouldn't turn out that the preference of the shadows over the ghosts is a brute, *ad hoc* metaphysical posit.

Solving this problem requires justifying one of two policies towards these ghosts: either 'elimination' – they don't exist, or 'discrimination' – they exist, but don't deserve the status

of ordinary material objects. But I won't be arguing for any particular solution. Instead, my strategy is to show that low-dimensionalism faces just the same kinds of ghosts; there is nothing distinctively *high-dimensionalist* about the issue. Thus, the high-dimensionalist can simply replicate whatever policy is adopted by the low-dimensionalist towards their own ghosts, and whatever justification they provide for this policy. I will demonstrate this strategy with three kinds of ghost which have featured in criticisms of high-dimensionalism.

2.5.1 Other-dimensional ghosts

Recall (§2.3) that all the proposed high-dimensionalist connections exploit a characteristic 'configuration-space mapping' between regions of the fundamental arena and regions of four-dimensional spacetime to describe how the goings-on at the former ground the goings-on at the latter. It is natural to wonder what privileges this particular mapping. After all, many other projections from the high-dimensional arena onto different derivative spaces are available. For example, instead of grouping dimensions of the fundamental arena into three N -tuples, as the high-dimensionalist's characteristic mapping does, we could group them into N three-tuples, yielding an $N+1$ -dimensional derivative space. So what distinguishes the *four*-dimensional shadow-world of spacetime from all these other-dimensional ghost-worlds?²⁹

²⁹ Hawthorne (2010: 152-3) raises this issue in the context of Many Worlds.

Moreover, even given a four-dimensional derivative space, there remain myriad alternative constructions to consider. For example, Monton (2002, 2006) and Lewis (2004) consider constructions defined by permuting the dimensions of the fundamental arena, or by translating the contents of the fundamental arena in a given direction. Again, the question is what privileges the high-dimensionalist's particular choice of mapping over these alternatives.³⁰

The answer is the *dynamics*: the four-dimensional derivative space (constructed as explained above) is uniquely privileged as the space inhabited by material objects because of the dynamical laws. Call the space inhabited by material objects moving and interacting in their characteristic ways the 'material space'.³¹ It is the dynamical laws (defined on the fundamental arena) which determine the nature of this material space. In particular, the spatiotemporal relations between points of the material space are identified by the nomic roles that these relations play, given the laws on the underlying fundamental arena. For example, spatial distance is the relation which correlates with the sizes of interactive forces like electromagnetic repulsion and gravitational attraction in characteristic ways; two

³⁰Although, Monton and Lewis' worry that the high-dimensionalist's mapping relies on a preferred coordinatisation of the fundamental arena raises interesting issues which I cannot fully address here. Ultimately, the concern may be that the dynamics required to privilege the high-dimensionalist correspondence is itself implausibly *ad hoc*, given the nature of their fundamental arena. As noted above (§2.1), this is an important issue which is beyond my scope; my concern here is the alleged arbitrariness of the high-dimensionalist grounding connections, granting the high-dimensionalist fundamentals and dynamics. (Lewis himself prefers a view on which it is the intrinsic structure of the fundamental arena – rather than the dynamics – which privileges the correspondence.)

³¹ See Albert (unpublished-a, unpublished-b) for a much fuller presentation of the ideas that follow. Albert refers to this derivative space variously as 'the space of possible *interactive distances*' (1996: 282), an 'emergent geometrical space' (unpublished-a), and 'the space of ordinary material bodies' (unpublished-b). Similarly, Lewis (2013: 123) refers to 'the arena in which *spatial* phenomena play out', noting that 'the term *spatial* is intimately connected to the dynamical laws'.

regions of the material space are close to each other to the extent that they allow for significant interaction between objects occupying them. It is these correlations, between the geometry of the material space and the dynamical interactions of its occupants, which allow for stable objects that bounce off or stick to each other, and which ultimately underwrite the macro-regularities exploited by perceptual systems like our own. In short, these correlations allow the objects within the material space to ‘formally enact’ (in Albert’s phrase) a material world.

The Hamiltonian in Schrödinger’s equation yields the three spatial dimensions of this material space: three orthogonal directions along which certain sorts of interactions change in certain ways – along which, that is, material objects can approach each other. More generally, the dynamics constrains the fundamental ontology in such a way that gives rise to a material world distributed across three spatial dimensions and evolving through one-dimensional time.

This dynamics will not yield material worlds in other-dimensional derivative spaces (or, for that matter, in twisted versions of four-dimensional spacetime.) Insofar as we can make sense of objects inhabiting such spaces at all, they will be passing right through each other, splitting, jumping around, and behaving in all sorts of odd ways. If indeed they exist, they are strange and unnatural, and certainly won’t merit being described as ‘beings walking around’ (as Hawthorne (2010: n.14) imagines) or anything like that.

Crucially, the situation is identical for low-dimensionalist theories. They have their own other-dimensional ghosts to ponder; there will, for example, be myriad flattenings of the four-dimensional arena available too. We could take N particles inhabiting a three-dimensional space and construct, say, $3N$ particles inhabiting a one-dimensional space. And the fundamental arena can be expanded as well as flattened: we could also construct, say, a $3N$ -dimensional space containing a single particle (Albert's 'world-particle').³² Thus, the question arises of what, if anything, *these* constructions correspond to. And it is, in both the high- and low-dimensionalist case, the dynamics of the fundamental arena which privileges four-dimensional spacetime over the available alternatives.

Note, in particular, that the low-dimensionalist cannot privilege four-dimensional constructions merely on the grounds that the fundamental arena is itself four-dimensional.³³ The privilege in question, recall, is that of being the *material objects*. And we can readily imagine dynamical laws characterising fundamental ontology on a four-dimensional fundamental arena which would give rise to worlds in which nothing four-dimensional behaves anything like a material object. Indeed, we can imagine laws according to which it would instead be a certain *two*-dimensional flattening, and not the four-dimensional arena itself, that would seem to earn the status of material space.³⁴

³² Indeed, since low-dimensionalist theories also posit a wavefunction – whether as a law, a multi-field, or something else – all the same richness of constructions from it is (in principle) available to them.

³³ This is contrary to the apparently widespread view that, as Allori (2013-b: 14) puts it, low-dimensionalists 'do not have to explain the *appearance* of three-dimensionality, since the world *is* three-dimensional' (see also Lewis 2016: 163). The fundamental arena's being three-dimensional is not sufficient for the world's *appearing* three-dimensional – what we experience is the material space, as determined by the dynamical laws.

³⁴ For vivid thought-experiments along these lines see Albert unpublished-b.

Why assume that this material space automatically coincides with the fundamental arena? To borrow a metaphor of Albert's, the directions in which material objects can move need not match the directions in which the fundamental story of the world can develop. If anything, once fundamental arena and material space are carefully distinguished, their putative coincidence begins to appear just that: coincidental. And indeed, the core of the high-dimensionalist explanation of quantum weirdness is precisely the claim that they in fact (dramatically) come apart.³⁵

2.5.2 Displaced ghosts

Maudlin (2019: 126) points out that in addition to Albert's shadows, there are hordes of ghosts related to them by spatiotemporal translations. Take the translation which shifts everything 'three feet to the North'. The resulting ghost-world conforms to the same structure – both geometrically and dynamically – as the shadow-world it is constructed from. For example, whenever ghost-billiard balls collide, they bounce off each other. Unlike other-dimensional ghosts, these displaced ghosts do seem to have the dynamic credentials to enact ordinary material objects.

Now, *prima facie*, this is a problem for low-dimensionalism as much as high-dimensionalism: these constructions are available whether the fundamental ontology is low- or high-dimensional. And it can be solved the same way in both cases – by positing a grounding connection describing how the material objects inherit their location from the

³⁵ See Ismael unpublished; Albert unpublished-b.

fundamental ontology. That is, we should not posit that there are tables wherever there is anything ‘playing the table-role’ (where this is a purely dynamical constraint), but rather only where there is an appropriate table-*realizer* (a table-shaped arrangement of particles, clump of high mass-density, or whatever) playing the table-role. High- and low-dimensionalists disagree about what, fundamentally, these table-realizers are – but each is entitled to make use of them in their account of tables.

This is contrary to some misleading suggestions of Albert’s and Ney’s to the effect that high-dimensionalism is (or ought to be) committed to some kind of purely functionalist understanding of material objects, according to which what it is to be such an object is understood in exclusively dynamical terms.³⁶ The question of how to ground ordinary objects is a complex one, and functional roles are likely to play a part on both low- and high-dimensionalist accounts; but both have additional resources to work with in order to specify which of the role-fillers are genuine.

Nonetheless, one might worry that there is an important disanalogy between the connections posited in the two cases. According to low-dimensionalism, pointers are straightforwardly located where the particles are. This seems like a maximally simple and natural inheritance principle. According to high-dimensionalism, on the other hand, pointers are located in some *distinct* space from the fundamental ontology that they ultimately inherit their location from. Positing that tables are located wherever the table-


³⁶ See, for example, Albert 2015: 129; Ney 2012: 545. See also Chen 2017: §2.2; Maudlin 2019: 123-4. This pure functionalism is explicitly renounced by Albert (unpublished-a) in a footnote addressing the issue of displaced ghosts.

shaped clumps of high mass-density are doesn't help – the real issue concerns what determines the locations of the clumps themselves! High-dimensionalist location inheritance seems bound to be radically more abstract and indirect here. The low-dimensionalist's location inheritance principle seems the obvious choice in a way that the high-dimensionalist's does not: since the high-dimensional fundamental ontology inhabits a distinct space, there seems to be no natural connection available, making any choice arbitrary.

This disanalogy, however, is only apparent. Spacetime is indeed a distinct space from the high-dimensionalist's fundamental arena but crucially, it is not *metaphysically* distinct: it is a derivative 'material space' which is itself grounded in – which owes its very existence and nature to – the fundamental arena (together with the dynamical laws pertaining to it). Recall the mapping G from spacetime regions to regions of the fundamental arena exploited by high-dimensionalist connections (§2.3). The region $G(p)$ grounds the spacetime point p ; the topological and geometric relations holding between spacetime points p_1, \dots, p_n are determined by the relations between their corresponding regions $G(p_1), \dots, G(p_n)$. According to high-dimensionalism, then, spacetime points themselves are derivative entities grounded in certain regions of the fundamental arena – namely, corresponding sets of hyperplanes.

This mapping – privileged by the dynamics, as outlined above – describes the construction of a new space out of the fundamental arena (not merely a correspondence between two *pre-existing* spaces). Given this grounding connection, there *is* a natural and obvious

location inheritance principle available to the high-dimensionalist: what is going on at a given point of the derivative space is grounded in what is going on at the region of the fundamental arena which grounds it. This is illustrated by the following diagram:



The straightforward thing to say, that is, about how local beables inherit their location from the fundamental ontology which grounds them is just that their locations are grounded in the locations of that fundamental ontology. For example, the undulations of the wavefunction across hyperplanes of the form $(x_{3i-2} = a, x_{3i-1} = b, x_{3i} = c, T = t)$ ground a mass-density spike derivatively located – just as one should expect – at the spacetime point $\langle a, b, c, t \rangle$. And similarly, the wavefunction’s table-wise* undulations across certain regions of the fundamental arena will ground mass-density arranged table-wise (and hence, a table) at the corresponding derivative region of spacetime.

Indeed, it would seem strange to claim that spacetime points were grounded in their corresponding regions if this wasn’t the case. It would be unnatural and bizarre (perhaps even incoherent) to posit instead that what is going on at a given spacetime point is determined by what is going on in the region of the fundamental arena which grounds the

point three feet to the south – just as it would be unnatural and bizarre for the low-dimensionalist to suppose that what is going on vis-à-vis the derivative ontology at a point is determined by what is going on vis-à-vis the fundamental ontology at a point three feet to the south. The sense that there is something arbitrary or stipulative or *ad hoc* about the high-dimensionalist's connection is illusory; both high- and low-dimensionalism face the problem of displaced ghosts – and both can solve it by positing natural and non-arbitrary location inheritance principles.

You may be thinking by now: perhaps high-dimensionalist connections are natural, but are they really *as* natural as the low-dimensionalist's? After all, the low-dimensionalist doesn't need to mess around with projections or derivative spaces at all – the connection between the locations of the fundamental ontology and the material objects is simply identity! It doesn't get more natural than that.

However, supposing that the locations of ordinary objects could be *identical* to the locations of the fundamental ontology grounding them is far too simplistic. Presumably, for example, the low-dimensionalist's fundamental ontology contains nothing located at a table-shaped region. To find such a region, we must fuse the points occupied by some fundamental particles, or at which certain fundamental field-values are instantiated. And in each case, we must explain what makes the particular collection we are fusing apt to correspond to the table. Indeed, it seems that the kind of location inheritance principle for ordinary objects which the low-dimensionalist implicitly endorses will be of exactly the same kind as that described on behalf of the high-dimensionalist above! For example,

suppose some table-wise arranged particles ground a table. The locations of these particles form a scattered set of points, and the location of the table will, at least on the most straightforward proposal, just be the scattered region grounded in those points. For low- and high-dimensionalist alike, then, the most natural location inheritance principle available is that grounded objects occupy locations which are distinct from but grounded in the locations of their grounds.

Moreover, common sense has it that ordinary objects like tables occupy continuous regions. Compare a cloud of dust-particles which happen momentarily to form a table shape. Intuitively, the cloud occupies a scattered constellation of points. Even if one thinks that, strictly speaking, this turns out to be true of the table too, one should still want to accommodate in some way the (seemingly principled) common-sense distinction between the dust-cloud's location and the table's.³⁷ This requires attending to the dynamical distinction between the dispositions of the dust-particles and the table-particles; the locational contrast is surely informed by the table's being solid and stable, resisting penetration, and supporting other objects.

Hence, any inheritance principles which are sensitive to our common-sense conception of location – whether high- or low-dimensionalist – are likely to appeal in some way to the dynamics, and to look somewhat complex and indirect. The fact that the low-dimensionalist avoids positing, in addition to these complex and indirect inheritance

³⁷ See §2.6 for discussion of this kind of 'accommodation'.

principles for ordinary objects, the high-dimensionalist's natural and straightforward inheritance principle for local beables, does not seem to be a significant advantage.

2.5.3 Alternative local beable ghosts

Maudlin (2007: 3161-2; 2019: 123) observes that there are alternative ways of deriving local beables from high-dimensional ontology. For example, in addition to the mass-density proposal detailed above, there is the 'flash' proposal, which maps the wavefunction's jumps (given GRW's stochastic dynamics) to unstructured events or 'flashes' at corresponding spacetime points, constellations of which are supposed to underwrite material objects.³⁸ If both these mappings successfully enact material objects, then what could privilege one over the other?

One thing that certainly couldn't distinguish them is any experiment. There are differences between the proposals, of course: for example, the mass-density field permeates spacetime, whereas the flashes occupy a region of measure zero. But these differences could not (even in principle) be detected by experiment because the proposals make identical predictions about the macroscopic positions of all pointers: with overwhelmingly high probability, constellations of flashes are located exactly where clumps of high mass-density are.³⁹ So not only do these alternative local beables both enact material objects, but they enact objects which behave identically in all circumstances.

³⁸ See Bell 1987: ch.22).

³⁹ See Albert 2015: ch.4 for an illuminating discussion of experimental distinguishability.

There are (at least) two attitudes one might reasonably take in response to this observation: a hard (staunchly realist) line and a soft (more deflationist) line.

For hard-liners, the world's metaphysical structure determines, somehow, which of these constructions corresponds to the real material objects – indeed, the other construction may simply not exist at all. Such privilege may seem arbitrary from our perspective, but we needn't always have epistemic access to metaphysical structure. Indeed, it seems hubristic to suppose that where we lack the means to decide between alternative metaphysics, the world itself fails to decide.

For soft-liners, since both constructions are equally credentialed to enact the macroworld, there is no fact of the matter about which 'really' corresponds to material objects. This is just one more kind of indeterminacy in our ordinary talk: just as talk of tables doesn't precisely specify any spacetime regions or fusions of particles, so it doesn't decide between 'mass-tables' and 'flash-tables'.

Hardness and softness each have their distinctive disadvantages: the former embraces arbitrariness, whilst the latter seems to unexpectedly double (or perhaps, overdetermine) our world, by positing parallel, non-interactive material realms.

This makes for a difficult choice – but the crucial point for our purposes is that it is not a distinctively *high*-dimensionalist choice. Perhaps the 'larger' the gap between fundamental ontology and material objects, the more 'leeway' there is for alternative connections, and

the harder it becomes to decide between them. But the problem arises for many fundamental theories, whether high- or low-dimensionalist: there are often alternative connections available which seem equally good.

To see this, consider a low-dimensional universe consisting fundamentally of some particles together with a gravitational field. How might we identify the material objects in such a world? The particle-construction grounds the objects in the particles' trajectories; the field-construction grounds them in the contours of the gravitational field. These constructions are empirically equivalent: they necessarily agree on the macroscopic positions of all pointers. Again, a difficult choice looms, and there is a hard line and a soft line available: one might insist that only one of these connections holds, or one might concede that the material world is equally enacted by both.

Perhaps this just shows that there is something objectionably redundant about such a theory, compared to a particle-free or field-free alternative. But we are at the mercy of the physics here; if the dynamics requires both particles and fields, then scrapping either may be unavailable or at least unattractive. Besides, there are more familiar metaphysical questions about the connections between low-dimensional fundamentals and ordinary objects. For example, suppose that fundamentally there are particles which endure – that is, which are wholly present whenever they exist. Then we can ask about the persistence of ordinary objects (assuming that they exist) – do they endure too, or do they have temporal parts? These options correspond to alternative, empirically equivalent

constructions from the fundamentals – mirroring the high-dimensionalist’s choice between alternative local beables.

In short, metaphysics has been hard long before high-dimensionalism came along. Low-dimensionalist connections also involve tricky metaphysics. Sometimes connections can be empirically adjudicated, but often extra-empirical virtues must be considered. And in some cases, a somewhat deflationist attitude may indeed be warranted: there may not be any uniquely right way of identifying the material world.

Drawing the line between good and bad metaphysical questions is difficult. The present point is just that it seems utterly *ad hoc* to draw it between the issues raised by high- and low-dimensionalist theorizing. If you think the question of whether tables are really mass-tables or flash-tables is good, then you ought to think the question of whether they are really field-tables or particle-tables is equally good. Whatever arbitrariness is involved in answering such questions, it afflicts high- and low-dimensionalism alike.

2.6. Conspiracy Theories

To take stock: I have argued firstly that any asymmetry in scrutability between high- and low-dimensionalist grounding connections is conceptual – the connections themselves are equally substantive (§2.4), and secondly that there is no asymmetry in arbitrariness – high- and low-dimensionalism face the same kinds of choices between alternative connections

(§2.5). Hence, neither of these criteria provide good reason to prefer low-dimensionalist connections.

It is illuminating to see the concerns I have been considering as stemming from a common source in a particular conception of metaphysics, as articulated by Allori (2013-a: 63): ‘Once the scientist sets up the theory, the metaphysical picture it provides has already been defined, and there is very limited freedom of reinterpreting the formalism.’ On this conception, the theory’s metaphysics has been ‘fixed a priori’ by physicists (Allori 2013-a: 63), and the manifest image ought to be ‘implicit’ within it, falling out more or less straightforwardly ‘as a purely *analytical* consequence’ (Maudlin 2007: 3161). Hence, the metaphysician’s task is reduced to old-fashioned conceptual analysis: they are simply to interpret what the physicists mean.

If grounding connections really did follow analytically or *a priori* from a physical theory’s fundamental description of the world, then we should indeed expect them to be scrutable and non-arbitrary. However, this conception of metaphysics is importantly misleading. Whilst there is no doubt that anyone interested in the metaphysical structure of the world ought to pay close attention to our best physics, the business of extracting metaphysics from it is not mere conceptual analysis but itself a matter of substantive theorizing. As I have argued, even low-dimensionalist grounding connections are theoretical posits, akin to the other components of fundamental theories.

Some historical context provides a useful corrective here. High-dimensionalist proposals belong to a long history of fundamental theorizing which suggests that what it is to theorize in this way is to posit ‘conspiracy’: some underlying ontology arranged just as it needs to be to give rise to the world that we experience.⁴⁰ Fundamental theories are thus ‘conspiracy theories’, obliged to somehow accommodate the intersubjectively agreed upon platitudes that describe our experience of the world, including the outcomes of our experiments.⁴¹

‘Accommodate’ is deliberately vague: they must either explain the truth of these platitudes, or at least their appearance-as-if-true.⁴² The distinction between these kinds of accommodation may be somewhat fuzzy and metaphysically shallow; it turns in part on meta-semantic issues about the content of ordinary platitudes, and perhaps also on epistemological issues about what constitutes our evidence. Take the statistical mechanical ‘reduction’ of heat flow to molecular motion. Does this entail that heat really does flow, or that heat flow is an illusion explained by molecular motion? Or consider the platitude that tables are solid. Does the particulate model explain this solidity, or explain it away as illusory? Since it depends, in part, on the pre-theoretical content of these platitudes, it is unclear – but what matters is that, either way, they have been adequately accommodated.

⁴⁰ On this point, see North 2013: 200.

⁴¹ Nonetheless, there is a clear difference between ‘conspiracy theories’, in this importantly virtuous sense of the term, and mere *skeptical scenarios*: only the former aspire to offer genuinely systematic and confirmable explanations of the phenomena. This important distinction is sometimes ignored by critics of high-dimensionalism (for example, Monton 2006: 784).

⁴² Ney (2013: 173) makes a similar distinction between ‘eliminative’ and ‘retentive’ reductions.

High-dimensionalism should be regarded as the extension of this familiar mode of theorizing – the conspiratorial accommodation of ordinary platitudes – to a particularly primordial component of our everyday experience: the notion that our world has three spatial dimensions and one temporal dimension. According to high-dimensionalism, this platitude is true if it refers to the motion and interaction of material objects, but false if it refers instead to the topology of the fundamental arena (although its appearance-as-if-true is explained, via the naïve conflation of the fundamental arena with the material space). Plausibly, the ordinary notion of dimensionality doesn't clearly distinguish these alternatives. The surprising thing from a pre-theoretical standpoint, if high-dimensionalism is right, is just that they fail to coincide.

Fundamental theories have always required surprising, non-trivial metaphysics. The Ancient Greeks posited that material objects are combinations of Earth, Water, Air and Fire in certain ratios and arrangements. Subsequent theories grounding material objects in particles and/or fields have all made metaphysical claims which are radical not only from a top-down perspective – given our naïve conception of material objects – but also from a bottom-up perspective – given the natures of the fundamentals being posited. Pre-theoretically, the idea that tiny whizzing hard bits yield the stable material world seems just as shocking as the idea that combinations of the Greeks' four basic elements do.

Fundamental theories are obliged to accommodate the tables and chairs of our experience. But surely it is no constraint on such theories that we are able to simply *find* these objects readymade, clear and comprehensible, in their fundamental ontology. The business of

extracting them from a true fundamental theory is likely to be far messier and more abstract and more involved than that.

Imagine a hardcore Cartesian rationalist objecting to Schrödinger's equation that an analysis of the concept *WAVEFUNCTION* revealed this law to be inscrutable and arbitrary. Such complaints seem implausibly anthropocentric: we shouldn't expect dynamical structure to privilege our idiosyncratic and limited conception of the world. Should we expect *metaphysical* structure to privilege this conception?⁴³

Just as the Copernican revolution diminished our status within the dynamical order, we should be open to the quantum revolution diminishing our status within the metaphysical order. We should be open – at least if the phenomena lead us that way – to incorporating into our fundamental theories radically unfamiliar and unintuitive metaphysics, and to discovering that the world we appear to inhabit is far removed from the fundamental arena that underlies it.

⁴³ Schaffer (2017: 14) makes a similar comparison.

3. Two Approaches to Metaphysical Explanation⁴⁴

3.1. Introduction

We come to metaphysics with a wide-ranging and rich conception of the world. Metaphysicians seek to impose some order on this conception (perhaps after some revision). They ask: how can we make sense of the world thus conceived? How do its various aspects ‘hang together’? This can be seen as demanding a distinctively metaphysical kind of *explanation*: metaphysicians want to know how various aspects of reality can be metaphysically explained in terms of others.

This demand presupposes that the world has a certain structure: it is not just a disconnected mess of brute details. Two observations seem to justify this assumption. Firstly, reality seems to be certain ways in virtue of being other ways; we capture this general idea by saying that some aspects of reality *ground* others.⁴⁵ Secondly, certain aspects of reality seem superficial, derivative, or mundane, whereas others seem deep, basic, or ultimate; we capture this general idea by saying that some aspects of reality are more *fundamental* than

⁴⁴ Thanks to Karen Bennett, David Builes, Alex Skiles, the Spring 2020 Rutgers Metaphysics Group, two anonymous referees, and especially to Verónica Gómez, Jonathan Schaffer, and Ted Sider for many discussions (and insights into the minds of generators and reducers) which have greatly benefitted this paper.

⁴⁵ *N.B.* I use ‘ground’ throughout as a generic term for the explanatory connection between more and less fundamental, as opposed to any specific conception of this connection. It can be taken to signify either metaphysical explanation itself (cf. ‘unionists’: e.g. Fine 2001, Dasgupta 2017), or a relation backing explanation (cf. ‘separatists’: e.g. Audi 2012, Schaffer 2012).

others. And these two features appear intimately connected: more fundamental aspects seem to ground less fundamental aspects.

In this way, the world appears to support metaphysical explanation: we can explain its less fundamental aspects in terms of the more fundamental aspects which ground them. And we might hope, or even expect, that the world supports an especially satisfactory form of metaphysical explanation, by providing a foundation of absolutely fundamental aspects which both ground all others and are themselves ungrounded. These, at a highly abstract level, are the roles of grounding and fundamentality in explanatory metaphysics (in the broadest sense of these terms).⁴⁶

As I see it, the contemporary metaphysics literature exhibits a crucial but neglected divide between two overarching conceptions of what plays these roles, and correspondingly two substantively different approaches to metaphysical explanation. On the ‘generation’ approach, the fundamental generates, or gives rise to, less basic features of reality. This generation is analogous in many ways to causation; in something like the way causation connects earlier to later, generation connects more basic to less basic. On the ‘reduction’ approach, by contrast, explanatory metaphysics seeks to reinterpret our conception of reality in perspicuous, or less metaphysically distorted, terms. There is merely a

⁴⁶ Although I will speak as if all metaphysical explanation revolves around fundamentality and ground, I recognize that there may be other kinds of metaphysical explanation (especially revolving around *essence* — see Fine 2015, Glazier 2017, Correia & Skiles 2017). But it bears emphasis that I am employing *broad* notions of fundamentality and ground: some kinds of metaphysical explanation may revolve around these notions whilst being disconnected from narrower versions of them. (A similar remark applies to those who would disconnect metaphysical explanation from these notions altogether e.g. Baron & Norton 2019.)

representational (rather than worldly) difference between fundamental and non-fundamental: non-perspicuous ways of representing reality reduce to, or collapse into, others.

The divide between these approaches is rarely explicitly articulated and discussed.⁴⁷ But it casts a shadow over ‘first-order’ debates: many disputes within both metaphysics and neighboring disciplines look quite different depending on which approach is taken. In this paper, I aim to articulate and clarify the core differences between generation and reduction (§3.2), to demonstrate the distinction’s wide-ranging significance (§3.3), to suggest how we might adjudicate between the two approaches (§3.4), and to argue that explanatory metaphysics needs both (§3.5).

3.2. The Two Approaches

Let me introduce the distinction between generation and reduction by means of some familiar examples of metaphysical explanation.

First, composition. According to a natural view, a composite object’s existence and nature are metaphysically explained by the existences and natures of its parts. For example, a table exists and has a certain mass at least partially because the atoms which make it up exist and have certain masses. On the generation approach to this explanation, the existences and masses of certain atoms generate – or ‘metaphysically cause’ – the table’s existence

⁴⁷ Williams (2012) and de Rosset (2017) are notable exceptions.

and mass. The table-portion of reality and the atomic portion are distinct but intimately connected: one gives rise to the other.⁴⁸ On the reduction approach to this explanation, there is only one portion of reality, and two ways of representing it. Our talk of the table's existence and mass is true, but the state of affairs it represents would in some sense be better represented in terms of the existence and masses of certain atoms. Only the atoms are 'really there', but they support our second-rate talk of tables (unlike talk of unicorns, say).⁴⁹

Second, mentality. On a standard understanding of physicalism, the mental is metaphysically explained by the physical. For example, the legend goes that we feel pain in virtue of our C-fibers firing. On the generation approach to this explanation, my C-fibers firing (a physical event) generates my pain (a mental event): the events are distinct but intimately connected. (This is a natural interpretation of 'non-reductive physicalism'.)⁵⁰ On the reduction approach, the pain *just is* the C-fibers firing: our talk of pain is a distinctive way of latching onto what is 'really' just some neural activity. (This is a natural interpretation of 'reductive physicalism'.)⁵¹

Third, determinables. An object's determinable properties are naturally explained in terms of its more specific determinate properties: for example, the apple is red because it is scarlet. But again, there are two quite different ways of understanding such explanations.

⁴⁸ E.g. Russell 2003:92, Fine 2010: §IX.

⁴⁹ E.g. Cameron 2010, Sider 2013c.

⁵⁰ E.g. Shoemaker 2007, Schaffer 2021.

⁵¹ E.g. Loar 1990, Papineau 2002.

On the generation approach, determinable properties are just as real as their determinates, with the facts involving the former generated or determined by the facts involving the latter.⁵² Alternatively, on the reduction approach, objects only really instantiate maximally specific properties, with determinables being mere ‘shadows’ of linguistic abstraction. On this view, there is a single state of affairs with respect to the apple, and more or less abstract ways of describing it.⁵³

These examples suggest that the distinction I am concerned with is very general: many proposed metaphysical explanations can naturally be conceived in two different ways. The rest of this section aims to present these two approaches as clearly as possible. Since my presentation aims to be neutral and abstract, particular examples, general principles, and the choice of primitives are all contentious. Nonetheless, we can ‘triangulate’ onto the distinction by considering a range of reasonably natural cases and connections to surrounding notions. What is important is that a clear conception is conveyed, which readers may then adapt to their own preferred frameworks. I hope that the reader’s grasp of the distinction will strengthen throughout this paper, as further cases and connections arise.

3.2.1 Facts vs. truths

The two approaches share the idea that certain ‘aspects’ of reality are more or less fundamental and stand in relations of ground to one another. But adherents of the

⁵² E.g. Wilson 2009, Rosen 2010: §11.

⁵³ E.g. Armstrong 1997:53, Heil 2003:207.

generation approach (‘generators’) take these aspects to be worldly, akin to objects or events, whereas adherents of the reduction approach (‘reducers’) take them to be representational, akin to names or sentences. I will speak in terms of facts on the generation side and truths on the reduction side.⁵⁴

Truths, as I use the term, are just true truth-bearers: perhaps sentence-tokens or sentence-meanings/propositions. (Since ‘proposition’ is sometimes used in a worldly sense, I will reserve it as a term which is neutral between facts and truths). ‘Fact’ is sometimes used as synonymous with my term ‘truth’. However, my use of ‘fact’ follows a tradition which recognizes a more inflationary (or ‘thick’) notion, on the side of reality rather than representation (Russell 1918, Wittgenstein 1922, Armstrong 1997).⁵⁵ On this conception, facts are entities which constitute the world — often identified with obtaining states of affairs — rather than true ways of representing the world.

Truths stand to facts as sense stands to reference: they are guises by which facts are represented.⁵⁶ Just as different names can refer to the same object, and different predicates can denote the same property, so different truths can represent the same fact (Correia 2010:

⁵⁴ Although I set this aside, both generator and reducer may also target ‘sub-propositional’ aspects (e.g. Schaffer 2009 on entity-grounding, Fine 2015 on generic grounding, and Dorr 2016 on identifications).

⁵⁵ Cf. Correia (2010:256) on “worldly facts” vs. “conceptual facts” and Jago (2011:42) on facts vs. propositions.

⁵⁶ Following Fine (2013), we might distinguish two senses of ‘representing’: a truth may *describe* a fact, or it may *express* a fact, where a disjunction represents the same fact as its true disjunct in the first sense but not the second. I largely ignore this issue here, but it becomes relevant in §3.5.2 (see n.113).

§2, Audi 2012: §IV.1). For example, the following pairs of sentences plausibly correspond to different truths but the same facts:

‘Hesperus is rising’, ‘Phosphorus is rising’;

‘Snow is white’, ‘Snow is white and snow is white’;

‘London is north of Paris’, ‘Paris is south of London’.

In each case, the difference between the sentences seems merely representational, rather than worldly: they are anchored to the same portion of reality, but tied by different ropes.

Of course, there are many ways of carving up both facts and truths, and I cannot explore the issues in any depth here. The crucial point is just that the criteria for individuating them differ, with truths individuated more finely: if p and q correspond to the same truth, they represent the same fact, but perhaps not conversely. In particular, certain paradigmatically ‘opaque’ contexts — which, intuitively, are sensitive to merely representational features — may distinguish truths, but not facts. For example, it might be held that ‘Hesperus is Hesperus’ is analytic, a priori, and believed by the Babylonians, whereas ‘Hesperus is Phosphorus’ is synthetic, a posteriori, and disbelieved by the Babylonians. This indicates that these sentences correspond to distinct truths, but not that they represent distinct facts.

For the sake of regimentation, let me introduce some official primitives.⁵⁷ I assume that for every true sentence *p*, there is *the truth that p*, which I denote by ‘<*p*>’, and which *represents* some fact.⁵⁸ I will also use the expression ‘*the fact that p*’, which I abbreviate by ‘[*p*]’. On one way of using this expression, ‘[*p*]’ denotes whichever fact <*p*> represents. However, following somewhat standard terminology (e.g. Rosen 2010:115), I adopt a more restricted use, where ‘[*p*]’ only successfully refers if reality provides a matching fact. I explicate this usage below, connecting it to perspicuity.

I am speaking as if ground is a relation, with the two approaches conceiving its relation differently. This makes for a vivid picture, but it should be remembered that ground needn’t be conceived this way (on either approach). An alternative conception expresses grounding claims using a sentential operator, with no commitment to any entities corresponding to the expressions that flank it (Correia & Schnieder 2012:10; Fine 2012:47; Dasgupta 2017: §2). Accordingly, an alternative way of presenting the two approaches eliminates my talk of facts and truths in favor of sentential operators (such as *...is factually equivalent to...*). The difference between my ‘reifying’ presentation and this non-reifying alternative will not matter for the discussion which follows.

⁵⁷ In taking these notions as primitive, I simply mean that I will not be defining them; I do not mean that they are primitive in any metaphysical sense, or that they cannot be defined.

⁵⁸ This is idealized in several ways. For example, since truths correspond to sentences in contexts, ‘<...>’ should be understood as a context-dependent expression, and since ordinary sentences are almost always imprecise, their corresponding truths do not determinately represent any particular fact. Also, there may be ‘non-factual’ truths which do not represent any fact e.g., <torture is wrong>, on an expressivist view. (Conversely, truth may be stronger than required for the reducer’s representational aspects; some weaker positive status such as ‘aptness’ may suffice.)

3.2.2 Generation vs. reduction

Generators take ground to be a matter of worldly aspects *generating* (giving rise to, or building) others. Introduction by paradigm cases is fraught in this context, since it is contested not only whether a given case involves ground at all, but also, assuming that it does, which notion of ground (generation or reduction) it is. Nonetheless, the grounding of sets by their members, and of wholes by their parts, provide intuitive examples. [Socrates exists] is naturally taken to generate [Socrates exists], and the fact that the hydrogen atoms are bonded is naturally taken to generate the fact that the H₂ molecule exists. In each case, the facts in question seem clearly distinct: in the first, one involves a person and the other a set; in the second, one involves two atoms and the other one molecule. Nonetheless, the facts seem intimately connected, with one both metaphysically necessitating and wholly relevant to the other. Thus, these cases exhibit the characteristic features of generation: they involve a tight explanatory connection between distinct facts.

The model here is causation — as Jonathan Schaffer (2012:122) puts it: ‘Grounding is something like metaphysical causation. Roughly speaking, just as causation links the world across time, grounding links the world across levels.’ Developing this analogy, Schaffer (2016) observes that grounding and causation are similarly connected to laws, necessity, counterfactuals, and explanation.⁵⁹ In a similar vein, Karen Bennett (2011a, 2017) characterizes a family of ‘building’ relations, operating variously on objects, properties, and states of affairs, with composition and constitution as paradigm examples. Building relations are unified by their licensing ‘generative’ talk: ‘builders *generate* the built’

⁵⁹ See also A. Wilson (2018).

(2017:58).⁶⁰ And Gideon Rosen glosses grounding as ‘metaphysical dependence’ (2010), ‘a much more intimate form of dependence’ than causation (2011:123), in which facts ‘*give rise to or generate*’ others (2011:130).⁶¹

Generators standardly understand fundamentality in terms of generation: generation orders the world into a hierarchy — ‘the great chain of being’ (Schaffer 2009:376) — and fundamentality is a matter of position within this hierarchy.⁶² If one aspect of reality generates another, it is thereby more fundamental than it,⁶³ and the absolutely fundamental aspects of reality are those that are ungenerated (Schaffer 2009:351; Rosen 2011:124; Bennett 2017: ch.5). I will use ‘basicness’ to denote this conception of fundamentality; thus, I will speak of more basic facts generating less basic facts.

By contrast, reducers take ground to be a matter of truths reducing to (consisting in, or collapsing into) others. With the same caveat about paradigm cases as above, intuitive

⁶⁰ Bennett reserves ‘grounding’ for a particular building relation, but notes a more generic sense which may be synonymous with ‘building’ (2011a: n.10; 2017:12).

⁶¹ Rosen’s credentials as a generator might seem undermined by his ‘grounding-reduction link’ (2010: §10), according to which, if <p> ‘reduces’ to <q> — in that <p>’s being the case ‘consists in’ <q>’s being the case — then [p] is grounded in [q]. However, since Rosen posits a worldly distinction between the facts in question, reduction in his sense is not reduction in my sense: it is not a relation among sentences ‘but rather among the facts ... those sentences purport to describe’ (2011:122).

⁶² One might view Kit Fine as a generator who does not understand fundamentality in terms of generation: in addition to his primitive notion of ground, he posits a primitive, absolute notion of ‘fundamental reality’, denying that it has ‘a relational underpinning’ (2001:25). However, as discussed below, I see Fine’s notion of fundamentality as paradigmatic of the reducer’s notion of perspicuity, and hence prefer to interpret his rich framework as combining the two approaches. Indeed, he (2001:15) explicitly distinguishes statements of ground, which sometimes connect propositions which ‘hold in reality’, from ‘stricter’ statements of reduction, which imply ‘the unreality of what is reduced’.

⁶³ Analyzing relative fundamentality in terms of generation is non-trivial: see Bennett 2017: §6.5.

examples are the reduction of truths about water to truths about H₂O and of truths about heat to truths about molecular motion. It is natural to hold that the truths in question represent the same facts, where this equivalence is directed: the water in the glass ‘is really just’ a collection of H₂O molecules, and the air getting hotter ‘is really just’ air molecules moving faster. This directedness seems to support explanation: truths about water obtain because truths about H₂O obtain, and likewise truths about heat obtain because truths about molecular motion obtain.

Reducing truths seem to ‘match’ facts which the reduced truths seem to ‘distort’ (and it is this that provides explanatory asymmetry). When discussing the examples of metaphysical explanation above, I spoke of what may be said truly in a ‘second-rate’ language, or what is merely a ‘shadow’ of language, being explained in terms of how things ‘really’ are, or what holds ‘in reality itself’. My preferred regimentation of this way of speaking uses a predicate of truths: ‘<p> is *perspicuous*’ means that <p> is not only true but reflects reality’s structure in a deeper sense.⁶⁴ I propose to understand reduction in terms of perspicuity as follows. ‘<p> reduces <q>’ means that:

- i) <p> and <q> represent the same fact, and
- ii) <p> is perspicuous and <q> is not.⁶⁵

⁶⁴ Cf. (O’Leary-)Hawthorne & Cortens 1995: §3, Fine 2001:3, Turner 2010:8–9, Sider 2013c:252.

⁶⁵ In the more general many-one case, ‘<p_{1n}> reduces <q>’ means that:

- i) <q> represents the same fact(s) as the conjunctive truth <p₁ & ... & p_n>,
- ii) <p₁>, ..., <p_n> are each perspicuous and <q> is not.

For example, ‘<the mean molecular energy of the air in the room is x> reduces <the room temperature is y>’ means that: i) these truths represent the same fact, and ii) <the mean molecular energy of the air in the room is x> is perspicuous and <the room temperature is y> is not.

We can get a grip on the key notion of perspicuity by considering the familiar example of ‘the average family has 2.2 children’. Granting the statistic’s accuracy, this sentence is true. But, in another sense, it clearly fails to ‘match’ the fact it represents: there is no family bearing the having relation to ‘each’ of 2.2 children. A sentence which appears to better approximate the corresponding fact is: ‘the total number of children divided by the total number of families is 2.2’. (Of course, this sentence may itself be non-perspicuous, e.g. if numbers don’t exist.)⁶⁶

As I see it, the notion of perspicuity appears in many guises throughout analytic metaphysics, and there are a number of routes to understanding it.⁶⁷ Following Kit Fine (2001: 25), one might view perspicuity as a matter of a truth’s matching the structure of

⁶⁶ This example may be superficial: the structure of the corresponding sentence-meaning may not match that of the sentence (Kennedy & Stanley 2009). Nonetheless, the mismatch between the sentence itself and the fact represented vividly illustrates non-perspicuity, and it will be helpful in what follows to proceed under the simplifying assumption that sentences’ structures by and large resemble that of their corresponding truths.

⁶⁷ Historically, it is central to the logical atomists’ notion of analysis and Quine’s notion of paraphrase. More recently, see van Inwagen (1990) on ‘real’ vs. ‘virtual’ objects, Field (1994) on ‘factually defective discourse’, Dorr (2007) on ‘fundamental’ vs. ‘superficial’ uses of sentences, Cameron (2010) on ‘Ontologese’, Williams (2012) on ‘requirements on reality’, Russell (2015) on ‘objective matters of fact’, deRosset (2017) on conciliatory irrealist views, and Bacon (2020) on fundamentality. I cannot undertake any detailed exegesis here, but — despite importance differences between their frameworks — I understand all of these authors as invoking versions of perspicuity.

the corresponding fact.⁶⁸ On this conception, $\langle p \rangle$ is perspicuous just in case for each representational constituent of $\langle p \rangle$, $[p]$ has a corresponding worldly constituent, and for each structuring relation between the constituents of $\langle p \rangle$, the constituents of $[p]$ are correspondingly related.

For example, a simple atomic truth of the form ‘a is P’ is perspicuous if and only if there is some object a and property P, corresponding to the name ‘a’ and predicate ‘P’ respectively, such that a instantiates P. Suppose there are facts involving determinate properties of the form *has x kg mass* but not involving the determinable property *has mass*. Then $\langle \text{Joe has mass} \rangle$ would not be perspicuous since there is no correspondingly structured fact $[\text{Joe has mass}]$, only a fact of the form $[\text{Joe has x kg mass}]$. Obviously this requires a structured conception of facts as built up from worldly constituents, such as objects and properties, together with a suitably ‘inflationary’ understanding of these constituents and their relations (Rosen 2010:114; Audi 2012:686).⁶⁹

Alternatively (though not incompatibly), one might understand perspicuity in terms of a sub-propositional notion. Following Ted Sider (2011: 116), one might require that the

⁶⁸ As Fine (2001:3) puts it, truths which fail to ‘hold in reality’ do ‘not *perspicuously* represent the facts—there will be some divergence between how the facts are “in themselves” and how they are represented as being’. Cf. Bennett (2001:147) on ‘truths that directly reflect the metaphysical situation’.

⁶⁹ Resistance to the idea that facts have sentence-like structures has surfaced recently, especially on the basis of the Russell-Myhill paradox (e.g. Dorr 2016, Goodman 2017). Whilst I cannot adequately address this issue here, we should distinguish the general idea that facts are structured from the view that truths with different structures represent different facts. For example, one might hold that $\langle \text{It is not not the case that snow is white} \rangle$ and $\langle \text{Snow is white} \rangle$ are factually equivalent (thus rejecting the latter view) whilst maintaining that the fact represented is structured. See Bacon 2020: §4 for a view along these lines.

constituents of a perspicuous truth are each ‘structural’. This notion extends Lewisian naturalness ‘beyond the predicate’, to names, sentential operators, quantifiers, *etc.* Structural notions ‘carve nature at its joints’, in ways that matter for objective similarity, explanation, and confirmation, and which confer a kind of epistemic value that is independent of truth. For example, ‘is grue’ and ‘is taller than’ do not denote perfectly natural properties/relations; hence <this emerald is grue> and <Trump is taller than Obama> are not perspicuous truths. On the other hand, perhaps ‘is negatively charged’ does denote a perfectly natural property; hence <Sparky is negatively charged> may be perspicuous (insofar as the name ‘Sparky’ is itself structural!)⁷⁰

I have been employing a binary notion of perspicuity. However, reduction is naturally viewed as a gradual process of analyzing or unpacking a non-perspicuous truth, with the underlying fact better approximated as various layers of distortion are removed. One might therefore recognize a relative notion of perspicuity, with reduction a matter of transforming some non-perspicuous truth into a more perspicuous (but perhaps not perfectly perspicuous) truth. For example, one might view <the mean molecular energy of the air in the room is x > as more perspicuous than <the room temperature is y >, but still in need of further reduction to yet more perspicuous truths about the masses and velocities of the particular air molecules. Only once we reach perfectly perspicuous truths — those which

⁷⁰ Along similar lines, one might require that the constituents of a perspicuous truth be ‘metaphysically primitive’, in the sense that they have no real definition. We can use Dorr’s (2016) framework of ‘identification’ to understand the relevant notion. On Dorr’s preferred view (2016: §8), identification obeys a non-circularity constraint: nothing can be a ‘non-logical complication’ of itself. For example, since grue is defined in terms of green (together with something non-logical), green cannot itself be defined in terms of grue. Thus, we can understand a real definition of x as an identification of x with some complex notion which involves a non-logical notion other than x itself (though this notion cannot be defined in Dorr’s preferred language).

perfectly reflect reality itself — is reduction complete.⁷¹ Although I take the absolute notion of perspicuity to be my official primitive here, there is an important project of characterizing the relative notion (which may or may not be prior to the absolute notion).⁷²

Just as generation can be seen as ‘metaphysical causation’, so reduction can be seen as ‘metaphysical semantics’: it provides something like the ‘metaphysical meaning’ of the reduced truth (Sider 2011: §7.4). Since it requires perspicuity, reduction is more constrained than ordinary semantics — in Sider’s (2011:112) words, it should ‘show how what we say fits into fundamental reality’. This is no constraint on ordinary meanings, which needn’t be ‘metaphysically privileged’ over the object-language itself. In another way, however, reduction is less constrained: it needn’t integrate with the cognitive science of language use and acquisition (at least as directly as ordinary semantics).⁷³

Both generators and reducers hold that ground is (at least typically) irreflexive: grounds and grounded are distinct. But where the generator sees a worldly distinction between facts,

⁷¹ One might even deny that there are any perfectly perspicuous truths. On one picture, there are no maximally perspicuous truths: any truth can be further unpacked. On another picture, there are maximally perspicuous truths, but they themselves involve some distortion: reality has an inherently ‘non-linguistic’ structure.

⁷² Relative perspicuity might be understood in terms of relative naturalness, or via Dorr’s (2016: §9) notion of ‘metaphysical priority’.

⁷³ The details here are subtle. Reduction can’t be completely divorced from semantics: for example, water reducing to H₂O is not independent of ‘water’ referring to H₂O. We might see the difference in the following way: ordinary semantics aims to describe truths, whereas metaphysical semantics aims to explain them. One plausible constraint is therefore that a sentence and its meta-language equivalent in a good semantic theory ought to have the same reduction.

the reducer merely sees a representational distinction between truths.⁷⁴ It is important to note that, although claims about reduction pertain directly to representation, they are also (albeit, indirectly) about the reality being represented: in particular, to say that $\langle p \rangle$ reduces to $\langle q \rangle$ is to say that the fact represented by $\langle p \rangle$ has a structure matching that of $\langle q \rangle$. I suspect that some feel an instinctive aversion to the reduction approach, on the basis that metaphysics is concerned with reality itself, independently of how we think or speak.⁷⁵ But I don't think that reducers should be seen as lapsed metaphysicians who conflate representation and reality: rather, they approach reality *via* representation. This approach is no more confused than the idea that in 'reverse-engineering' a landscape from an impressionist painting we ought to take into account the stylistic distortions of the artist. Indeed, from a reducer's perspective, it is generators who conflate representation and reality — it is as if, in comparing the painting with a photograph of the same landscape, they conclude that there must be two landscapes: the photographic one, and the impressionist one it generates!⁷⁶

⁷⁴ Jenkins (2011) challenges the irreflexivity of 'metaphysical dependence' on the basis of cases such as pains both depending on and being identical with brain states. However, such cases may be better interpreted as involving reduction.

⁷⁵ For example, Dorr (2016:44) expresses the conviction that 'any operators we might need to appeal to in stating questions that are central to the subject matter of metaphysics should be transparent'. Since perspicuity/reduction are opaque, this would make questions stated in terms of them peripheral at best.

⁷⁶ A distinction is sometimes drawn between 'representationalist'/'conceptual' grounding and 'worldly'/'metaphysical' grounding (e.g. Correia 2010:258-9; Correia & Schnieder 2012:21; Correia & Skiles 2017:656). Those introducing the distinction often disavow the representationalist notion as conflating 'mere shadows of language with real features of the world' (Kramer & Roski 2015:60).

As I see it, this distinction is internal to the generation approach — concerning the granularity of facts — rather than corresponding to the broader distinction of interest here. The key difference between the two distinctions is that no distinctive notion of fundamentality is associated with representationalist as opposed to worldly grounding, whereas reduction is tied to perspicuity as opposed to basicness.

3.3. Significance

I have been speaking as if metaphysicians divide into two camps: generators and reducers. But of course, things aren't so neat. Many metaphysicians do not clearly have either conception in mind (and sometimes seem to veer between them). Even those who consistently employ one approach rarely do so in self-conscious opposition to the other. Indeed, I suspect that most would regard the two approaches as nuanced variations on the same core idea.

This section argues that the distinction between generation and reduction deserves careful attention. It enables us to better understand what is at issue in certain metaphysical debates and the arguments given by each side (many of which implicitly trade on the underlying approach). I focus on some particular case studies, but I hope my discussion will indicate that the distinction's impact on metaphysics is pervasive.

3.3.1 Ontology

In paradigmatic ontological disputes, e.g. concerning ordinary objects, or abstracta, articulating the disagreement can be difficult. As Bennett (2009) describes, both sides 'minimize the difference', playing down their commitment to/rejection of the controversial entities.

Take mereology. 'Conciliatory nihilists' like Cameron (2010) and Sider (2013b) hold that the table before me exists, but only in a non-fundamental sense: it exists because there are

some simples (objects with no smaller parts) arranged table-wise. These simples being arranged table-wise necessitates the existence of the table, but the connection needn't be analytic or a priori: those who are competent with the relevant concepts may doubt it. This creates a puzzle: what makes this view 'nihilist'? After all, many 'conciliatory anti-nihilists' agree that the table's existence is non-fundamental, and explained by some simples being arranged table-wise, via a necessary a posteriori connection.

Distinguishing two approaches to metaphysical explanation allows us to understand the underlying dispute. Conciliatory nihilists say things like: 'reality itself is merely a swarming of simples; it lacks any compositional structure matching our ordinary conception.' They hold that truths about composite objects reduce to truths about simples. For their conciliatory anti-nihilist opponents, by contrast, facts involving composites are generated by facts involving simples. Our talk of composites is not merely true on this view, but reflects reality itself.

Many, I believe, have an inchoate grasp of disagreements like these, and are only skeptical insofar as they cannot clearly articulate them. It is a point in favor of the distinction that it allows us to do so. Moreover, the distinction enables us to better understand the arguments given on each side: the arguments given by 'believers' in the entities in question typically suit a generation approach, whereas those given by 'deniers' typically suit a reduction approach. This is unsurprising: the attractive conciliatory versions of these positions — according to which the entities in question exist non-fundamentally — are only available

once the corresponding approach is recognized. Recognizing both approaches allows us to properly evaluate the arguments.

On the believers' side, take Schaffer's (2009:360) insistence that the existence of controversial entities is (almost always) obvious: ordinary material objects, numbers, properties, and even fictional characters all exist in the one true non-deflationist sense, as 'full-blown "heavyweight"' entries on the roster of entities'. On Schaffer's view, the interesting question is how these entities are generated by more basic entities. He (2009:358) motivates this 'permissivism' by various Moorean arguments, such as:

My body has proper parts (e.g., my hands).

Therefore, there are things with proper parts.

This argument does not impress conciliatory nihilists, who deny the perspicuity but not the truth of its conclusion. What, then, underlies Schaffer's dismissive attitude? It is that his framework leaves no room for conciliatory nihilism; since he understands fundamentality as basicness, he interprets the claim that bodies exist 'non-fundamentally' as entailing that they are derivative entities which exist in the one true sense! Thus, the only nihilist opponent he recognizes is revisionary.⁷⁷ The real issues emerge once we have reduction in sight. For example, are we justified in upholding the perspicuity of Moorean claims? And

⁷⁷ Cf. his description of Sider's approach to the non-fundamental as 'radically eliminative' (2013:736).

are we justified in believing that these claims are true, as opposed to ‘apt’ in some weaker sense?⁷⁸

On the deniers’ side, take Sider’s (2013b: §1) argument that nihilism is more ideologically parsimonious since it ‘allows us to eliminate ‘part’ from the ideology of our *fundamental* theories... If one’s theory of fundamental matters included an ontology of composite objects, then that theory would presumably also need a predicate of parthood to connect those composites to their parts...’. This won’t impress conciliatory anti-nihilists: their ‘fundamental theory’ only talks about simples, not the derivative fusions they generate! Sider is evaluating parsimony by the notions which figure in a perspicuous theory, whereas the anti-nihilist I am imagining evaluates parsimony by the constituents involved in the basic facts. Thus, applying the distinction reveals the key issue (to which I return shortly): which sense of fundamentality — basicness or perspicuity — is relevant to parsimony?⁷⁹

3.3.2 Meta-ontology

For ‘highbrow’ ontologists, ontology concerns what exists fundamentally:

⁷⁸ See Sider 2013c: §4 on the metaphysical shallowness of the meta-semantic issue of truth. This becomes important if the English quantifier and the perspicuous quantifier cannot differ in meaning; see Dorr 2014.

⁷⁹ It is tempting to argue that the anti-nihilist’s complete fundamental theory must include not only the basic facts but the principles of generation connecting them to derivative facts, and parthood (or the like) is needed to state these connections (Sider 2011: §8.2.1). Sider’s view avoids the analogous issue since he takes truths about reduction to be reducible. But the anti-nihilist might defend the parallel view that facts involving generation are themselves generated by basic facts not involving parthood (Bennett 2011b, deRosset 2013; cf. Bennett 2017:226).

What we debate in the ontology room is ... what there is *in the most fundamental sense* (Dorr 2005:24).

... 'a really (or fundamentally) exists' is true iff a is an element of our ontology (Cameron 2008:7).

The ontological question is... "Are there Fs?" where 'there are' is understood as having a fundamental sense (Sider 2011:171).

On this view, when nihilists and anti-nihilists debate the existence of composite objects, nominalists and platonists debate the existence of abstracta, and so on, their disagreement concerns what exists *fundamentally*.

Opposed to this elitism are 'lowbrow' ontologists, who separate ontology from fundamentality:

Of course not every entity is a fundamental entity... (Schaffer 2008:18)

Fundamentality does not lead to greater reality. (Hofweber 2016:329)

...flatworldism is the result of taking to its extreme the thought that nonfundamentalia are less than fully real. I myself have never seen the appeal of this thought... (Bennett 2017:216).

On this view, ontology (even ‘serious’ ontology) concerns plain old existence — and non-fundamental things exist too.⁸⁰

These two views are *meta*-ontological: they characterize ontological disagreements rather than addressing them. But of course meta-ontology guides ontology; insofar as one is drawn to the moderate claim that some controversial entity exists non-fundamentally, one will be a denier/believer if one is a highbrow/lowbrow ontologist.

The disagreement between highbrow and lowbrow ontologists can be perplexing. On one hand, it is liable to seem verbal, with the disputants merely employing the technical term ‘ontology’ differently. Both sides distinguish three categories: fundamental, non-fundamental, and non-existent (Kriegel 2015). The disagreement is then whether ontology concerns what goes in the first category or what goes in the last. But both are good questions; it doesn’t matter which we call ‘ontology’! Indeed, even the disputants themselves sometimes suggest this reconciliation.⁸¹

On the other hand, the disputants’ understanding of ontological disputes as substantive belies the idea that the background meta-ontological dispute is merely verbal. For example, Schaffer (2009) and Bennett (2017) are believers in ordinary objects, whilst Cameron (2010) and Sider (2013b) are deniers. There is an obvious reconciliation of these

⁸⁰ Along similar lines, von Solodkoff & Woodward (2013:568) distinguish ‘inflationism’ about the non-fundamental (~lowbrow ontology) from ‘deflationism’ (~highbrow ontology).

⁸¹ See Schaffer’s (2008:18) response to Cameron, and Sider’s (2013a:761) response to Schaffer.

‘ontological’ positions: Cameron and Sider address the highbrow question of whether ordinary objects exist fundamentally, whereas Schaffer and Bennett address the lowbrow question of whether they exist *simpliciter*.

Distinguishing generation and reduction clears this up. Highbrow ontologists understand fundamentality in terms of perspicuity —their view is that ontology concerns what exists perspicuously (that is, which objects the perspicuous truths are about).⁸² Lowbrow ontologists understand fundamentality in terms of generation — in shunning elitism, they extend ontology to derivative entities (which are no less real).⁸³ It is because they have different background conceptions of fundamentality that existing non-fundamentally is ‘really’ existing for lowbrow but not highbrow ontologists. Since they mean different things by ‘fundamental’ — not ‘ontology’ —the meta-ontological disagreement is verbal after all. This reconciliation accommodates substantive ontological disputes: deniers hold that non-fundamentalia only exist in a non-perspicuous sense, whereas believers hold that they are derivative entities (cf. Sider 2011:170; 2013c:253).⁸⁴

Distinguishing the two approaches also exposes an underlying disagreement on the closely connected question of ontological parsimony. Cameron, Sider, Schaffer and Bennett

⁸² Not: which truths of the form ‘there are Xs’ are perspicuous. Table-believers might recognize perspicuous truths about tables whilst denying that ‘there are tables’ is perspicuous, because either the predicate ‘table’ or quantification itself may be non-structural.

⁸³ Although, see Rosen (2011).

⁸⁴ Highbrow ontology is an instance of the broader idea that metaphysics concerns what is fundamental, as criticized by Barnes (2014) and Bennett (2017: §8.3). Whichever conception of fundamentality we adopt, Barnes and Bennett are surely correct about the broader point: important metaphysical questions needn’t concern either basic facts or perspicuous truths.

appear to share the view that only commitment to fundamental entities counts against the ontological parsimony of a theory (leading Cameron (2010: n.2) to speculate that their dispute is ‘merely terminological’). But in fact, their views are importantly different: what counts for Cameron (2010:250) and Sider (2011:198) is what exists in a perspicuous sense, whereas Bennett (2017: §8.2.2) and Schaffer (2015) only count the basic entities. This disagreement is a significant driver of ontological disputes: deniers typically reject the believers’ claim that commitment to derivative entities is non-costly. (More generally, we can consider what counts against ‘ideological parsimony’: the constituents of perspicuous truths or of basic facts?)

Two initial points concerning the underlying disagreement are salient. First, insofar as parsimony is tied to explanatory efficiency — the principle being: ‘minimize the unexplained’ (Schaffer 2009:361; Bennett 2017:221) — then derivative structure seems non-costly. Why should it matter that some truth is irreducible, if the corresponding fact can nonetheless be explained via generation? Second, though, an important motivation for denying that non-fundamental structure is costly is the idea that it is ‘nothing over and above’ the fundamental (Lewis 1991:81), or ‘no addition to being’ (Armstrong 1997:12). For generators, since derivative entities/facts are distinct from the basic entities/facts generating them, these slogans are awkwardly misleading.⁸⁵ Reducers, by contrast, can make straightforward sense of the idea. If the table’s existing reduces to some particles being arranged table-wise, then: (i) the table is nothing over and above the particles in that

⁸⁵ Rosen (2017: n.3) proposes the replacement ‘hardly anything over and above’; cf. Audi 2012:709, Bennett 2017:222.

there is no worldly distinction between truths about the table and truths about the particles, and (ii) the table is no addition to being in that it does not exist perspicuously.

3.3.3 Realism

An important theme emerges from the discussion thus far: in a particular sense, generators take a realist (or inflationist) attitude towards the non-fundamental, whereas reducers take an anti-realist (or deflationist) attitude.

This joint between the two approaches is striking, but spelling it out takes care: ‘realism’ is a slippery word which is used in many ways. To my mind, the difference in question is best articulated in terms of perspicuity: when reducers target a non-fundamental proposition, they treat it as a non-perspicuous truth, whereas when generators target a non-fundamental proposition, they treat it as a worldly fact — and hence, view the corresponding truth as perspicuous.

This is vivid in the ontological case. Generation enables a ‘permissivist’ position, according to which the entities in question (e.g. mereological wholes) are ‘real’ but non-basic: truths about them are perspicuous. By contrast, reduction supports a ‘minimalist’ position, according to which the relevant entities only exist in a non-perspicuous sense. In this way, generators take non-fundamentalia more seriously: they accommodate non-fundamental truths about them by positing corresponding entities in reality itself, whereas reducers accommodate the same truths without positing such entities.

But the difference is more general than this: generators are not just realist about what exists non-fundamentally but, more broadly, about how things are. Consider nominalists who take certain ‘worldly’ predicates, such as *has mass*, more seriously than ‘made-up’ predicates, such as *is grue*. This is not an ontological matter: it does not concern whether the corresponding properties exist, since this is denied in each case. Rather, it is ‘ideological’: they hold that only some predicates reflect ‘real’ (not merely representational) aspects of the world. This difference in attitude may manifest in their approach to non-fundamental propositions expressed using these predicates. For example, they might naturally say that [x has 1kg mass] generates [x has mass], but not that [x is green and observed before t] generates [x is grue], since the world provides no distinctive fact involving x’s grueness. As a ‘made-up’ predicate, *grue* is an unsuitable means of specifying targets for generation; truths about grueness are not ‘self-standing’ but collapse into truths about color and observation times.⁸⁶

There is thus a general (not merely ontological) sense in which generators are more ‘realist’ about the non-fundamental. This realism shouldn’t be confused with conservatism: the view that our ordinary conception of the world is largely correct. Realism (in the sense I have been discussing) concerns the status of the non-fundamental explananda one accepts, whereas conservatism concerns the prior question of which putative explananda should be accepted. Insofar as neither generators nor reducers regard the acceptance of non-

⁸⁶ Since a non-reifying version of the generation approach is available, the dispute is not best cast as an ontological question concerning the fact [x is grue].

fundamental explananda as detracting from the simplicity of their theory, conservatism carries a strong pull for both.⁸⁷

Nonetheless, I suspect that generators' realism may to some extent be attributed to a certain reverence for our ordinary conception which is not shared by reducers. Contrast 'Weak Mooreanism', the methodological claim that commonsensical Moorean propositions should be regarded as true (or at least correct), with 'Strong Mooreanism', the claim that there should be a presumption in favor of Moorean propositions being perspicuous. To a reducer's eyes, Strong Mooreanism likely seems misguided: perspicuity is appropriately evaluated by wielding Occam's razor in the light of our best scientific theories (Sider 2011:13) — or by exploiting its connections to surrounding notions — not merely by reflecting on common sense.⁸⁸ However, something like Strong Mooreanism — together with the idea that many Moorean propositions are non-fundamental — may implicitly guide the realist commitments of some generators.

This methodology may be motivated by reflecting on certain 'super-Moorean' propositions:

I exist.

I am conscious.

⁸⁷ Although, highbrow ontologists may defend radical-sounding ontological theses whilst maintaining the correctness of mundane existence claims.

⁸⁸ Moore himself would likely oppose Strong Mooreanism, given his openness to the 'analysis' of propositions about the external world in terms of sense-data (1925: §IV).

There is an external world (things outside me).

Time passes (the past was present, the future will be present).

I persist (the person that stands up is the one that sat down).

I control my actions.

Some actions are morally wrong.

These propositions are commonsensical in the extreme way of being utterly central to our shared ‘sense of reality’, so that denying them would, as Jerry Fodor (1990:156) puts it, be ‘the end of the world’. The idea that such beliefs distort the way things really are — that they are somehow metaphysically off-track — can be profoundly disturbing.⁸⁹ It is therefore tempting to regard them as deep insights rather than mundane truisms; when we reflect on ourselves, or on time, we seem to make direct contact with reality itself. This conviction can be seen to underlie a swath of views which are realist in the sense discussed above: non-reductive theories of persons and mentality, A-theories of time, endurantist theories of persistence, libertarian accounts of free will, non-naturalist versions of moral realism, and so on. These views reject the idea that some super-Moorean beliefs are reducible to any austere, unfamiliar truths.

This lends Strong Mooreanism some credibility; if super-Moorean propositions may be presumed perspicuous, regarding other Moorean propositions as irreducible ‘rounds out’ the picture in a natural way. But the analogous position with respect to generation — that Moorean propositions are presumably ungenerated — is unmotivated; nothing seems to

⁸⁹ Thanks to a referee for raising this issue.

suggest that our super-Moorean beliefs make direct contact with *basic* reality. Here too, therefore, distinguishing reduction and generation affords a better understanding of the disagreement between realist and anti-realist views and their motivations.

3.4. Adjudication

The choice between reduction and generation matters. How, then, can we tell which of the two approaches is better suited to a given case?

One strategy is to exploit the connections to surrounding notions that I used to introduce the distinction. Most prominently, if $\langle p \rangle$ reduces $\langle q \rangle$, then they latch onto the same fact, whereas if $[p]$ generates $[q]$, then (at least typically) $[p]$ and $[q]$ are distinct. Thus, a theory of fact identity would provide a significant means of adjudication. Likewise, given the connection between perspicuity and Siderian structure described above, a theory of the latter will also be powerful: if $\langle p \rangle$ reduces $\langle q \rangle$, then $\langle q \rangle$ must involve some non-structural notions.

However, fact identity and structure are complex and controversial matters, and it is by no means obvious that they can or should be settled prior to adjudicating between the generation and reduction approaches. For example, one might deny that truths latch onto the same facts as their self-conjunctions on the basis that conjunctive facts are generated by their conjuncts, or one might deny that conjunction is a structural notion on the basis

that conjunctive truths reduce to their conjuncts. It is likely that several packages of views about fact identity, structure and the scope of generation/reduction are attractive.

This section aims to provide some independent resources for adjudication. The overall theme which emerges is that, since generation concerns the way the world itself works whereas reduction concerns the way language fits the world, we can discriminate between them on the basis of both the explananda targeted and the nature of the explanatory connections provided.

3.4.1 Explanatory targets

When we try to explain some proposition using the generation approach, we treat it as a worldly fact, outputted by a worldly process. Hence, taking the generation approach towards indeterminate and/or context-relative propositions involves positing some form of worldly indeterminacy and/or context-relativity. Insofar as reality is plausibly determinate and absolute, this provides good reason to favor the reduction approach in such cases.

Take, for example, the truth that Tibbles the cat exists. It is uncontroversial that there is no precise cat-like object uniquely worthy of the name 'Tibbles'; for many particles around Tibbles's boundary, it is indeterminate whether they are part of Tibbles (think of hairs in the process of being shed). Hence, if <Tibbles exists> latches onto a single fact, it must involve an imprecise cat. But imprecise cats are implausible posits: following Lewis (1999)

and McGee & McLaughlin (2000), we should locate the indeterminacy in our language rather than reality itself.⁹⁰

Hence, before we can apply the generation approach, we must first associate this truth with some set of facts: its admissible precisifications. For example, we might associate it with some facts concerning the existence of a range of cat-like micro-aggregates of particles in Tibbles's vicinity. This is an instance of reduction: <Tibbles exists> is being reductively explained in terms of its admissible precisifications.⁹¹ The association of Tibbles with some cat-like micro-aggregates — as opposed to some cat-like spacetime regions, or some cat-like undulations across the gravitational field, say — is an explanatory hypothesis about what in reality supports <Tibbles exists>.

Similarly, take Lewis's (1983b) counterpart-theoretic account of de re modality. On this view, x's possibly being P is metaphysically explained by x's bearing some counterpart relation R to some y which is P. According to Lewis, our de re modal talk is indeterminate between many candidate counterpart relations. For example, there is no determinate answer to the question: how tall could Trump have been? Candidate counterpart relations differ on the height of the tallest Trump-counterpart. For this reason, it would be awkward to view Lewis's account as describing the generation of worldly modal facts, including

⁹⁰ For defense of worldly indeterminacy, see Barnes & Williams 2011, Wilson 2013. Defending the general idea is one thing, defending its application to particular cases quite another.

⁹¹ Depending on the semantics for indeterminate referring expressions like 'Tibbles', it might be that each admissible precisification fully explains <Tibbles exists>, or that all the precisifications collectively do. Either way, it is indeterminate how <Tibbles exists> reduces, since it is indeterminate which precisifications are 'admissible' (i.e. which truths latch onto the same fact as it).

‘indeterminate facts’ about Trump’s possible height. It is much more naturally viewed in terms of reduction: as providing a way to understand modal talk in terms of an entirely non-modal reality. On this interpretation, it is not that a modal truth latches onto some indeterminate aspect of reality, but that it is indeterminate which aspect of reality it latches onto.

Another reason to view Lewis’s account of modality as exemplifying reduction is that he takes the counterpart relation to be context-relative: in one context, our modal talk is underwritten by counterpart relation R1, and in another, it is underwritten by relation R2. Taking the generation approach in this case would therefore require the strange idea that different modal facts are generated in different contexts — or perhaps, that the modal facts fragment into several realms, each generated by the facts involving a corresponding counterpart relation. The reduction approach, by contrast, requires only facts involving the various counterpart relations, with the context-relativity built into how our modal talk latches onto them: the world remains both absolute and unfragmented.

For further illustration of context-relativity, consider claims about simultaneity, such as that the clock strikes five in New York as the clock strikes ten in London. In light of the Special Theory of Relativity, reality itself has no absolute simultaneity structure, and in this sense such statements do not straightforwardly correspond to any facts. A hot-headed reactionary might hold that this makes simultaneity claims either false or meaningless, but another view counsels cool-headed conciliation: physics just tells us something important

about what makes these claims true.⁹² In particular, simultaneity talk is in fact supported by a three-place relation between two events and a contextually supplied reference frame. The context-relativity needn't be built into our linguistic understanding, since in ordinary contexts our frames of reference (determined by our motion) agree on what is (approximately) simultaneous: we move through practically the same specious presents.⁹³

We could understand this view as fragmenting reality into separate realms of simultaneity facts, one per reference frame, with the facts within each realm generated by the facts about which pairs of events stand in the relevant three-place relation R to the corresponding reference frame (as suggested by Fine 2005: §10).⁹⁴ But it is much more naturally understood in terms of reduction, with truths within a context reducing to truths about pairs of events standing in R to the contextually supplied reference frame. Instead of fragmenting reality itself, this approach simply understands simultaneity talk as latching onto reality in a context-relative way, with reduction revealing how the context contributes to determining which fact is latched onto.

⁹² Even supposing that all such statements fall short of truth, there is surely a distinction to be made between those which are apt or 'approximately true', such as the example in the text, and those which are straightforwardly false. The explanatory demand might then be reconceived as concerning this weaker status.

⁹³ More generally, appeal to context is much less costly in theories of reduction than in ordinary semantics, since it needn't be accessible to competent speakers.

⁹⁴ For similarly fragmentalist proposals, see Hare 2009, Lipman 2016, and Spencer 2016.

3.4.2 Explanatory goodness

When evaluating the plausibility of an explanation, we are guided by certain epistemic principles. Two constraints seem especially important. First, explanations should be systematic: they should invoke reasonably concise general principles that subsume many cases at once, rather than proceeding on an ad hoc, case-by-case basis. For example, there is a function which maps, for each object x , the truth about x 's color to the truth about x 's location. But since this function does not obey any simple rule, it is not plausibly taken to describe an explanatory connection. Second, explanations shouldn't be arbitrary: the general principle invoked ought to be privileged in some way over relevant alternatives. For example, the mereological principle that composition always and only takes place on Mars is implausible because it unjustifiably assigns Mars a special status.⁹⁵ (Note that this principle seems acceptably systematic, so systematicity and non-arbitrariness come apart.)⁹⁶

These constraints apply differently in the case of generation and reduction, however. Generative connections describe how worldly entities are built up from others, whereas reductive connections describe how our non-perspicuous conception projects onto reality. Consequently, what it is for these connections to be systematic and non-arbitrary looks quite different in the two cases. This provides another means of adjudicating between the

⁹⁵ For arbitrariness arguments, see Lewis 1986a:212-3; van Inwagen 1990:126-7; Sider 2001:§4.9. For discussion see Korman 2010, Fairchild & Hawthorne 2018, Builes 2021.

⁹⁶ It is an open question how they are connected. From one perspective, systematicity is a matter of being subsumed under a simple theory, and non-arbitrariness is a matter of being subsumed under a theory which is simpler than the relevant alternatives.

approaches: an explanation may be plausibly systematic and non-arbitrary on one approach but not on the other.

With respect to systematicity: in the generative case, this constraint reflects the assumption that the world works systematically, with respect to both dynamical evolution through time and metaphysical generation across levels (Kment 2014:5). We should expect that, as in the causal case, all generation ultimately derives from some concise list of ‘root principles’ (perhaps: principles of logic and/or connecting determinates to determinables, and set formation and/or mereology). Hence, generation should be ‘nomicallly’ systematic in the way that causation is — backed by some sparse ‘lawbook’ of simple principles.⁹⁷

By contrast, in the reductive case, the systematicity constraint reflects the assumption that language works systematically, with respect to both ordinary meaning and metaphysical analysis (Sider 2011:118). We should expect that, as with ordinary semantics, reduction may be subsumed under a reasonably systematic meta-semantic account, which explains why a notion reduces in the way it does in terms of its meta-semantic profile (e.g. its history of use, its causal/nomic connections, etc.). Hence, reduction should be ‘meta-semantically’ systematic in the way that semantics is — backed by some ‘dictionary’ of metaphysical analyses (together with some principles of meaning composition).

⁹⁷ Witness the widespread practice of formulating general principles describing how facts of certain kinds are grounded. Schaffer (2017b) argues that metaphysical explanation requires ‘laws of metaphysics’, and Wilsch (2016) offers an account of metaphysical explanation in terms of these laws. These authors are best interpreted as focusing on generative explanation.

To illustrate the difference between nomic and meta-semantic systematicity, consider the task of explaining conscious experiences in physical terms. Suppose this task is accomplished by pairing phenomenal properties with functional properties: *x* is in pain because *x* is in a state which plays nomic role R1, *x* sees red because *x* is in a state which plays nomic role R2, and so on. If describing generation, this connection seems problematically case-by-case: the list cannot be subsumed under any root principle of functional-to-phenomenal generation, since the outputted phenomenal properties cannot be concisely defined in terms of the inputted functional properties.⁹⁸ However, if describing reduction, it may be suitably systematic: the list might be subsumed under a meta-semantic account which explains, in terms of the meta-semantic profiles of our phenomenal concepts, why they reduce to the corresponding functional properties. Phenomenal truths would then reduce compositionally via the metaphysical meanings of the predicates they involve.⁹⁹

With respect to arbitrariness: in the generative case, it is costly to posit a principle describing how the world works which fails to stand out over relevant alternatives. For example, it is costly to posit a mereological law which specifies how tightly some parts must be stuck together in order to compose, just as it is costly to posit a dynamical law which specifies some threshold mass-value above which particles no longer move

⁹⁸ As Levine (1983:357) remarks: ‘there seems to be nothing about C-fiber firing which makes it naturally “fit” the phenomenal properties of pain, any more than it would fit some other set of phenomenal properties’.

⁹⁹ For further discussion, see my ‘Against Grounding Physicalism’ (ms).

according to 'F = ma'. Generative connections should have a kind of 'objective' privilege, conferred by reality itself.

By contrast, in the reductive case, there should be some account of why it is that we represent one aspect of reality rather than some relevantly similar aspect. Such an account may appeal to somewhat parochial features of our own 'window on the world', where this would be an inappropriate source of privilege in the context of generation. For example, any reduction of our talk of composite objects which specifies how tightly some parts must be stuck together will be implausible because nothing in that talk seems to distinguish any particular degree of tightness. Reductive connections should have a kind of 'perspectival' privilege, which partly reflects our own vantage point.

By analogy, compare two constructions of the natural numbers. They can be built up from powers of ten: for example, 476 is $(4 \times 10^2) + (7 \times 10^1) + (6 \times 10^0)$. This construction is perspectivally but not objectively privileged: it is highlighted by our decimal system of representation, but nothing in the numbers themselves distinguishes it. Alternatively, they can be built up from their prime factors: for example, 476 is $2^2 \times 7 \times 17$. This construction is objectively but not perspectivally privileged: it reflects the deep nature of the numbers themselves, but appears strangely random from our own perspective.

To illustrate the difference between objective and perspectival privilege, consider the best system account of laws (Lewis 1983a, Loewer 1996). According to this account, laws are the theorems of the deductive system which best summarizes the 'Humean mosaic': the

complete collection of ‘local matters of particular fact’. A system is good to the extent that it balances various virtues, especially simplicity (roughly, length in fundamental terms) and informativeness (roughly, amount of worlds excluded). Thus, its being a law that *p* is metaphysically explained by *p*’s being a theorem of the best system (which is explained in turn by the Humean mosaic itself).

Arbitrariness provides strong reason to view this explanation as reductive.¹⁰⁰ As Lewis (1994a:479) admits, how the virtues are defined and traded off may partly be a psychological matter: the account is to be fleshed out in accordance with the role that laws play in our own cognitive lives and in the practice of science. What counts as the ‘best’ system may differ for aliens with a different psychology who practice science differently. We may therefore distinguish the ‘human’ best system from the ‘alien’ best system, with neither providing an ‘objectively correct’ notion of law.¹⁰¹ Thus, it seems problematically arbitrary to posit that *p*’s lawhood is generated by its belonging to the human best system, since this explanatory connection lacks objective privilege: why would nomic facts be attuned to our conception rather than the alien conception? But there is no analogous issue with positing that *p*’s lawhood reduces to its belonging to the human best system, since this

¹⁰⁰ Indeed, Lewis himself seems to have reduction in mind: he describes Humean Supervenience as the view that ‘all there is in the world is a vast mosaic of local matters of particular fact’ (1986b: ix).

¹⁰¹ Lewis hoped that the human best system would be best ‘under any standards’ (for critical discussion, see Massimi (2017) and Gómez Sánchez (forthcoming)). If this were plausible, the corresponding notion of law might be objectively privileged enough to be generated rather than reduced.

reduction is perspectively privileged: we would have an explanation — in psychological terms — for why it is that we employ this notion of lawhood rather than the alien notion.¹⁰²

3.5. Redundancy

So far, I have presented the distinction between generation and reduction, illustrated the significance of its impact on metaphysics, and suggested some ways of adjudicating between the two approaches. This final section argues that both approaches have a role to play in explanatory metaphysics; we can't rely on one of them alone.

3.5.1 Against pure generation

According to pure generators, all metaphysical explanation is (or is backed by) generation amongst facts.¹⁰³ I know of no explicit defenses of this extreme position, and it might be that no generator would, upon reflection, endorse it. Nonetheless, the rhetoric of many generators, together with their readiness to widely apply their framework, suggests that it may not be uncommon as an implicit presupposition.¹⁰⁴ Sociology aside, discussing the

¹⁰² In order to avoid arbitrariness, generators might posit plenitudinous generation: for example, in the case of the best system account, they might posit many principles of the form 'if p is a theorem of the best _{i} system, then p is a law _{i} '. This strategy still requires supplementation by (merely perspectively privileged) reduction: once we have generated the various laws _{i} , it remains to specify which of them our own notion represents.

¹⁰³ At least, all explanation revolving around fundamentality and ground (see n.46).

¹⁰⁴ Indeed, I suspect that many generators would subscribe to Schaffer's (2017b:303) characterization of metaphysical explanation as 'an explanation backed by grounding relations' (where 'grounding' refers to generation).

view is theoretically important for my purposes: it helps to bring out what makes reduction distinctively important.

Now, I have already argued above (§3.4.1) that certain imprecise truths, such as <Tibbles exists>, and certain context-relative truths, such as those concerning simultaneity, are better treated in terms of reduction than generation, at least insofar as there are no corresponding imprecise or context-relative facts. In each case, generators must find some precise, absolute facts associated with the truth in question. But this prior step is reductive: the non-perspicuous imprecise/context-relative truths seem metaphysically explainable in terms of perspicuous truths which reveal the associated facts. These kinds of cases may have already convinced you that pure generation is untenable. Nonetheless, in this section, I wish to add a further kind of example, before considering what I take to be the pure generator's most natural response.

Consider a case of 'cross-level conflict': non-fundamentally, this table occupies a continuous, table-shaped region; fundamentally, this region is largely empty space, inhabited only by a sparse network of particles. How can it be both a non-basic fact that the region is occupied and a basic fact that it is not? Their obtaining at different 'levels' seems irrelevant — as Martin Lipman (2018:597) puts it: 'stacking the facts does not help remove the conflict.'¹⁰⁵ Thus, in order to explain the region's being occupied, generators must first associate it with an appropriate fact, by clarifying the sense in which the region

¹⁰⁵ Lipman (2018) argues that cross-level conflict makes trouble for 'fundamentality-based metaphysics' in general, but does not distinguish generation from reduction.

is ‘occupied’ despite being largely empty. But again, this prior association of the explanandum with a fact appears reductive: it reveals, in more perspicuous terms, what supports the truth in question.¹⁰⁶

Pure generators might insist that, in each of the cases I have highlighted — involving indeterminacy, context-relativity, and cross-level conflict — the association with a fact is not part of the explanation itself, but rather provides a suitably clear gloss on the explanandum being targeted. The idea is that the semantics in question does not amount to reduction: it is merely ‘descriptive’ rather than genuinely explanatory.

This response doesn’t do justice to the phenomena in question. Take Tibbles. The association of ‘Tibbles’ with its precisifications seems importantly different from merely being taught the name, or having the name disambiguated, neither of which seems to yield the same kind of metaphysical understanding. If you were to say, upon learning that ‘Tibbles’ refers to some cat in the vicinity, ‘Ah, now I see – Tibbles exists in virtue of *that cat existing!*’, you would not have advanced the cause of explanatory metaphysics. Nor should the understanding yielded by precisifying ‘Tibbles’ be conflated with that attained when an ordinary Frege puzzle is resolved. Having realized that Tibbles is the very same cat already known as ‘Hamish’, you would not satisfy metaphysicians by informing them that for Tibbles to exist is just for Hamish to exist. The notion of perspicuity allows us to capture the important difference between these ways of understanding <Tibbles exists>: it

¹⁰⁶ Again, there’s a radical response (akin to making reality imprecise and/or context-relative): pure generators might embrace an inconsistent reality.

is only when it is glossed in more perspicuous terms that metaphysical explanation is achieved.

Similarly, pure generators might see the context-relative ‘explanation’ of truths about simultaneity as akin to the merely descriptive association of a context-dependent meaning with straightforwardly indexical truths, such as <I exist>. It does not seem especially natural to think that <I exist> may be understood in any metaphysical sense in terms of <x exists>, where x is the speaker/thinker in the relevant context. It seems more natural to say that these are just the same truth (in the context), and so no explanatory progress has been made.¹⁰⁷ By contrast, however, it is natural to think that the truth that two events are simultaneous may be metaphysically understood in terms of the truth that they are simultaneous relative to a contextually determined reference frame, and not at all natural to think that these are just the same truth (as opposed to representing the same fact in the relevant context). It is a substantive discovery about reality (not merely about language) that our simultaneity talk is misleading in this way.

Finally, on the proposed response, pure generators would conflate the resolution of cross-level conflict with ordinary polysemy: in one sense of ‘door’, she walked through the door, but in another (related) sense, she didn’t. In this case, clarifying that what she walked through was the airy gap within the doorframe pedantically spells out the speaker’s original intention. The task of explicating the sense in which the table-shaped region is occupied is

¹⁰⁷ Although, even here it might be argued that it is a genuine metaphysical discovery (albeit one typically made early in life) that the I-concept picks out something which may also (and indeed, more perspicuously) be regarded third-personally.

importantly different. When we discover that this occupation is a matter of the dynamical laws underwriting the exclusion of other (macroscopic) objects, we haven't merely described what was meant all along: we have explained what in reality supports the truth in question.

Some generators might still prefer to use 'explanation' in a narrower 'worldly' sense, according to which the reduction step may be a necessary prelude to explanation but is not itself The Real Thing. I don't wish to quibble about terminology; the substantive point is that characterizing worldly generation relations is not the only way to satisfy the aims of 'explanatory' metaphysics. On this narrower usage of 'explanation', neither the association of truths about pain with truths about C-fibers nor the association of truths about laws with truths about the Humean mosaic would constitute proposals for metaphysical explanation. But these hypotheses are surely intended as contributions to explanatory metaphysics (indeed, as crucial parts of the defense of overarching metaphysical positions such as physicalism and Humeanism), and if true, they yield the kind of understanding that is its distinctive aim.¹⁰⁸

¹⁰⁸ Let me briefly outline a further argument against pure generation. I argued above (§3.4.2) that generation ought to be backed by a concise lawbook of general principles. But it is hard to see how such principles could back the generation of truths involving names and/or atomic predicates, unless supplemented by some auxiliary reductions. For example, by what suitable general principle could the existence of Obama be outputted? Perhaps there is a suitable principle governing the generation of mereological fusions; but this principle is only applicable once Obama has been 'located' within the mereological hierarchy. This location step amounts to a reduction of <Obama exists> in perspicuous mereological terms. (See Rosen 2010, Glazier 2016 and Wilsch 2016 for discussion of metaphysical laws requiring 'auxiliary identities'.)

This discussion has revealed that the reduction approach has resources which the generation approach lacks. Reduction can advance our understanding by focusing on features of representation itself, uncovering the ways in which it makes contact (or fails to make contact) with reality.¹⁰⁹ Reducers can appeal to semantic indeterminacy, avoiding an imprecise reality, to the representational context, avoiding a context-relative reality, and to fundamental and non-fundamental uses of an expression,¹¹⁰ avoiding an inconsistent reality. These resources have wide potential application: our ordinary conception of the world abounds with imprecision, and may turn out, upon analysis, to be thoroughly context-relative and to conflict radically with fundamental reality.

3.5.2 Against pure reduction

According to pure reducers, all metaphysical explanation is (or is backed by) reduction amongst truths.¹¹¹ The most direct way to challenge this view is to identify certain irreducible (i.e. perspicuous) truths which plausibly represent generated facts. This section argues that, even allowing for the success of an extremely ambitious reductive project — in which all truths are reduced to logically simple truths about ontologically simple entities — such cases are hard to avoid. In conceding so much to reducers, this argument departs significantly from the existing arguments for the importance of generation in the literature

¹⁰⁹ Cf. Sider (2011: §7.9) on the advantages of semantic ascent in the context of grounding.

¹¹⁰ Cf. Dorr 2007:33.

¹¹¹ As with pure generation, I am not here targeting a position that has (at least to my knowledge) been explicitly defended. My discussion aims to foreground the distinctive importance of generation.

(such as those in Schaffer 2009, Rosen 2010, and Bennett 2017), which either do not consider the reduction approach or effectively assume that it is much more limited in scope.

Schaffer (2004) argues that some ‘macro-properties’, such as being an H₂O molecule, ‘serve as the ontological basis for linguistic truths’. (In his words: ‘Molecules aren’t merely manners of speaking.’) On this view, the truths about such properties are perspicuous, and yet correspond to higher-level facts which are generated from microphysical facts. Pure reducers must reject this picture. This requires faith in an ambitious enterprise: all truths, stated both in ordinary vocabulary as well as that of ‘higher-level’ sciences, will need to be shown to be reducible to fundamental physics. Embarking on such an enterprise is certainly brave, but may not be foolish (for words of encouragement, see Sider 2011: §7.11.1).¹¹²

However, we needn’t defend anything as anti-reductionist as Schaffer’s picture to refute the pure reducer. Instead, we can grant that all truths can be successfully reduced to fundamental physics, and argue that even some ‘fundamental’ physical facts are plausibly generated.

One remaining obstacle to pure reduction is logical complexity. If reduction requires necessary equivalence, it is hard to see how all disjunctive truths, or all quantificational truths, could reduce to logically simple truths. But if there are logically complex facts — such as [Fa or Fb] and [there is an F] — then they would cry out for explanation in terms

¹¹² For further discussion, see Schaffer (2013) and Sider (2013a: §3).

of logically simple facts — such as $[Fa]$ — and this explanation is naturally taken up by the generation approach. Hence, pure reducers need an alternative conception of reduction which does not require necessary equivalence, and it is unclear that any such conception is legitimate.¹¹³

Set this challenge aside, however: suppose that reality consists in particular details, non-perspicuously described by logically complex patterns.¹¹⁴ A further obstacle remains: ontological complexity. The perspicuous truths plausibly involve composite entities of some sort, such as mereological fusions. For example, to completely describe spacetime, we plausibly need to specify the pattern of topological properties and/or distance relations. As Sider (2013c: §11) points out, the topological notion of ‘openness’ is a predicate of regions, and distances in general relativity are path-dependent (the distance between two points is the length of the shortest path between them). So our perspicuous description plausibly includes sentences of the form ‘region R is open’ and ‘the distance along path P is x’. But the corresponding facts are not plausibly basic: facts involving composite entities

¹¹³ Pure reducers might adopt Fine’s (2013) ‘D-project’, which aims to ‘describe’ everything there is to be described, rather than ‘expressing’ everything there is to be expressed. For example, disjunctions and existentials might be taken to describe what their true disjuncts/instances describe. (See Sider 2013b:742 for criticism.) More generally, they might look to forms of semantics ‘beyond the biconditional’, such as truthmaker semantics, which gives meanings in terms of ‘verifiers’ and ‘falsifiers’ (see e.g. Fine 2017, Elgin forthcoming). The potential advantage of such a strategy is that the relations connecting truths to reality can themselves play a significant explanatory role, alongside the reality being represented.

¹¹⁴ This may require an atomic totality fact of the form ‘ $T(a, b, \dots)$ ’ to reduce generalizations (Russell 1918:503; Fine 2012:62). Pure reducers might regard negative truths as perspicuous, with the corresponding negative facts ungenerated.

like regions and paths are plausibly generated by facts involving the sub-regions and sub-paths (and perhaps ultimately the points) which make them up.¹¹⁵

Pure reducers face a dilemma here. They can either (i) deny these plausible generation claims and insist that regions and paths are fundamental in the same sense that points are, or (ii) provide alternative primitive notions which apply to points directly (rather than regions or paths), and which all truths about open regions and path-dependent distances reduce to.

Proponents of (ii) could, for example, take openness and distances to be captured by applying plural predicates to continuous infinities of points. For example, truths of the form ‘region R is open’ would reduce to infinitary truths of the form ‘points p, q, ... are open’. However, since it requires perspicuous plural predicates, this exchanges ontological complexity for a significant jump in ideological complexity, and — following Sider (2011: §9.15) — we may doubt that this sort of tradeoff is worthwhile.¹¹⁶

Besides, the explanatory connections we were worried about survive the transition from ontological to ideological complexity. The openness of certain regions is determined by their being the union of certain other open regions, and the distances along paths are determined by the distances along various sub-paths, irrespectively of whether these facts

¹¹⁵ Similarly, the richness of set theory suggests that mathematical reality cannot be completely described without using sets (or the like), and yet, facts involving sets would appear to be generated by facts involving their members.

¹¹⁶ Sider is concerned with plural quantification, but analogous points apply to plural predication.

involve monadic properties of composite entities or plural properties of their constituents. These pervasive patterns indicate explanatory connections, with the more ‘local’ truths (pertaining to smaller regions/paths) plausibly explaining the less local truths which they determine (pertaining to larger regions/paths). But these explanatory connections are better treated in terms of generation than reduction, since ‘locality’ seems irrelevant to perspicuity: it seems implausible to claim that $\langle p, \dots \text{ are open} \rangle$ and $\langle q, \dots \text{ are open} \rangle$ are perspicuous whereas $\langle p, q, \dots \text{ are open} \rangle$ is not!¹¹⁷

Sider (2011:146) points out that, at least in some cases, this sort of generation would not be well-founded: in a continuous space there are no smallest open regions and no shortest paths. But it is unclear that this is reason to deny the generation in question, especially in this context, where we are supposing that it merely supplements the (well-founded) explanatory structure of reduction. Methodologically speaking, we must decide whether general principles like well-foundedness should be prioritized over intuitions about particular cases. But besides, denying any generation amongst the atomic facts would preserve well-foundedness at the cost of violating another (to my mind, at least as plausible) general principle: that the fundamental facts ought to be non-redundant (i.e., free of necessary connections). If all openness facts are ungenerated, for example, then the basic facts involve pervasive redundancy.

¹¹⁷ The point here is not that there are necessary connections between the atomic truths which must be explained by invoking generation. Pure reducers may instead seek to explain necessary connections via some reduction of modality. The point is rather that there are general patterns amongst the atomic truths, of the kind which indicate explanatory connections (cf. Kim 1993:167). This is consistent with the necessary connections explaining the generational connections, or something else explaining both — for example, following Sider 2011: §12.5 and Wilsch 2016 respectively, we might account for both necessity and generation in terms of ‘metaphysical laws’. Thanks to a referee for helping me to clarify this argument.

Recognizing reduction and generation offers an attractive diagnosis of the conflict between well-foundedness and non-redundancy. An intuitive motivation for well-foundedness is that all truths have some ultimate basis in reality: this seems to fit reduction better than generation.¹¹⁸ Some collection of geometric facts ought to serve as foundation for all the geometric truths, but positing some bottomless generative hierarchy within this foundation seems no more problematic than, in the causal case, a universe without initial conditions (cf. Fine 2001:27). An intuitive motivation for non-redundancy is that necessary connections between contingent truths ought to be explicable, and that ground (either through their grounding one another, or sharing some common ground) provides a plausible explanation. This tells against necessary connections between basic geometric facts, but not between perspicuous geometric truths, since the latter connections may be explained in terms of the generation of the corresponding facts.

In sum: even granting that all truths reduce to physical, logically atomic truths about simple entities, it still seems plausible that some of the corresponding facts are generated. Generative structure is a hardy plant, surviving even the harshest desert.

In characteristic style, reducers might countenance truths about generation whilst denying that reality itself features any generative structure. Indeed, this kind of position is defended by Sider (2011:145): ‘we reductionists can take the same attitude towards metaphysical

¹¹⁸ Strictly speaking, it fits reduction as supplemented with the representation relation, which bridges the gap between facts and truths. As mentioned above (n.28), there may be no perfectly perspicuous truths. But we may still demand that all truths be supported by some foundation of facts, even if these facts have no perfectly perspicuous representation.

causation [i.e. generation] as towards everyday causation: it reduces in some way to facts that don't involve metaphysical causation'. For example, Sider suggests, we might reduce generation to metaphysical laws (and reduce these laws in the style of the best system account). Thus, reducers might concede that logically/ontologically/ideologically complex facts are 'generated', but deny that this way of talking is perspicuous.

I suspect that generators would not regard this 'Humean' position as a genuine reconciliation: they typically view truths about generation as perspicuous (although they may not put it that way). Nonetheless, it does amount to a retreat from pure reduction, in favor of a hybrid approach to metaphysical explanation. First, Humeans treat reduction and generation alike, since they cannot plausibly hold that truths about reduction (which feature non-perspicuous truths as constituents) are themselves perspicuous. On this picture, all explanatory structure disappears from reality. Second, the crucial question as I see it is: what makes metaphysical explanations explanatory? Humeans recognize two sources of explanatoriness: generation amongst facts and reduction amongst truths. For the Humean — unlike the staunch pure reducer — not all metaphysical explanation is tied to perspicuity.

3.6. Conclusion

I have laid out the distinction between two broad approaches to metaphysical explanation, and have demonstrated some of the ways in which it matters. I hope to have made it clear

that whenever one is engaged in explanatory metaphysics, the question of approach should be front and centre.

I have also suggested some ways of adjudicating between the two approaches, and argued that neither can sustain metaphysical explanation alone. This motivates a ‘hybrid’ approach which employs both. The hybrid approach recognizes two dimensions of relative fundamentality: perspicuity amongst truths, and basicness amongst facts. But it does not make explanatory metaphysics a disconnected enterprise. Rather, generation and reduction complement each other: reduction ‘lays bare’ the facts which generation targets, whilst generation provides the ‘raw material’ for reduction. For example, an explanation of water’s existence in terms of hydrogen’s existence might involve two components: water’s existence reduces to H_2O ’s existence, and hydrogen’s existence generates H_2O ’s existence.¹¹⁹

The hybrid approach is a broad church, with room for disagreement about where to place the line between generation and reduction (and, correspondingly, their relative importance to the overall project). At one extreme is the view that our conception is more or less perspicuous: reduction is largely just ‘tidying around the edges’, with most interesting action involving the formulation of ‘heavyweight’ principles of generation. At the other extreme is the view that our conception is radically non-perspicuous: the main task is to recover this conception from an unfamiliar reality, with the only generation principles being ‘lightweight’ in character (purely logico-mathematical, perhaps).

¹¹⁹ In future work, I plan to extend this hybrid approach to causal and mathematical explanation.

Recognizing both approaches allows us to raise new questions and to sharpen some old questions. For example, many contemporary metaphysicians have considered whether fundamentality is fundamental, and how ground should be grounded. But generators are asking: ‘are facts involving generation and basicness basic, and if not, how are they generated?’, whereas reducers are asking the very different question: ‘are truths about reduction and perspicuity perspicuous, and if not, how are they reduced?’ We can also ask whether there are basic facts involving perspicuity, and whether truths about generation are perspicuous.

The answers to these questions are likely to mark significant divisions within the hybrid approach, and no doubt further important divisions remain to be made. Nonetheless, what unifies the hybrid approach — as I see it — is the recognition of two approaches to metaphysical explanation which are both different in important ways and, in different ways, important.

4. Generalism Without Generation¹²⁰

4.1. Introduction

According to generalism, the world is fundamentally general — ultimately, there are no individuals: no particles, no spacetime points or regions; nothing. Heuristically, we can think of generalism as positing that the world is fundamentally described by a single complex sentence of the form:

$$\exists x \exists y \dots Q(x, y, \dots)$$

where Q is some qualitative description (one which involves no essential reference to any individuals).¹²¹ Generalists hold that descriptions like these correspond to complete ways for the world to be. This is only a heuristic — in fact, generalists have sought to give non-quantificational fundamental descriptions. But I will stick to it in what follows; nothing will turn on the subtleties I thereby neglect.

Generalism is revisionary. To put it mildly, individuals are central to our conception of reality: our talk and thought is more or less entirely about them. Even when we use general

¹²⁰ Thanks to the Fall 2018 Rutgers Metaphysics Group, the Spring 2020 Rutgers Dissertation Seminar, Karen Bennett, Verónica Gómez, Daniel Murphy, Jonathan Schaffer, Alex Skiles, Jason Turner, three anonymous referees, and especially Ted Sider.

¹²¹ This allows for Q to feature identity and distinctness, which may not seem purely qualitative. In future work, I plan to explore ‘pure qualitativism’, on which fundamental descriptions involve neither individuals nor identity.

descriptions which do not mention any particular things, we understand these descriptions as concerning things.¹²² Orthodoxy assumes that our fundamental theories will be similarly individualist (albeit concerning only fundamental things).

Why, then, take generalism seriously? One important reason, as Jason Turner (2011) and Andrew Bacon (2019) suggest, is just that it can be illuminating to see what happens when we reject one of our most entrenched assumptions — even if it only teaches us why that assumption is justified. However, I believe that there are powerful reasons — presented by Shamik Dasgupta (2009) — to view generalism as more than just radical experimentation.¹²³

First, Dasgupta argues that fundamental individuals would be both physically redundant and undetectable (much like absolute velocities in Newtonian mechanics). If there are fundamental individuals, then there are distinct physically possible worlds differing only by permutations of these individuals. But Dasgupta argues that the laws of physics are blind to such differences (and hence, no physical instrument could detect them): if a world is physically possible, then its qualitative duplicate is also. Thus, by the principle that we should eliminate physically redundant and undetectable structure where we can, we should

¹²² For this reason, quantificational fundamental descriptions may seem inappropriate for generalists — see Dasgupta 2009, Turner 2011 and Russell 2018.

¹²³ Historically, suspicion of individuals has come in many forms which I will not directly discuss here: for example, the concern that they are objectionable ‘bare substrata’, that they are merely a projection of the noun-verb structure of our language, that they lead to problematic metaphysical questions (involving identity and ontology), and that no two metaphysical possibilities could differ merely by a permutation of their individuals.

eliminate fundamental individuals.¹²⁴ (We can think of Dasgupta’s argument as an empirical reinforcement of the intuitive idea that purely haecceitistic distinctions are ‘distinctions without a difference’.)

A second reason to take generalism seriously is defensive in nature: it can explain the radical divergence between our individualist conception of reality and its fundamental generality. A broad point here is that, insofar as our conception has been shaped by its usefulness for navigating our environment, there seems to be little reason to expect it to match the way things are fundamentally. We are non-fundamental creatures navigating a non-fundamental environment; fundamental metaphysics is therefore akin to the more theoretical parts of physics and mathematics, where our ordinary conceptual scheme has little authority. Arguably, for example, our naïve conceptions of time and of sets should be supplanted by less intuitive theories (the former because of its conflict with relativity, and the latter because of its inconsistency).

Generalists can also offer a more targeted defense. Dasgupta (2009: §4.4) proposes that agents inhabiting a fundamentally general reality would find it pragmatically essential to employ referring expressions as representational ‘hooks’ on which to hang general information. To illustrate, consider belief-updating. Suppose we start with the belief:

“something which is F bears R to something else”.

¹²⁴ Dasgupta (2009) argues further that we can eliminate fundamental individuals, by outlining his ‘algebraic generalism’. The remaining question is whether the apparent gain in parsimony proves worthwhile — see Sider 2020: §§3.14–6.

Now we discover that ‘the something else’ is a G. To represent this new information, we must transition to the whole new belief (which supersedes our previous belief, and cannot be decomposed into simpler beliefs):

“something which is F bears R to something which is G”.

This holism makes generalist representation extremely impractical: information-gathering requires the constant wholesale replacement of a single extremely complex global belief. By contrast, suppose we introduce some individualist labels ‘a’ and ‘b’ to represent our original information as:

“a is F”;

“a bears R to b”.

This atomistic approach facilitates information storage and revision. For example, we can replicate the transition above by merely adding the belief:

“b is G”.

The resulting stock of atomistic beliefs adequately represent the general situation so long as we receive and convey information using rules analogous to the \exists -elimination and -introduction rules of predicate logic. Thus, agents naturally construct an ‘ersatz world’ of

individuals in order to render a general reality representationally tractable. This suggests that generalism would be deeply counterintuitive regardless of its truth.

For these reasons and others, several metaphysicians and philosophers of physics have recently explored various forms of generalism (Dasgupta 2009, 2014a, 2015; Turner 2011, 2017, forthcoming; McKenzie 2014; Russell 2016, 2018; Bacon 2019; Dewar 2019; Glick 2020; Sider 2020). This paper aims to contribute to this undertaking, by addressing generalism's core challenge: recovering non-fundamental individuals from a fundamentally general reality.

I consider generalism a serious and interesting enough hypothesis to make my proposal worthwhile in itself. But it also serves more broadly as a case study in two importantly different approaches to the task of explaining the non-fundamental in terms of the fundamental: one based on 'generation', the other on 'reduction'. Part of my aim is therefore to demonstrate a broader methodological point: the evaluation of theories about what is fundamental is crucially sensitive to the approach to metaphysical explanation that is taken.

4.2. Two Versions of Generalism

There are two importantly different versions of generalism, which may be introduced by analogy with 'atomism': the view that fundamental reality consists in the existence and nature of mereological atoms (objects with no parts).

According to atomism, all composite objects (molecules, tables, people, etc.) are non-fundamental: their existence and nature may be metaphysically explained in terms of an underlying purely atomic reality. But there are different senses in which these objects may be non-fundamental.

‘Permissive atomism’ holds that composite objects are generated by, or built out of, their atomic constituents. *Generation* is a worldly ‘determination’ relation much like causation; the two relations are similarly tied to laws, counterfactuals, and explanation.¹²⁵ Generation differs from causation by holding across ‘levels of reality’ rather than across time, and with metaphysical rather than natural necessity. According to permissive atomism, the most basic level of reality features only atoms; composite objects belong to derivative levels of reality which are generated from this basic level. This view is ‘permissive’ in holding that composite objects are just as ‘real’ as atoms; reality is a rich hierarchy, featuring both composite objects and the atoms which make them up.

Another version of atomism — ‘strict atomism’ — holds that reality is mereologically sparse; the atoms are all there ‘really’ is. Nonetheless, this sparse reality supports our ordinary talk ‘about’ composite objects; this talk is made true by the ways atoms are arranged.¹²⁶ This view countenances no worldly relation between atoms and composite objects — there are no composite objects in reality to stand in such a relation! Instead, the sense in which composite objects are non-fundamental must be understood via

¹²⁵ See Rosen (2010) on ‘metaphysical dependence’, Fine (2012) and Schaffer (2016) on ‘grounding’, Bennett (2017) on ‘building’, and A. Wilson (2018) on ‘metaphysical causation’.

¹²⁶ This can be viewed as a form of ‘conciliatory nihilism’, à la Dorr & Rosen (2002), Cameron (2010), Sider (2013c). (Conciliatory nihilists may disagree about whether claims concerning composite objects are strictly true as opposed to ‘apt’ in some weaker sense.)

representation: only truths about atoms reflect reality's ultimate, intrinsic, objective structure, or 'carve nature perfectly at its joints', whereas talk of composite objects distorts this structure, in ways that reflect our own interests and perspectives 'outside of the metaphysics room'. As I will say, only truths about atoms are 'perspicuous'.¹²⁷ Non-perspicuous truths about composite objects *reduce* to perspicuous truths about atoms, in that talk of composite objects has something like a semantic analysis or paraphrase in terms of atoms.¹²⁸

All atomists agree that the existence of the table before me is metaphysically explained by the way various atoms are arranged. But this is consistent with two quite different pictures. The permissive atomist picture involves two 'portions of reality': a basic portion, consisting of atoms, and a distinct but collocated derivative portion which this basic portion generates, consisting of the table. The strict atomist picture involves two ways of aptly representing a single portion of reality: a non-perspicuous way, in terms of the table, and a perspicuous way, in terms of the atoms, to which this non-perspicuous way reduces.

By analogy, we may distinguish two versions of generalism. 'Permissive generalism' holds that all individuals are non-basic. This view is permissive in holding that individuals belong to reality itself. But all individuals (not just the composite ones!) are confined to the derivative levels of reality: their existence and nature are generated by purely general facts

¹²⁷ Cf. (O'Leary-)Hawthorne & Cortens (1995: §3), Turner (2010:8–9). See also Fine (2001) on propositions holding 'in reality', Sider's (2011) notion of a 'fundamental truth', Russell's (2015) 'objective matters of fact', and deRosset's (2017) 'conciliatory irrealism'.

¹²⁸ The analysis won't belong to ordinary semantics since competent speakers need not have any access to it; see Sider (2011: §7.4) on 'metaphysical semantics'.

which involve no individuals. ‘Strict generalism’, on the other hand, is the view that all talk of individuals (not just talk of composite individuals) is non-perspicuous. Individualist talk is useful and may be true, but it distorts reality’s underlying structure. A purely general reality supports individualist talk in that individualist truths are reducible to general truths.

Permissive and strict generalists agree that the existence and nature of all individuals may be metaphysically explained in purely general terms. But they approach this task quite differently. The permissive generalist takes the generation approach, seeking to describe how facts involving individuals are generated from basic generalist facts. The strict generalist takes the reduction approach, seeking to describe how truths about individuals reduce to perspicuous generalist truths.

There is plenty to say at an abstract level about the distinction between these two approaches to metaphysical explanation.¹²⁹ The distinction rests on the notion of perspicuity (or, reality; or, worldliness) — the idea that a truth may either reflect or distort the structure of the fact it latches onto. I hope to have at least given you a preliminary grip on the two corresponding versions of generalism. Things should crystallize as we proceed: I will argue that permissive and strict generalism differ substantively, and indeed that we have good reason to prefer the latter.

¹²⁹ I explore the distinction in ‘Two Approaches to Metaphysical Explanation’ (ms).

4.3. Against Permissive Generalism

This section presents two challenges for permissive generalism: it is hard to see how generalist facts could generate individualist facts in a suitably systematic and non-arbitrary manner.

4.3.1 Systematicity

Explanation should be systematic: the explanans should be connected to the explanandum via law-like general principles. These ‘connections’ not only subsume the particular case in question, but extend to many relevantly similar cases (both actual and non-actual). For example, when physicists seek to explain why two magnets attract, and why dropped objects fall, they look for general principles which systematize the behavior of magnetic and free-falling bodies respectively. The explanatory project would be deemed a failure if no such principles could be found. If the behavior in question turned out to be utterly haphazard and non-systematic, we would infer that no explanation could be given (at least, no good explanation).¹³⁰

Metaphysical explanation is no exception: a candidate metaphysical explanation is only taken seriously when it may be subsumed under some systematic general principle.¹³¹ For

¹³⁰ Schaffer (2017b) provides several arguments that explanation in general—and metaphysical explanation in particular—requires ‘laws’ in the inclusive sense of ‘counterfactual-supporting general principles’. Note that this constraint does not entail any particular account of explanation (e.g. the deductive-nomological account): for example, it is compatible with holding that explanation is a matter of unification (Kitcher 1981), or of tracking dependence relations (Kim 1994).

¹³¹ Witness the widespread practice of attempting to formulate general principles describing how facts of certain kinds are grounded: e.g., Bennett 2004, Rosen 2010, Fine 2012, Schaffer 2016.

example, consider the claim that a certain ball is red in virtue of being scarlet. One reason this explanation seems plausible is that it may be subsumed under a systematic general theory, according to which it is not only the object in question which is red in virtue of being scarlet, but all scarlet objects, irrespective of their size or shape, and according to which, more generally still, the same holds for any determinate of any determinable, not merely for scarletness and redness. By contrast, consider the proposal that any object has the color which it does in virtue of having the shape which it does. For example, this ball is red in virtue of being round, others are blue in virtue of being round, and yet others are green in virtue of being round (and likewise with objects of other shapes). This theory is implausibly non-systematic: it determines colors on a case-by-case basis, rather than by a general principle.¹³²

To take another example, explaining the existence of singleton sets in terms of the existence of their members seems suitably systematic: the existence of {Socrates} is determined by the existence of Socrates, {Obama} by Obama, and so on. We have a concise rule by which the outputted set can be defined in terms of the inputted entity. By contrast, consider the following haphazard ‘rule’: the existence of {Socrates} is determined by the existence of Cleopatra, {Obama} by Picasso, ‘and so on’. One problem with this connection would be its lack of systematicity: the outputted sets cannot be concisely defined in terms of the inputs.

¹³² It has a further problem with arbitrariness (discussed below). These problems are distinct: a theory can be systematic whilst providing arbitrary explanations e.g. one which explains all objects’ determinable properties in terms of their determinate properties together with the fact that Socrates exists.

These examples illustrate a key feature of systematic connections between fundamental and non-fundamental facts: they are characterized by way of ‘introduction rules’ for sub-factual constituents. In the case of determinables, the sub-factual constituent in question is the determinable property, and the introduction rule is that x instantiates the determinable P if x instantiates one of P ’s determinates. In the case of singleton sets, the sub-factual constituent in question is the singleton-forming operator (which takes an individual and forms its singleton set), and the introduction rule is that the singleton exists if the individual from which it is formed exists. It is hard to see how non-fundamental facts could be outputted systematically if not via connections which are ‘indexed’ by their sub-factual constituents in this way.

Permissive generalists regard individuals as non-fundamental constituents of facts of the form ‘ a is P ’ (more generally, ‘ a_1, \dots, a_n are R ’). Hence, in order to systematically explain these facts, they ought to provide introduction rules for individuals. These introduction rules will have to take the schematic form: ‘if ... P ..., then a is P ’. Since generalists only have qualitative facts out of which to construct individuals, ‘... P ...’ must stand for some purely qualitative description. How are permissive generalists to provide such rules?

The natural strategy is to appeal to a ‘qualitative essence’ $Q[a]$ which can serve as a surrogate for the individual a in the underlying qualitative facts: the a -involving facts will vary in accordance with the $Q[a]$ -involving facts. For example, this essence should plausibly satisfy the introduction rule: ‘if something is $Q[a]$, then a exists’. The natural extension to facts about a ’s nature is: ‘if something is $Q[a]$ and P , then a is P ’. (We can

think of these rules as governing the introduction of an ‘individual-forming operator’ i , which takes in some qualitative essence Q and forms the individual $i(Q)$, where $i(Q[a]) = a$.)

However, this strategy fails in worlds containing qualitatively indiscernible but distinct individuals. For a familiar illustration take TWINS, a world containing two duplicate fundamental particles, Castor and Pollux, floating one mile apart in an otherwise empty universe.¹³³ Let $Q[\text{Castor}]$ be Castor’s essence and $Q[\text{Pollux}]$ be Pollux’s essence. Given the proposed introduction rules, Pollux instantiates $Q[\text{Pollux}]$ (since something does). Since Castor and Pollux are qualitatively indiscernible, Castor also instantiates $Q[\text{Pollux}]$. But then it follows that something instantiates $Q[\text{Pollux}]$ and is identical to Castor, and hence, from our introduction rule for Pollux, that Pollux is identical to Castor! By similar reasoning, it follows that Pollux is floating one mile apart from itself. Thus, the proposed strategy fails to recover the non-qualitative properties of individuals, as well as the qualitative relations that they bear to one another.

When applied to symmetric worlds, this approach also faces a version of what Dasgupta (2014a:12) calls the ‘differentiation problem’. Castor and Pollux presumably have the same qualitative essence. (If not, what could explain the asymmetry?) Thus, their existences are each generated by the fact that something instantiates this essence. But it is a plausible principle that if two facts are distinct, it must be possible for them to be generated in

¹³³ C.f. Black 1952, Adams 1979. The fundamental description of TWINS takes the form:

$$\exists x \exists y (Px \ \& \ Py \ \& \ Rxy \ \& \ Ryx \ \& \ x \neq y \ \& \ \forall z (z = x \ \vee \ z = y))$$

different ways.¹³⁴ At least, the facts that Castor exists and that Pollux exists do not seem plausible counterexamples to this principle: the worldly difference between these non-basic individuals should somehow be reflected in their basis.¹³⁵ Unlike individualists, generalists cannot appeal to a brute non-qualitative distinction between them.

Excluding symmetric worlds like TWINS from the domain of the generalist's systematic explanatory principles would seem ad hoc: why shouldn't their account of the generation of individuals extend to such worlds? Granting that they belong to this domain, then, generalists seem forced to abandon the idea that individualist facts can be explained by matching individuals to qualitative essences in the manner suggested above. But how else are they to be systematically explained? For the explanation of individualist facts to be systematic, their constituent individuals must co-vary with some corresponding feature of the underlying generalist facts, as determinables co-vary with their determinates, and sets with their members. It's hard to see what this feature could be, if not some form of qualitative essence.

¹³⁴ Perhaps [p or q] and [p or r] are each actually generated by [p], but the first is possibly generated by [q] whereas the latter isn't. Perhaps [{Obama} exists] and [{Obama}] exists] are each necessarily generated by [Obama exists], but the first generates the second and does not generate itself.

¹³⁵ Perhaps the fusion of Obama with Trump and the set {Obama, Trump} are each necessarily generated by Obama and Trump's existences. But if so, this is because there are two different kinds of 'construction-operation' at work on the same raw material. The case of Castor and Pollux is not plausibly like this; presumably, there is only one kind of construction-operation at work in the generation of Castor and Pollux.

4.3.2 Arbitrariness

Metaphysical explanation should not be arbitrary. For example, suppose we want to explain the existence of composite objects in terms of their parts being ‘stuck together’. How tightly must they be stuck together? Any answer would seem implausibly arbitrary.¹³⁶ Or suppose we want to explain the existence of natural numbers in terms of pure sets. Should we explain them in terms of the von Neumann ordinals or the Zermelo ordinals? Any choice seems arbitrary: infinitely many sequences of sets implement the natural number structure equally well, with nothing privileging one over the others (Benacerraf 1965).

It is not obvious how this ‘arbitrariness’ constraint should be analyzed. It seems to concern the choice of a metaphysical basis failing to be suitably ‘privileged’. Leaving this at an intuitive level suffices for my purposes; as the cases illustrate, it seems clear both that there is such a constraint and that we have some grip on it.¹³⁷

When facing problems like those described above, we can either embrace extremity or indeterminacy. Take the mereological case. The ‘hard-liner’ embraces an extreme connection which revises our intuitions. For example, the claim that all collections of objects have a fusion seems attractively non-arbitrary, but revisionary in its commitment to gerrymandered mereological sums (e.g. Lewis 1986a:213). Conversely, the claim that no (or only very special) collections of objects have a fusion trades arbitrariness for a

¹³⁶ For arguments along these lines, see Lewis 1986a:212-3; van Inwagen 1990:126-7; Sider 2001:§4.9.

¹³⁷ For discussion of arbitrariness in metaphysics, see Korman 2010, Fairchild & Hawthorne 2018, Builes 2021.

revisionary denial of ordinary composites like tables (e.g. van Inwagen 1990:127). The ‘soft-liner’ embraces the idea that our notion of composition is indeterminate between many eligible candidate connections, thus avoiding the arbitrary choice between them.¹³⁸ Where these connections disagree on the outputted mereological truths, those truths are indeterminate. For example, when the cement between some bricks is drying, the point at which the wall exists is indeterminate.

An analogous situation arises for permissive generalists: to avoid arbitrariness in their connection between generalist and individualist facts, they must embrace extremity or indeterminacy.

Take the fact that Joe the electron exists. How is this generated from purely general facts? What must generally be the case for it to obtain? It is easy to see how answering such questions involves apparently arbitrary decisions. Consider all the qualitative situations vis-à-vis electrons, corresponding to the various ways of threading electron-trajectories through spacetime. How are we to label these trajectories with names? What makes it the case that one is occupied by Joe, and another by Anne?

Focus on worlds containing a single electron. Do any contain Joe? If so, which? Any choice of the trajectories ‘open to’ Joe seems objectionably arbitrary: for any cut-off, we should wonder why Joe’s existence depends on following a trajectory on one side rather than the other. But if none of these single-electron worlds contain Joe, then why not? How many

¹³⁸ Hirsch (1999) defends a view along these lines.

other electrons must there be before Joe's existence is generated? Any answer to this question would also seem arbitrary.¹³⁹

Extreme views appear to avoid arbitrariness, at the cost of rejecting our intuitions about how individuals vary across generally described situations. At one extreme, permissive generalists might adopt the plenitudinous view that there is an electron for every 'trajectory-choosing function', where a trajectory-choosing function is a function whose domain consists of some generalist worlds, and which maps each world in its domain to some electron-trajectory which is occupied at that world.¹⁴⁰ This avoids 'global arbitrariness' in the distribution of electrons across worlds: at each world, there are simply as many as there can be (given the constraint of no more than one trajectory per electron per world).¹⁴¹ Any given trajectory T at world w is occupied by very many electrons: one for every trajectory-choosing function which maps w to T. There is a clear cost associated with this plenitude. Even if non-basic individuals do not count against the view's ontological parsimony (as Bennett (2017: §8.2.2) and Schaffer (2015) argue), it is surely

¹³⁹ Compare Chisholm (1967:6) on the arbitrariness of assigning essential properties to Adam.

¹⁴⁰ Trajectories need to be understood here as qualitative properties, e.g. occupying a given trajectory is a matter of occupying a path with i) a certain shape, and ii) certain distance relations to other occupied paths. (In symmetric worlds, this will collapse distinct paths into a single 'trajectory', so trajectory-choosing functions will not settle relations of co-location. We might wonder: how are the facts about co-location non-arbitrarily settled at such worlds?)

¹⁴¹ This mirrors Bennett's (2004:355) suggestion that the modal difference between Statue and Lump is primitive and yet non-arbitrary since modal profiles are instantiated plenitudinously. For discussion of plenitude, see Fairchild (2019).

counter-intuitive to think that the world contains so many co-located electrons.¹⁴² (I discuss this view further in §4.4.2.)

At the other extreme, they might adopt the view that each electron only exists at a single world. Again, this seems to avoid global arbitrariness in that each individual is treated equally — they are all maximally fragile — and in this sense, their generation is principled.¹⁴³ But there is a clear cost: it is surely counterintuitive to think that all electrons depend for their existence on the most specific details of their qualitative situation, many of which seem to be irrelevant.¹⁴⁴ Intuitively, for example, that some faraway electrons follow the trajectories they do has nothing to do with Joe's existence!¹⁴⁵

If permissive generalists do not wish to endorse an extreme view of the kind I have described, an alternative way to avoid arbitrariness is to embrace indeterminacy. For example, they might say that, in single-electron worlds, the electron's identity is indeterminate: it is indeterminate whether such worlds contain Joe or Anne (although they determinately do not contain both). Embracing some such indeterminacy seems natural for

¹⁴² Another issue with this view is that it seems to unduly restrict which generalist worlds are possible: for example, any world featuring one electron features infinitely many co-located electrons!

¹⁴³ Although one might still wonder about 'local arbitrariness': why do the actual generalist facts generate *these* individuals rather than some others?

¹⁴⁴ This claim about dependence is distinct from Lewis's (1986a) view, in the context of his modal realism, that individuals are 'world-bound'. The latter concerns mereological overlap between spatiotemporally isolated concrete worlds, not how individuals' existence depends on qualitative facts.

¹⁴⁵ The intuition being violated here is not merely modal but concerns explanatory relevance directly. Hence, the issue would not be addressed by detaching modality in some way from the explanatory connections between generalist and individualist facts.

generalists, but it is uncomfortable for permissive generalists. Permissive generalists treat individuals like Joe and Anne as real (albeit non-basic) constituents of the world; their connections output worldly facts, not merely representational truths. Hence, if these outputs are indeterminate, the indeterminacy pertains to reality itself, not merely to our representation of it. This is a serious cost: it is at best bizarre—and at worst nonsensical—to think that reality features individuals but that it is indeterminate which.¹⁴⁶

The choice appears stark: if permissive generalists wish to avoid arbitrary connections between generalist and individualist facts, they must either embrace counterintuitively extreme connections or else worldly indeterminacy.

4.4. Strict Generalism

We have been considering how purely general facts could generate an individualist reality. But strict generalism denies that there are individuals in reality itself; instead, our talk of individuals non-perspicuously describes purely general facts. Hence, strict generalists must answer a different question: how can we use names (and other referential devices) to truly represent a general reality?

This ‘reduction’ approach targets individualist *truths*: ways of truly representing the world which are particular in form and which are ‘about’ individuals. There are different

¹⁴⁶ There have been several proposals for making sense of worldly indeterminacy (e.g. Barnes & Williams 2011, Wilson 2013) but none, I think, make its application to the case at hand especially attractive.

conceptions of truths, but for concreteness I focus on sentences. Hence, the task is to make it plausible that the truth of all true individualist sentences can be explained in purely general terms.¹⁴⁷

I conceive of this task in Sider's (2011:112) terms, as that of providing a 'metaphysical semantics' which shows 'how what we say fits into fundamental reality'. In particular, I aim to give generalist 'metaphysical truth-conditions' for individualist sentences, such that their truth may be explained by the obtaining of the corresponding metaphysical truth-conditions. My proposal is intended to serve as a 'proof by construction': I am more interested in illustrating the broad advantages of the reduction approach than in the contentious details of its implementation.

I help myself to the complete collection of generalist truths, including many non-fundamental truths (e.g. that something is tall). Thus, I am not providing a complete generalist explanation for individualist truths in fundamental terms, and it remains open that the complete explanation will involve the generation of some generalist facts by others. But the crucial question for generalists is how individualist truths may be explained in general terms, and I am showing how reduction is a better tool for this critical step.

¹⁴⁷ This is somewhat simplified. Actual languages are only capable of referring to a tiny portion of the individuals there are, whereas generalists should be obliged to show that they can (in principle) explain what is true about *all* the individuals there are. Moreover, their theory of how individualist truths are determined should apply in worlds containing no language. Thus, it may be more accurate to characterize their task more abstractly: they must show that they can explain the truths expressed by individualist sentences in all 'hypothetical' languages. This might be tantamount to explaining all true individualist *propositions*; I avoid this formulation since the nature of propositions is contentious.

4.4.1 The proposal

In order to give individualist sentences systematic truth-conditions, we ought to assign ‘metaphysical contents’ to names, which systematically contribute to metaphysical semantics, in the way that ordinary contents systematically contribute to ordinary semantics.¹⁴⁸

I propose that a name’s metaphysical content is its referent’s ‘qualitative profile’: the most specific qualitative property which its referent instantiates. Take some complete generalist description of a world in quantificational form. The (monadic) qualitative profiles witnessed at that world correspond to open formulae generated by removing one of the existential quantifiers in this description. For any individual a , let ‘ Q_a ’ denote a ’s qualitative profile, i.e. its most specific qualitative property.¹⁴⁹ The metaphysical content of a name n , which I denote ‘[[n]]’, will be Q_a , where n refers to a . For example, the metaphysical content of ‘Obama’, [[‘Obama’]], is Q_{Obama} , i.e. Obama’s qualitative profile.¹⁵⁰

¹⁴⁸ For simplicity I focus on the core case of names, ignoring other referring expressions such as pronouns, demonstratives, *etc.* The question of a name’s metaphysical content must be separated from debates within the philosophy of language concerning its ordinary content. For strict generalists, the metaphysical contents of names must be purely qualitative, but this does not entail descriptivism about their ordinary contents: the Millian claim that a name’s content is its referent can be accommodated as a (non-perspicuous) truth about ordinary semantics.

¹⁴⁹ Of course, individualist talk of names and their referents is needed to communicate a generalist theory of metaphysical truth-conditions. At the end of this section I ‘kick away the ladder’ by showing how the theory applies to itself.

¹⁵⁰ Ordinary indeterminacy concerning a name’s referent yields indeterminacy concerning its metaphysical content: if ‘Obama’ is indeterminate between various micro-aggregates, then [[‘Obama’]] will be correspondingly indeterminate between the qualitative profiles of these micro-aggregates.

I assume that for every qualitative property, the generalist language I am using to give metaphysical truth-conditions has some corresponding (perhaps complex) predicate, and I use the italicized ' Q_a ' to abbreviate the predicate corresponding to the qualitative profile Q_a . Thus, to say that something instantiates Q_{Obama} in our generalist language we may write: ' $\exists x Q_{Obama}(x)$ '.

Here then is a natural first pass for monadic qualitative truths — atomic truths about a single individual instantiating a qualitative property:

‘a is P’ is true iff $\exists x (Px \ \& \ Q_ax)$.

Thus, for example, ‘Obama is tall’ is true just in case something is Q_{Obama} and tall.

This first pass is extensionally adequate in the sense that the biconditionals it entails are true, and moreover, the truth-condition on the right varies systematically with the sentence on the left.¹⁵¹ But it does not seem appropriately explanatory: the truth-conditions it provides include much irrelevant detail. For example, Obama’s tallness is explained in terms of the complete details of Obama’s qualitative profile, most of which are surely irrelevant. Indeed, this profile is only instantiated at the actual world, whereas Obama remains tall at many nearby worlds.¹⁵²

¹⁵¹ Left-to-right: if ‘a is P’ is true, then a is P and Q_a (since a is Q_a), and hence something is P and Q_a . Right-to-left: if something is P and Q_a (where P is qualitative), then anything which is Q_a is also P, and hence a is P.

¹⁵² In modal terms, we might require that metaphysical truth-conditions entail *necessary* biconditionals.

To fix this, I propose that the metaphysical truth-conditions for individualist truths be specified using a counterpart relation, much like in Lewis's (1983b) account of *de re* modal truths. Lewis proposed that actual individuals stand in a counterpart relation to relevantly similar possible individuals. Similarly, I propose that actually instantiated qualitative profiles stand in a counterpart relation to relevantly similar possibly instantiated profiles. An individual's qualitative profile has as its counterparts all those qualitative profiles which, intuitively, it would have instantiated had the world been different in some qualitative respect.

This counterpart relation should obey three constraints: (i) it relates a profile Q_a to profiles similar enough that a could instantiate them; ii) if it relates Q_a to Q_b , then Q_b is more similar to Q_a than any other profile instantiated in Q_b 's world, and iii) if it relates Q_a to Q_b , then Q_a is more similar to Q_b than any other profile instantiated in Q_a 's world.¹⁵³

As with Lewis's counterpart relation, which of many candidate relations best satisfies these constraints is indeterminate and context-dependent. The first constraint should be applied in accordance with our *de re* modal intuitions: for example, we think that Obama could have been one inch shorter, and could have survived various extrinsic changes, but are less inclined to believe that he could have been a poached egg, or made of silicon.¹⁵⁴ The second and third constraints should be applied in accordance with our intuitions about relative

¹⁵³ Unlike with Lewis's counterpart relation (1983b:29), this third constraint ensures that no two actual qualitative profiles share any counterparts.

¹⁵⁴ We are not, however, in the business of fixing modal truths here; it is an open question how (if at all) the theory of *de re* possibility connects to the counterpart relation in question here.

similarity. For example, we think that Q_{Obama} is more similar to a situation just like it except for the positions of some electrons than it is to a likewise modified version of Q_{Trump} . Intuitively, the counterparts of Q_{Obama} are *distinctively Obama-ish*: they are similar enough to this profile that Obama could have instantiated them, more similar to it than any other profiles witnessed at their world are, and more similar to it than to any other actual profile.

Let the ‘expansion’ of a qualitative profile Q , Q^R , be the qualitative property generated by disjoining all those profiles related to Q by the counterpart relation R .¹⁵⁵ (Thus, $Q^R x$ iff $\exists F Fx \ \& \ R(Q, F)$.) We can then give metaphysical truth-conditions for monadic truths as follows:

‘ a is P ’ is true iff $\exists x (Px \ \& \ Q_a^R x)$.

This evades the problem with the first pass: a qualitative profile’s expansion is far less specific than the profile itself, and accordingly the fact that something instantiates it omits details which are irrelevant to the individualist truth in question. For example, ‘Obama is tall’ is plausibly only true in situations where something instantiates some ‘distinctively Obama-ish’ qualitative profile.

There is a new problem, however. World w features two alternative versions of Earth. The most similar individual to Trump Sr at w inhabits Earth 1, and the most similar individual to Trump Jr at w inhabits Earth 2. It follows by the proposal above that ‘Trump Sr inhabits

¹⁵⁵ This parallels Lewis’s (1983b:35) notion of a thing’s essence.

Earth 1' and 'Trump Jr inhabits Earth 2' are each true at w . But intuitively, this is incorrect: w hosts no interplanetary relatives.

This calls for holistic truth-conditions. First, for any individuals a_1, \dots, a_k , let $Q_{\langle a_1, \dots, a_k \rangle}$ be their 'joint qualitative profile' — the most specific k -place qualitative relation that they instantiate — and let this be the metaphysical content of a corresponding sequence of names $\langle n_1, \dots, n_k \rangle$.¹⁵⁶ Second, extend the counterpart relation to these joint profiles, so that it relates them to relevantly similar joint profiles in a manner which obeys constraints parallel to those set out for monadic profiles.¹⁵⁷ To illustrate, the metaphysical content of $\langle \text{'Trump Sr'}, \text{'Trump Jr'} \rangle$ is Trump Sr and Trump Jr's joint profile (the maximally specific binary qualitative relation that Trump Sr and Trump Jr instantiate). Since this relation is 'paternal' — x 's bearing it to y involves x being the father of y — it should only have similarly paternal relations as its counterparts (insofar as the father-of relation intuitively holds with necessity).

With this generalized counterpart relation in place, we can give generalized truth-conditions for atomic individualist sentences as follows:

¹⁵⁶ In symmetric worlds, this joint profile is not determined by the monadic qualitative profiles in question. Hence, one might worry that metaphysical contents fail to be 'compositional': the content of a sequence of names is not determined by their contents in isolation. But since the metaphysical content of a sequence of names is defined as the joint profile of their referents, it is determined compositionally by their ordinary contents (assuming these contents fix reference). This seems to be compositionality enough.

¹⁵⁷ Lewis (1983b:44) observed an analogous issue for de re modal truths, and proposed an analogous solution: determine joint possibilities by assigning counterparts to n -tuples of individuals.

‘ $Pa_1 \dots a_n$ ’ is true iff $\exists x_1 \dots \exists x_n Px_1 \dots x_n \ \& \ Q_{\langle a_1, \dots, a_n \rangle}^R x_1 \dots x_n$

Thus, for example:

‘Trump Sr is Trump Jr’s father’ is true iff

$\exists x \exists y x \text{ is } y\text{'s father} \ \& \ Q_{\langle \text{Trump Sr}, \text{Trump Jr} \rangle}^R xy$

i.e. some distinctively $\langle \text{Trump Sr}, \text{Trump Jr} \rangle$ -ish things instantiate the son-of relation.

Hence, we may suppose that Trump Sr is Trump Jr’s father (and that they inhabit the same planet) at w , since no interplanetary pair is $\langle \text{Trump Sr}, \text{Trump Jr} \rangle$ -ish.¹⁵⁸

My proposed metaphysical semantics for individualist truths presupposes the truth of certain linguistic claims which are themselves individualist and so must be given a generalist reduction. To illustrate, let S be the following sentence-token: ‘Obama is tall.’. Let N be the name-token which S features. The proposal above provides the following generalist truth-conditions:

¹⁵⁸ To ensure consistency between the truth-conditions for various individualist truths, the counterparts for profiles of various adicities must mesh appropriately. For example, if the counterpart at w of $Q_{\langle \text{Trump Sr}, \text{Trump Jr} \rangle}$ is instantiated by some pair of Earth 2-inhabitants, then the counterpart at w of $Q_{\text{Trump Sr}}$ should be instantiated by the first of this pair. Otherwise, the proposed truth-conditions might make ‘Trump Sr and Trump Jr each inhabit Earth 2’ true but ‘Trump Sr inhabits Earth 2’ false. The required constraint is as follows:

At any world w , $\forall x_1 \dots \forall x_n (Q_{a_1}^R x_1 \ \& \ \dots \ \& \ Q_{a_n}^R x_n \leftrightarrow Q_{\langle a_1, \dots, a_n \rangle}^R x_1 \dots x_n)$.

An admissible counterpart relation trades off the constraints applying to profiles of different adicities.

‘S is true’ is true iff $\exists x (x \text{ is true} \ \& \ Q_S^R x)$.

‘N refers to Obama’ is true iff $\exists x \exists y (x \text{ refers to } y \ \& \ Q_{\langle N, Obama \rangle}^R xy)$.

Of course, these semantic clauses themselves presuppose further linguistic truths about names and sentences, but these may be similarly analyzed.¹⁵⁹

4.4.2 Progress

I am now in a position to show how the proposal above makes progress, by addressing the problems with systematicity and arbitrariness that I raised for permissive generalism.

Firstly, recall that permissive generalists cannot outfit individuals with qualitative essences in a way that would make the generation of individualist facts appropriately systematic, since this strategy fails in symmetric worlds like TWINS. The proposal above is systematic: it provides a concise rule by which individualist truths vary systematically across generalist situations in accordance with the metaphysical contents of the names used to expressed them. But it has no analogous difficulty with symmetric worlds such as TWINS.

¹⁵⁹ The ensuing regress is non-vicious because the linguistic truths presupposed in stating the truth-condition for a sentence are not themselves involved in *reducing* the truth that it expresses. Compare: that ‘Obama’ refers to Obama is relevant to the truth that ‘Obama is tall’ means that Obama is tall, but not to the truth that Obama is tall.

Suppose we live in a symmetric world, and introduce the name ‘Obama*’ for the qualitative duplicate of Obama on the other side of the universe. Then the truth that Obama is distinct from Obama* reduces to a general truth of the form:

$$\exists x \exists y (x \neq y \ \& \ Q_{\langle \text{Obama}, \text{Obama}^* \rangle}^R xy)$$

i.e. there are some distinct <Obama, Obama*>-ish things. Similarly, the truth that Obama is far away from Obama* but not from himself reduces to a general truth of the form:

$$\exists x \exists y (Dxy \ \& \ \sim Dxx \ \& \ Q_{\langle \text{Obama}, \text{Obama}^* \rangle}^R xy)$$

i.e. there are some <Obama, Obama*>-ish things such that the first is far away from the second but not from itself.

The key to this resolution of the systematicity problem is the holistic nature of the metaphysical semantics described above: the collective semantic contribution made by some names cannot always be read off their individual contributions. The names ‘Obama’ and ‘Obama*’ have the same metaphysical content, and yet the pair <‘Obama’, ‘Obama*’> has a different content from the pair <‘Obama’, ‘Obama’>.

Now, the generation approach may be made similarly holistic: instead of assigning joint contents to collections of names, we could assign ‘joint essences’ to collections of individuals, where these joint essences are not determined by individual essences (cf. Fine

1994:65). Thus, for example, it is essential to $\langle \text{Castor, Pollux} \rangle$ that they are distinct, even though Castor is not essentially distinct from Pollux, nor Pollux essentially distinct from Castor.¹⁶⁰ Castor and Pollux's distinctness may then be generated from the generalist fact that there are some distinct things instantiating Castor and Pollux's joint essence.

I regard this parallel solution as significantly more costly than the proposal I have outlined. In order to give systematic explanations, both sides need some sort of 'defining properties' for both individuals and collections of individuals, which determine how and when they appear in various generalist situations. But for permissive generalists these must be real definitions, i.e. worldly essences, whereas strict generalists merely need metaphysical contents together with the counterpart relation (which might be thought to fix a kind of 'linguistic essence'). Positing joint worldly essences is to go out on a metaphysical limb in a way that is not required by the idea that names used in symmetric worlds may acquire holistic contents. Indeed, it is unclear how the facts about worldly essences are themselves to be systematically generated. By contrast, the reduction approach can explain how names get their metaphysical contents. For example, we can reduce the truths that (i) 'Obama' refers to Obama, (ii) 'Obama*' refers to Obama*, and (iii) Obama and Obama* instantiate $Q_{\langle \text{Obama}, \text{Obama}^* \rangle}$, and use these to explain why $Q_{\langle \text{Obama}, \text{Obama}^* \rangle}$ is the metaphysical content of $\langle \text{'Obama'}, \text{'Obama}^* \rangle$.

¹⁶⁰ Alternatively, we might suppose that Castor and Pollux's essences each mention the other. But reciprocal essences like these appear to yield unattractive explanatory circles: how Castor is would be partly explained by reference to Pollux, and vice versa (Fine 2015:297).

The reduction approach has a further advantage when it comes to the ‘differentiation problem’. Recall that Castor and Pollux plausibly share a worldly essence, and so are both generated by the fact that something instantiates this essence. This violates the plausible general principle that different facts ought to at least be possibly generated in different ways. Now, the reduction approach faces an analogous phenomenon of truths involving names used within symmetric worlds —such as ‘Obama exists’ and ‘Obama* exists’— reducing to the same general truth. However, the analogous principle that different truths reduce (or at least, possibly reduce) in different ways is not nearly as plausible: it is a familiar idea that sentences involving different names may differ representationally without differing in their worldly basis.¹⁶¹

Secondly, recall that permissive generalists face a stark choice: if they wish to avoid arbitrary connections between generalist and individualist facts, they must either embrace counterintuitively extreme connections or else worldly indeterminacy. The proposal above embraces indeterminacy: it avoids arbitrary cut-off points by denying that there are any sharp cut-offs. In particular, in any given context, it is indeterminate which relation plays the role of the counterpart relation: many admissible candidates each satisfy the relevant constraints equally well.

I favor a supervaluationist resolution: an individualist sentence is true iff it is true relative to all admissible counterpart relations, and indeterminate iff it is true relative to some and

¹⁶¹ Though see §4.5.2 below for a related concern: that the reduction approach eliminates singular thought.

false relative to others. Thus, many individualist sentences have indeterminate truth-values at many generalist worlds.¹⁶² For example, suppose our conception is that any particular electron could just as well exist alone as any other. Then, for any qualitative profile corresponding to a lonely electron, it will be indeterminate which of the profiles instantiated by actual electrons are its counterpart: on some admissible counterpart relations, Joe's profile is its counterpart, and on others, Anne's is. Thus, at any single-electron world, it is indeterminate whether any given actual electron exists (although it is determinate that at most one actual electron exists, assuming that the counterpart relation reflects the necessity of distinctness).

This result is natural from a generalist perspective. Where generalists recognize a single qualitative possibility, our individualist conception envisages a world where Joe exists alone and a distinct world where Anne exists alone. It is natural for generalists to deny that there is any determinate mapping from their range of fundamental qualitative possibilities into the second-rate pluriverse of individuals. On the reduction approach, this indeterminacy is linguistic. It is indeterminate what makes it true that Joe exists, since the mechanism by which our use of names extends to hypothetical scenarios doesn't nail down what it takes to be distinctively Joe-ish. Linguistic indeterminacy is familiar: for example, it is indeterminate which of many precisely-bounded Everest-candidates makes it true that Mount Everest exists, since the content of 'Mount Everest' fails to specify precise boundaries. The parallel move on the generation approach would be to suppose that Joe

¹⁶² Since the counterpart relation is determinately reflexive, this proposal does not make for any (extra) indeterminacy when it comes to the truth of individualist sentences at the actual world.

has an indeterminate essence, and that, at some generalist worlds, an indeterminate fact regarding Joe's existence is generated. Since facts are worldly, this is worldly indeterminacy, analogous to the odd idea that some mountain with imprecise boundaries makes it true that Mount Everest exists.

As discussed above, an alternative solution for the generation approach secures a mapping from generalist worlds to individuals which is both principled and determinate by embracing plenitude: roughly, for every consistent disjunction of complete qualitative profiles, there is an individual whose existence is generated by the instantiation of the corresponding property.¹⁶³ This proposal seems unattractive in light of the reduction approach, which ties referring expressions directly to qualitative aspects of reality, without detouring through corresponding individuals. Overpopulating the world with an unexpected plenitude of unfamiliar individuals is simply unnecessary: the posit fails to earn its keep, since we can explain individualist truths in a principled way without it. (In a sense, the reduction approach is also plenitudinous: not with respect to individuals themselves – of which, in reality, there are none! — but with respect to the corresponding consistent disjunctions of complete qualitative profiles.)¹⁶⁴

Compare: the mereological universalist posits a composite object for every collection of simples, using the resulting plenitude to provide semantic values for names. If names could

¹⁶³ Presumably it would be indeterminate which of these many properties actual names correspond to, and hence, indeterminate which of many co-located qualitative duplicates they refer to.

¹⁶⁴ The selection of these profiles by actual names is constrained both by the actual qualitative profiles of the individuals which the linguistic community can refer to and by the counterpart relations which the community can adopt.

instead be tied directly to the collections of simples themselves, then composite objects are redundant, and the mereological nihilist has a compelling parsimony argument against them. Analogously, I am saying that plenitudinous individuals *are* redundant for the generalist: instead of vindicating our individualist talk by generating individuals from general facts, they can give at least as adequate a theory directly in terms of the general facts themselves.

4.5. Remaining Issues

4.5.1 Proportionality

Dasgupta (2014a: §3) argues that generalists cannot provide proportional grounds for individualist facts: that is, grounds which are both specific enough to necessitate these facts and yet not so specific as to include intuitively irrelevant detail. Consider Obama's instantiating some intuitively local (and accidental) property:

(O) Obama is smiling (at time *t*).

Any qualitative fact intrinsic to the solar system can obtain in almost-symmetric worlds where, intuitively, it is an Obama-impersonator rather than Obama who is smiling.¹⁶⁵ Indeed, since we can repeat this argument for any region, the only qualitative facts which could necessitate (O) must be global facts about the entire universe. But such facts are

¹⁶⁵ Such worlds are reached by a two-step modification of the actual world: first, add a faraway duplicate solar system containing a smiling Obama-impersonator; second, modify the actual solar system so that Obama is not smiling.

intuitively irrelevant: surely what is going on around Alpha Centauri, for example, has nothing to do with Obama's smiling!

How does this issue affect the choice between permissive and strict generalism? I believe that, if anything, it provides further reason to favor strict generalism.

One reaction to this issue — proposed by Dasgupta (2014a) himself — is to retreat to a 'structuralist' implementation of generalism.¹⁶⁶ On this view, individualist facts have no explanation in isolation, but may only be explained as parts of a plurality of individualist facts. Whilst a global generalist basis may not be wholly relevant to any particular fact in this plurality, all of its details are relevant in some way to the entire plurality. Since this response abandons the ambition of explaining facts like (O), I view it as a last resort. However, it is worth noting that it is equally available to both permissive and strict generalists: just as permissive generalists may embrace a many-many relation of generation between facts, strict generalists may embrace a many-many relation of reduction between truths.

Moreover, adopting structuralism does not affect the issues with systematicity and arbitrariness discussed above. A structuralist implementation of permissive generalism will struggle to provide connections between generalist facts and pluralities of individualist facts which are systematic and non-arbitrary. It is hard to see how these holistic connections could be systematic in the absence of any systematic 'sub-connections' between generalist

¹⁶⁶ Murphy (ms) defends a similar view.

facts and particular individualist facts. And the shift to targeting pluralities of individualist facts does nothing to avoid the need for seemingly arbitrary decisions concerning the general situations in which a given individual exists. For the reasons discussed above, a structuralist implementation of strict generalism would avoid these problems.

A second — and to my mind, more promising — reaction denies our intuitions about relevance: the way things are globally may be wholly relevant to how they are with respect to a particular individual.¹⁶⁷ Here, I think that strict generalists may have an advantage: only they can give a plausible error theory for the intuition that individuals should have a local basis. In particular, this locality intuition may reflect the implicit assumption that individualist truths are perspicuous. If there really are worldly individuals, as permissive generalists believe, it seems bizarre for them to be non-locally generated: we expect worldly individuals to be ‘anchored’ to their own patch of reality. If, however, truths like (O) turn out not to involve individuals at all in their worldly basis, then it is less surprising that this basis turns out to be global. The non-locality reflects the way in which names latch onto an individual-free reality: by contrast with the local matter of simply being Obama, being distinctively Obama-ish is a global competition amongst the qualitative profiles instantiated at a world.¹⁶⁸

¹⁶⁷ A third reaction denies that the qualitative grounds for individualist facts must be necessitating (perhaps the only necessitate in conjunction with holistic ‘background conditions’). See Skiles 2015 for defense of the general idea that grounds need not necessitate the grounded. *Prima facie*, this response is equally available to strict and permissive generalists.

¹⁶⁸ Compare the case of simultaneity. Intuitively, whether e1 and e2 are simultaneous ought to be ‘internal’ to these events. But once it is revealed that truths of the form ‘e1 and e2 are simultaneous’ are non-perspicuous – since they hide relativity to a contextually specified reference frame – the intuition loses its force.

4.5.2 Singular thought

It might be thought that permissive generalism has an advantage over strict generalism when it comes to the phenomenon of singular thought, since it provides individuals in reality for such thought to latch onto. To make this idea vivid, suppose we are living in a symmetric world. On Earth, Trump is thinking that Obama is happy; on the faraway duplicate planet Earth*, Trump* is thinking that Obama* is happy. Trump and Trump* are clearly thinking thoughts with different contents. For example, suppose that Trump is brought unknowingly to Earth* overnight. When he meets Obama*, he is not justified in thinking that this man is happy, even if he is justified in thinking that Obama is happy.¹⁶⁹

Since Obama and Obama* are qualitative duplicates, how are strict generalists to make sense of this purely *de re* difference? On the proposal above, Trump and Trump*'s thoughts would have the same metaphysical truth-condition: that something is both Q_{Obama} and happy. By contrast, permissive generalists can view the two thoughts as latching onto distinct facts involving distinct individuals (though at the cost of making it mysterious how these distinct facts are generated).

In my view, strict generalists should distinguish the ordinary contents of our thoughts, which may be singular, from their metaphysical contents, which cannot be. For example, they can reduce the non-perspicuous truth that Trump is thinking about Obama whereas Trump* is thinking about Obama* to a generalist truth of the following form:

¹⁶⁹ See Turner (forthcoming: §8) for a version of this concern.

$$\exists x \exists x^* \exists y \exists y^* (T_{xy} \& T_{x^*y^*} \& Q_{\langle Trump, Trump^*, Obama, Obama^* \rangle} R_{xx^*yy^*}).$$

Moreover, they can account for the epistemological significance of ordinary contents. For example, they can reduce the truth that Trump is justified in thinking that Obama is happy but not justified in thinking that Obama* is happy as follows:

$$\exists x \exists y \exists y^* (J_{xy} \& \sim J_{xy^*} \& Q_{\langle Trump, Obama, Obama^* \rangle} R_{xyy^*}).$$

It is hard to shake — but perhaps harder to precisely articulate — the feeling that something has gone missing here. We can put it metaphorically by imagining a ‘Godseye’ perspective on the symmetric world. From this lofty position, we might imagine pointing to one of the two Obama-duplicates and asking: is *that* Obama, or merely his duplicate?

But the strict generalist should deny that this question makes sense. First, there are no particular individuals in the world for the demonstrative ‘*that*’ to identify: all we can ‘point at’ are instantiated qualitative profiles. Asking this question is like selecting a variable in the world’s quantificational description, and asking whether ‘*it*’ corresponds to Obama or his duplicate.

Second, it is crucial to Trump’s thinking about Obama rather than Obama* that his thought is itself embedded within the world, and so may bear different relations to Obama and Obama*. Once a thought is ‘brought outside’ the world, there is nothing to determine which of the qualitative duplicates it picks out. What remains, from the Godseye perspective, are

certain ‘penumbral connections’: for example, that Obama and Trump share a planet, and that Obama is denoted by the concept OBAMA. Thus, I suspect that the elusive dissatisfaction with the strict generalist’s description of reality has its source in a broader concern, arising for individualist and generalist alike, that any objective description misses the indexical truths about ‘who is where’.¹⁷⁰

4.6. Concluding Remarks

I have argued that the generation approach to generalism faces serious difficulties, and that these difficulties can be overcome by switching to the reduction approach. In conclusion, I want to return to the initial appeal of generalism. As I will argue, strict generalism is not only a more successful version of the view — it is also more attractive in the first place.

First, recall Dasgupta’s (2009) argument that individuals are physically redundant and empirically undetectable. Permissive generalism avoids physically redundant individualist facts: worlds differing on their individualist facts at some time must also differ on their general facts at that time, and hence on their general facts at later times. But compare the case of absolute velocities (after which Dasgupta’s argument is modeled). We could avoid physically redundant absolute velocities in a parallel way, by generating them from relative velocity facts. It seems clear that preserving non-basic absolute velocities in this way would be unmotivated and, in some sense, unparsimonious. The more natural response is to

¹⁷⁰ I leave the question of whether (strict) generalists face any special version of this problem to future work.

eliminate absolute velocity structure altogether (whilst allowing for non-perspicuous truths about absolute velocity, with respect to contextually specified reference frames). If Dasgupta's argument succeeds, I see no reason why our attitude towards individuals should be any different: positing individualist facts, whether basic or not, seems unmotivated.¹⁷¹ (Relatedly, strict generalism provides a much clearer vindication of the intuition that permuting individuals makes no 'real' difference than permissive generalism does; cf. Sider 2020:8.)

Secondly, recall the idea that our individualist conception is explained by its usefulness: agents inevitably construct an 'ersatz world' of individuals in order to render a holistic general reality representationally tractable. The vision this motivates is *strict* generalism: an explanation along such lines would make it understandable that we represent the world in a radically non-perspicuous way. If anything, it only makes permissive generalism *more* mysterious! After all, the ersatz world would appear to do its work irrespective of any derivative individualist facts: given that individualist language aptly represents the general facts, perspicuously representing some derivative reality seems to provide no additional expediency.¹⁷² From the generalist perspective, then, there seems little reason to expect reality to indulge our individualist conception this way — and, as I have argued, plenty of reason not to

¹⁷¹ A more detailed discussion of physical redundancy and fundamentality must be left for future work.

¹⁷² There is plenty more to say about the pragmatic advantages of representing perspicuously. The claim here is not that perspicuity is pragmatically irrelevant, but only that it is hard to see why an individualist conception should be more adaptive in a permissive generalist reality than in a strict generalist reality.

5. Against Grounding Physicalism¹⁷³

5.1 Introduction

According to physicalism, there is nothing ‘over and above’ the physical. The enormous success of physics made physicalism the dominant view in the second half of the twentieth century. Discussion in this time typically presupposed that physicalism involves some way of identifying mental phenomena with broadly physical phenomena, with the focus on questions such as: is there ‘type identity’ or merely ‘token identity’ (Putnam 1967, Davidson 1970, Fodor 1974)? Are phenomenal properties to be identified with functional properties or with their realizers (Block 1978, Lewis 1980, Kim 1998)? Are the identities necessary or contingent (Place 1956, Kripke 1980), a priori or a posteriori (Lewis 1994b, Block & Stalnaker 1999, Chalmers & Jackson 2001)?

However, many who are otherwise attracted to the elegance of the physicalist world-view are unsatisfied with identity physicalism, especially as it applies to consciousness. A common kind of complaint is simply that the view is committed to ‘bare and incredible identity assertions’ (Schaffer 2021: 203): physical properties and phenomenal properties are intuitively very different (and their identification has too often been justified by implausible conceptual analyses). Moreover, the idea that consciousness reduces in some sense to physics seems to amount to a radical eliminativism about something whose reality

¹⁷³ For discussion of previous versions of this paper, thanks to Karen Bennett, Ned Block, David Builes, Dave Chalmers, Christopher Frugé, Luke Roelofs, Ted Sider, the Fall 2022 Rutgers Metaphysics Group, the Fall 2022 NYU Philosophy of Mind Group, and especially Verónica Gómez and Jonathan Schaffer.

is supremely vivid and undeniable: our own subjective perspective on the world. As Nagel (1974:447) puts it, the identification of physical and phenomenal properties is surrounded by ‘an air of mysticism’. Levine (2006:148) goes further: ‘The idea seems unintelligible, in a way that the identity of water with H₂O doesn't.’

A less direct concern is that identifying consciousness with some broadly physical property renders its significance mysterious (Pautz 2017). When I am conscious of some object in my environment, that relation is very different to the one that I bear to other objects in my environment, in a way that underlies both my special ability and my special reason to have certain beliefs about the object in question. But how could this be if the conscious-of relation is some arbitrary-looking broadly physical relation (defined in terms of causation or biological function, say), with many intrinsically similar relations in its vicinity?¹⁷⁴

Another problem is more general: it applies to any high-level natural kinds, such as the property of being a carbohydrate, or being a tiger. It takes the form of a trilemma. Identifying high-level natural kinds with unified ‘first-order’ realizing properties gets their actual or possible extension wrong, by ignoring multiple realizability (Putnam 1967). Identifying these properties with potentially infinite disjunctions of their actual and/or possible realizing properties makes them disunified, rendering them unsuited to their nomic/causal role and broader explanatory significance (Fodor 1974). Finally, identifying them with functional properties unifies them but threatens both to strip them of their causal

¹⁷⁴ Lee (2018) embraces the idea that, since it is just some arbitrary-looking physical property, consciousness has no special significance.

efficacy (Kim 1998) and to return us to the dilemma between extensional inadequacy and disunification when we ask what functional properties are themselves to be identified with (Schaffer 2013).¹⁷⁵

For these reasons, many have sought a more moderate version of physicalism, which reconciles the unity, significance, causal efficacy, and apparent irreducibility of phenomenal properties with their non-fundamentality, thus avoiding the extremes of identity physicalism on the one hand and dualism (or idealism) on the other.¹⁷⁶ In line with a wider trend in contemporary metaphysics, *grounding physicalism* has recently emerged as the leading contender to occupy this position. According to grounding physicalism, no physical phenomenon is itself identical to consciousness, but physics nonetheless grounds (generates, gives rise to, or ‘metaphysically causes’) conscious experiences. This view attributes non-physical levels to reality itself (rather than our representation of it) and so there is a sense in which it is less austere than identity physicalism (and a sense in which, since these levels are merely derivative, it isn’t).¹⁷⁷

¹⁷⁵ The point here is that functional definitions will mention stimuli, behavior, and other psychological states, connected by relations such as causation, counterfactual dependence or high-level laws. Since none of these notions are plausibly fundamental, functional identifications will not vindicate a ‘pure’ identity physicalist view, which seeks to provide identifications in fundamental terms.

¹⁷⁶ For discussion of these and some other challenges for identity physicalism, see Pautz forthcoming: §2.3.

¹⁷⁷ For discussion of grounding physicalism, see Dasgupta 2014b and Schaffer 2021. I use ‘ground’ here to cover any relation which constitutes or backs metaphysical explanation and where there is a worldly (as opposed to merely representational) difference between explanans and explanandum, including e.g. determination (Yablo 1992), constitution (Pereboom 2011), and realization (Shoemaker 2007, Wilson 2011). If the ‘fundamental’ is understood as grounding all else, then fundamentality physicalism (‘only the physical is fundamental’) is a version of grounding physicalism (Rabin 2019).

Like identity physicalists, grounding physicalists hold that consciousness is ‘nothing over and above’ physics, but they view the connection as a matter of grounding rather than identity. For example, grounding physicalists would deny the following identification:

Joe’s being in pain at 2pm *just is* Joe’s C-fibers firing at 2pm.

But they might replace it by the following grounding claim:

Joe’s being in pain at 2pm *is grounded in* Joe’s C-fibers firing at 2pm.

On this approach, Joe’s pain (a mental event) is distinct from and yet intimately connected to — not ‘metaphysically distinct’ from — his C-fibers firing (a physical event).

This paper argues against grounding physicalism. In a nutshell, the argument is this: if the broadly physical realm is connected to the phenomenal realm by ground, then there ought to be some compact system of physical-phenomenal laws. But no such system seems to be available. Hence, if consciousness is nothing over and above the physical, the gap is bridged by identity rather than ground.

For recent comparison of grounding physicalism and identity physicalism (with a different perspective from my own), see Pautz forthcoming. The division approximates the more traditional distinction between ‘reductive’ and ‘non-reductive’ physicalism, though the latter is sometimes drawn in terms of the availability of a priori entailments, which (as I see it) is orthogonal to the former. For an overview of the various formulations and varieties of physicalism, see Stoljar 2021.

5.2 Grounding Physicalism

Let me begin by formulating grounding physicalism (as it applies to consciousness) a little more precisely. As I see it, the view has a positive component and a negative component. The positive component is that the phenomenal is grounded in the physical (this is what makes the view physicalist, as opposed to dualist or idealist). The negative component is that the phenomenal is not itself physical (this is what makes it grounding physicalism rather than identity physicalism).

To flesh out these components in a neutral way, I will speak in terms of relations between sentences.¹⁷⁸ First, following Kit Fine (2012: §1.4), let ‘<’ (pronounced ‘(fully) grounds’) relate some sentences on the left and a sentence on the right, as in:¹⁷⁹

‘Snow is white’, ‘Grass is green’ < ‘Snow is white and grass is green’.

This can be understood as saying that the truths expressed by the sentences on the left make it the case that the truth expressed by the sentence on the right obtains, or that the truth expressed on the right holds in virtue of the truths expressed on the left.

¹⁷⁸ This ‘meta-linguistic’ approach allows for simple formulations without use/mention confusions and without assuming a domain of fine-grained propositions. Of course, it has the somewhat artificial expressive limitation of being tied to language, but this won’t matter for our purposes.

¹⁷⁹ Strictly speaking, Fine uses this symbol as the sentential operator, and a triangle as the corresponding predicate.

Second, following Cian Dorr (2016: §3), ‘ \equiv ’ (pronounced ‘for ... to be true just is for ... to be true’) relates two sentences, as in:¹⁸⁰

‘Grass is green and snow is white’ \equiv ‘Snow is white and grass is green’.

This ‘identification’ can be understood as saying that there is no worldly difference between what the sentence on the left expresses and what the sentence on the right expresses: they differ merely representationally. In ‘facts’ language, they latch onto the same fact, or are ‘factually equivalent’.

The two components of grounding physicalism about consciousness can then be articulated as follows:

- i) For any phenomenal truth q , there are some physical truths pp such that $pp < q$.
- ii) For no phenomenal truth q is there a physical truth p such that $p \equiv q$.

A physical truth can be understood as a true sentence whose constituent terms are all (broadly) physical: paradigm cases are ‘All electrons are negatively charged’ and ‘Joe’s C-fibers are firing at 2pm’. Since my interest is in the case of consciousness, I will be expansive in my conception of which terms count as physical, to include, roughly, anything

¹⁸⁰ Again, Dorr uses this symbol as a sentential operator rather than a first-order predicate.

objective and non-phenomenal (including the vocabulary of higher-level sciences like chemistry and biology, logical/mathematical vocabulary, nomic/modal vocabulary, *etc.*). It is likely that grounding physicalists will endorse their view with respect to a narrower conception of the physical. But set aside any grounding within this broadly physical realm: the question is whether grounding connects this realm to the phenomenal.

Let a 'phenomenal truth' be a true sentence which directly concerns the distribution of conscious experiences across subjects and times by involving some phenomenal predicate. Paradigm examples are 'Joe is in pain at 2pm' and 'Everyone is having a reddish experience'. I assume that physical and phenomenal truths are delineated in such a way that there are no trivial logical entailments between them. Thus, 'It is not the case that Joe is in pain and not in pain' is not a phenomenal truth, and 'Either Joe's C-fibers are firing or he is in pain' is neither physical nor phenomenal.

Thus conceived, grounding physicalists (even those with a relatively coarse-grained outlook on identification in general) deny that there are any identifications between physical truths and phenomenal truths. In particular, they deny any identifications which derive from underlying 'analyses' of phenomenal properties in (broadly) physical terms, such as:

To be in pain just is to have firing C-fibers.

To have a reddish experience just is to be in a state which is caused by something's being red.

I will argue that the positive component of grounding physicalism conflicts with this negative component: if the phenomenal is grounded in the physical, then the phenomenal must itself be physical.

5.3 Against Grounding Physicalism

My argument runs as follows:

- i) If phenomenal truths are grounded in physical truths, then these physical-phenomenal grounding relations are covered by laws.
- ii) There are no physical-phenomenal laws.
- iii) If physical-phenomenal grounding relations are covered by laws, but not by physical-phenomenal laws, then there are some physical-phenomenal identifications.

Therefore, if phenomenal truths are grounded in physical truths, then there are some physical-phenomenal identifications.

(By 'there are some physical-phenomenal identifications' I mean that for some phenomenal truth q there is a physical truth p such that $p \equiv q$.)

The rest of this section will explicate and justify each of these premises. My argument is broadly parallel in structure to a well-known argument of Donald Davidson's (1970). Davidson argued from the claims that causation must be covered by laws and that there are no suitable laws involving mental predicates to the claim that if mental events interact causally with physical events, then mental events must themselves be physical. My argument invokes grounding in place of causation: I argue from the claims that grounding must be covered by laws and that there are no suitable laws involving phenomenal predicates to the claim that if phenomenal truths are grounded in physical truths, then phenomenal truths must themselves be physical.

5.3.1 From grounding to covering

The first premise of my argument is:

- i) If all phenomenal truths are grounded in physical truths, then these physical-phenomenal grounding relations are covered by laws.

By 'laws', I mean metaphysical laws, as opposed to laws of nature. Following Jonathan Schaffer (2017b: 305), I have in mind 'counterfactual-supporting general principles', which concern 'the not-causal-but-constitutive generation of a dependent outcome', and which correspond to functions from more fundamental inputs to less fundamental outputs.

For concreteness, I will treat these laws as (interpreted) sentences.¹⁸¹

¹⁸¹ This departs from Schaffer's (2017b) 'minimal' conception of laws as functions. For conceptions which are more similar to my own, see Glazier 2016 and Wilsch 2016. My discussion presupposes that laws are fine-grained enough to stand in entailment and explanation relations, and that it makes sense to quantify their simplicity/strength.

Some paradigm candidates to be metaphysical laws are:

The Fusion-Formation Law: For any things, there is a thing which is their mereological fusion.

The Truth-Making Law: For all p, if p, then the proposition that p is true.

The Disjunction Law: For all p and all q, if p, then p or q.

For example, The Fusion-Formation Law supports counterfactuals of the form ‘If xx didn’t exist, then the mereological fusion of xx wouldn’t exist’, concerns the not-causal-but-constitutive generation of fusions from their parts, and corresponds to a function from things (or the truths that they exist) to their fusion (or the truth that it exists).¹⁸²

I assume that — like laws of nature on the usual conception — these metaphysical laws form a kind of minimal basis for the necessary connections that they support i.e. a basis with as little redundancy as possible. For example, I assume that the following entailments of the candidate laws above are poor candidates to be metaphysical laws:

¹⁸² I don’t mean to suggest with these examples that laws must be universal generalizations — they may, strictly speaking, be universal generalizations prefixed by a distinctive sentential operator ‘It is a law that...’, or (as Fine 2016 and Glazier 2016 suggest) they may involve a distinctive kind of variable-binding. I do assume that they logically entail corresponding generalizations, though if this assumption is denied, I believe that my argument can be restated by modifying the definitions of ‘covering’ and ‘physical-phenomenal laws’.

For any rocks, there is a thing which is their mereological fusion.

For all p, if p, then p or snow is purple.

For all p and all q, if p, then the proposition that p or q is true.

(The first is undermined by being a special case of The Fusion-Formation Law; the second by being a special case of The Disjunction Law; the third by deriving from the combination of The Disjunction Law and the Truth-Making Law.) I view this assumption as a terminological stipulation (akin to the question, facing proponents of the ‘best system account’, of whether to reserve lawhood for the axioms of the best system or to extend it to their consequences). Some may prefer to call my laws ‘fundamental laws’, and extend the term ‘law’ to include all the necessary connections they support, such as the principles above. (Below, I address the worry that my usage of law is too strict for premise i) to be plausible.)

Intuitively, some laws ‘cover’ a grounding relation just in case they allow the entailment of the grounded (under some guise) by the grounds (under some guise). More precisely, suppose that $pp < q$, and let p be the conjunction of the sentences pp. Some laws LL *cover* this grounding relation just in case there is some p^* and some q^* such that:

$p \equiv p^*$; $q \equiv q^*$; and LL together with p^* logically entail q^* .

To illustrate, suppose that:

‘Ann exists’, ‘Bob exists’ < ‘There is a thing which the fusion of Ann and Bob’.

This grounding relation is straightforwardly covered by The Fusion-Formation Law, since this law, together with ‘Ann exists and Bob exists’, logically entails ‘There is a thing which the fusion of Ann and Bob’. Now let Chris be the fusion of Ann and Bob, so that:

‘Ann exists’, ‘Bob exists’ < ‘Chris exists’.

Then this grounding relation (intuitively, the same relation as before, under a different guise) is also covered by The Fusion-Formation Law, on the assumption that:

‘Chris exists’ \equiv ‘There is a thing which is the fusion of Ann and Bob’.

Given this understanding of ‘covering’, premise i) can be thought of as the claim that grounding physicalism entails the following ‘scrutability thesis’:

Assisted Scrutability: All phenomenal truths are logically entailed by some physical truths together with some laws and/or identifications.¹⁸³

If the grounding relations between phenomenal truths and physical truths are covered by laws, then — by definition — there are some laws which, together with the physical truths

¹⁸³ For discussion of various notions of ‘scrutability’, see Chalmers 2012. Chalmers focuses on a priori propositional entailment; I am employing logical sentential entailment.

and some identifications, entail the phenomenal truths. This claim is significantly weaker than the constraints associated with standard scrutability-based objections to physicalism, since it allows that there are some connecting laws and/or identifications which are not themselves scrutable from the physical truths. (My overall argument can be thought of as making the case that any inscrutable connecting laws must be accompanied by inscrutable identifications.)

Premise i) follows from a general principle, which might be called — in Davidson’s (1970) grand style — ‘The Principle of the Nomological Character of Ground’:

All grounding relations are covered by some metaphysical laws.

This principle merely states a necessary condition for there to be a grounding relation; it certainly does not amount to a ‘deductive-nomological account’ of grounding (or grounding-based explanation).¹⁸⁴ It is consistent with many views about what explains grounding facts, such as the view that q’s being grounded by pp is itself grounded in pp (Bennett 2011, deRosset 2013), or the view that grounding facts derive from essence facts (Rosen 2010, Fine 2012, Dasgupta 2014b).¹⁸⁵ (Indeed, these views are themselves naturally cast in terms of metaphysical laws concerning grounding facts!)

¹⁸⁴ For defense of such an account, see Wilsch 2016. Note that being covered by some laws (as I have defined it) could not suffice for there to be a grounding relation, since this would disastrously entail that $pp < q$ whenever pp logically entails q .

¹⁸⁵ Although, I consider an essence-based objection to premise i) below.

The justification for The Principle of the Nomological Character of Ground may be adapted from Schaffer's (2017b: §2) argument that metaphysical explanations require metaphysical laws.¹⁸⁶ We can apply this argument to the grounding-based approach to metaphysical explanation: we may assume that grounding relations are explanatory in that they either constitute or back explanations (depending on whether one is a 'unionist' or a 'separatist'). Hence, grounding relations must be covered by metaphysical laws since the explanations they either constitute or back must be covered by metaphysical laws.¹⁸⁷ Let me briefly summarize Schaffer's defense. (Schaffer does not explicitly argue that metaphysical explanations must be *covered* by metaphysical laws, but this, I believe, is implicit in his case.)

First, since grounding relations are explanatory, they ought to be subsumed under a more general pattern, which unifies the explanandum to relevantly similar cases, shows how to intervene on the explanandum, and provides an understanding of the explanandum. This is only possible if they are covered by some metaphysical law: 'Laws are the stable patterns which unify the phenomena, provide recipes for manipulation, and guide understanding.' (Schaffer 2017b:307). For example, without knowing the general principle which covers the grounding of some composite object's existence, we would not see what unified it with the existence of other composite objects, we would not know how to intervene on its

¹⁸⁶ See also Glazier 2016 and Wilsch 2016.

¹⁸⁷ 'Backing' might be thought to be holistic, so that pp's explaining q need not entail pp's grounding q. Nonetheless, we may assume that if all phenomenal truths are grounded in physical truths, then all phenomenal truths are explained by physical truths (even if not the same ones which ground them). Thanks to Noga Gratvol for raising this point.

existence by intervening on its parts, and we would not grasp the explanatory connection between its existence and its parts.

Second, the need for covering laws is suggested by the analogy between grounding-based explanation and causal explanation: just as the latter needs causal laws to connect explanans to explanandum, so the former needs metaphysical laws to play this connecting role.¹⁸⁸ For example, the causal explanation of a pattern of iron filings in terms of the presence of a magnet is not complete until it is supplemented by some causal law which connects the pattern to its cause. Similarly, the metaphysical explanation of some subject's experience in terms of the physical goings-on in their brain is not complete until it is supplemented by some metaphysical law which connects the subject's experience to its grounds.

Finally, the paradigm cases of grounding relations are covered by metaphysical laws. For example, the grounding relations between a composite object's existence and the existence of its parts is covered by The Fusion-Formation Law, the grounding relations between facts/truths and the truth of corresponding propositions is covered by The Truth-Making Law, and the grounding relations between disjunctive facts/truths and their disjuncts is covered by The Disjunction Law.¹⁸⁹

¹⁸⁸ Again, this does not assume a deductive-nomological account of causation/causal explanation. The claim is that explanation requires connecting laws — not that it is simply a matter of these laws — and it is neutral on how causation might itself be accounted for.

¹⁸⁹ In fact, logical entailments are trivially covered by any laws. This suggests that there is some stronger principle in the vicinity of The Principle of the Nomological Character of Ground e.g. in which 'covering' is understood in terms of subsumption under the function corresponding to some law(s).

For these reasons, I suspect that almost anyone who embraces the idea of grounding will happily accept that it exhibits general patterns of the sort that are described by laws — at least as understood in a suitably ‘lightweight’ sense.¹⁹⁰ We needn’t assume that these laws are ‘brute’ in the sense of having no metaphysical explanation, that they ‘govern’ their instances, or that they somehow correspond to ‘sui generis’ entities. The rest of my argument makes the case that there are not physical-phenomenal laws of the kind that grounding physicalism requires.

5.3.2 Against physical-phenomenal laws

The second premise of my argument is:

- ii) There are no physical-phenomenal laws.

By ‘physical-phenomenal laws’, I mean laws which bridge the logical gap between physical sentences and phenomenal sentences. That is, some laws LL are physical-phenomenal laws just in case there are some physical sentences pp and some phenomenal sentence q such that LL together with pp logically entail q. Physical-phenomenal laws are able to cover physical-phenomenal grounding relations without the aid of any physical-phenomenal identifications.

The sub-argument that there are no physical-phenomenal laws is as follows:

¹⁹⁰ Of course, grounding skeptics might reject the idea of metaphysical laws, but this is not relevant to my argument (which might be viewed as one way of developing grounding skepticism).

ii) a. The metaphysical laws ought to comprise a compact system.

ii) b. Physical-phenomenal laws do not belong in a compact system.

By ‘compact’, I mean both simple and strong. (A system is compact insofar as the conjunction of the laws comprising it is compact.) I do not have fully worked-out accounts of simplicity and strength to hand, but the intuitive ideas suffice for my purposes. Roughly, simplicity is a matter of being concisely expressed in joint-carving terms. For example, assuming that ‘electron’ and ‘negatively-charged’ are joint-carving, the law ‘All electrons are negatively charged’ counts as simple. By contrast, the law ‘Everything which is either an electron and observed before 3000 or a proton and not observed before 3000 is negatively charged’ is not simple, even in the superficially concise form ‘All schmelectrons are negatively-charged’. Roughly, strength is a matter of informativeness, or ‘how much’ a sentence tells us. For example, ‘All electrons carry a charge of -1.602×10^{-19} coulombs’ is more informative than ‘All electrons are negatively-charged’ and ‘This electron carries a charge of -1.602×10^{-19} coulombs’. If one sentence logically entails another, then it is at least as strong as it (but it is notoriously difficult to say anything general about sentences which do not stand in any entailment relation).¹⁹¹

¹⁹¹ Given the standard assumption that metaphysical laws are metaphysically necessary, the notion cannot be understood in terms of quantity of metaphysically possible worlds excluded (though perhaps it can be understood in terms of epistemically possible worlds instead).

The idea behind ii) a. is that compactness is what guides us in our search for the general principles by which the world ultimately works. This is the ‘T-shirt conception’ of laws – as David Chalmers (1996:214) articulates it:

Physicists seek a set of basic laws simple enough that one might write them on the front of a T-shirt; in a theory of consciousness, we should expect the same thing. In both cases, we are questing for the basic structure of the universe, and we have good reason to believe that the basic structure has a remarkable simplicity.

This compactness constraint is epistemic or methodological: the claim is that a non-compact system should thereby be regarded as implausible in practice, not that it is impossible in principle. (The constraint is consistent with, but by no means requires, a Humean account of laws on which their comprising a compact system is what makes them laws.)

Crucially, this methodology applies whether our theory of consciousness posits bridging laws of nature or bridging laws of metaphysics.¹⁹² To be sure, metaphysical laws differ from laws of nature in important ways: they are ‘tighter’ connections, which hold across levels of reality, and with maximal modal strength. But it is hard to see why these differences should invalidate the compactness constraint. When giving a theory of metaphysical laws, we are seeking to describe ultimate explanatory principles by which worldly generation works: this is part of the universe’s ‘basic structure’.

¹⁹² Cf. Pautz 2017:352; forthcoming: §4.

The idea behind ii) b. is that the psychophysical correlations we find in nature do not appear to lend themselves to a compact system. As described by Schaffer (forthcoming: §5.1), this has been a prominent criticism of naturalistic dualism: the naturalistic dualist requires fundamental psychophysical laws, and yet finding a suitably compact system of these laws looks very difficult.¹⁹³ I am merely applying this observation to the grounding physicalist case.¹⁹⁴

As I see it, there are two main obstacles to the compactification of psychophysical correlations. Firstly, there is the problem of compressing the space of phenomenal properties: there seems to be a huge range of possible conscious experiences, differing along many seemingly incommensurable dimensions (Adams 1987: 256-8). Visual, auditory, tactile and olfactory experiences each seem to have their own sets of basic parameters, belonging to distinct similarity spaces. Consider, moreover, the (potentially vast) array of alien experience-types which are inaccessible to human beings (Nagel 1974): should we expect the experiences that bats have when using sonar to fit into any phenomenal similarity space that we are familiar with? The basic phenomenal parameters seem likely to vastly outnumber the fundamental physical properties that we expect to feature in a final theory.

¹⁹³ See, for example, Adams 1987, Latham 2000, Bennett (forthcoming).

¹⁹⁴ See also Collins (2011: §III). This complaint can be viewed as a follow-up to Sider's (2011: §8.2.1) well-known objection that grounding connections must themselves be explained because they involve non-fundamental constituents (see Dasgupta 2014b and Glazier 2016 for responses). The objection here does not claim that metaphysical laws stand in need of explanation, nor that they cannot be explained. Instead, it only relies on the claim that, since they constitute basic explanatory structure, they ought to comprise a compact system.

Secondly, there is the problem of the physical correlates of consciousness, on the other side of the connection: these correlates are not themselves simply expressible in joint-carving terms (Latham 2000: 78–80).¹⁹⁵ For example, notwithstanding the philosophers' legend of C-fibers, the human correlate of pain appears to be a holistic pattern of activation in the central nervous system. This correlate is complex when stated in neural terms (let alone when stated in terms of fundamental physics). Famously, moreover, the human story fails to extend to Martians or sentient AI. Furthermore, there can be little hope for simple functional definitions of physical correlates: try listing the stimuli and behavioral dispositions which are causally connected to pain (together with the background conditions under which these connections obtain). Not only is the list long, but its entries are themselves unlikely to be simply expressible in joint-carving terms.

Putting these two obstacles together, it seems likely that there would have to be many physical-phenomenal laws (at least one for each basic phenomenal parameter), and that these laws would each have to be complex (at least as complex as the physical correlates they feature). Strictly speaking, this allows that one or two sufficiently simple physical-phenomenal laws may feature in reasonably compact systems. But, given the incommensurability of the many distinct phenomenal parameters, and the variety of their physical correlates, any such laws would pertain to their own particular corner of psychophysics, incapable of supporting the rest of the physical-phenomenal laws. Hence,

¹⁹⁵ Schaffer (forthcoming) focuses on the closely related issue that these correlates are either high-level or disunified, rendering them unsuitable to feature in basic laws of nature. This particular problem does not obviously extend to grounding physicalism, since we might expect high-level properties to feature in the level-connecting metaphysical laws.

it is implausible that any such principles would add enough strength to these systems to make the cost in simplicity worthwhile. Moreover, it would seem arbitrary for one or two such principles to be lawful: why, for example, would there be a law concerning human pain, but not concerning Martian pain, or phenomenal redness?

5.3.3 From covering without physical-phenomenal laws to identity

So far, I have argued that the laws alone do not bridge the logical gap between physical truths and phenomenal truths, and yet, if there are physical-phenomenal grounding relations, then the laws together with some identifications must bridge this gap. The idea behind my third premise is that some identifications must therefore play a bridging role:

- iii) If physical-phenomenal grounding relations are covered by laws, but not by physical-phenomenal laws, then there are some physical-phenomenal identifications.

A little more carefully, suppose that some laws LL cover some physical-phenomenal grounding relation. Then, by definition, there are some physical truths pp such that LL together with pp and some identifications logically entail some phenomenal truth q. Suppose also that LL are not physical-phenomenal laws. Then LL together with pp do not logically entail q. There are two (non-exclusive) ways that LL could cover the gap between pp and q without bridging it: there could be some auxiliary identification on the side of the physical truths and/or some auxiliary identification on the side of the phenomenal truth.

Strictly speaking, in neither case must there be a physical-phenomenal identification (that is, an identification between a physical truth and a phenomenal truth). It could instead be

that there is some intermediary kind of truth, X, such that: there is some physical-X identification, yielding an X-phenomenal grounding relation (case 1), and/or some X-phenomenal identification, yielding a physical-X grounding relation (case 2).

Now, it would already be a strong result that grounding can only connect the broadly physical to the phenomenal by embracing some identifications involving an intermediary realm. However, it is hard to see what could plausibly be substituted for ‘X’ in either of the cases above.

Case 1 requires that some X-truths ground the phenomenal truth q. But it is plausible that the only kinds of truths that could ground phenomenal truths are broadly physical truths or other phenomenal truths. (Moreover, it is plausible from the grounding physicalist perspective that all metaphysical laws connect non-fundamental truths to broadly physical truths.)

Case 2 requires that some X-truth be identified with the phenomenal truth q. But again, it is hard to see what other kind of truth a phenomenal truth could turn out to be factually equivalent to, if not a broadly physical truth (they certainly don’t seem to be factually equivalent to normative or aesthetic truths, even if they do bear some interesting connections to such truths).¹⁹⁶

¹⁹⁶ Could the intermediary truths be representational — that is, truths about mental states having certain contents? (Thanks to Avi Sommer for raising this suggestion.) On the most plausible ways of implementing this suggestion, the representational truths turn out to either be broadly physical or phenomenal. Following Dretske (1995) and Tye (1995), one might understand representation itself in broadly physical (causal/informational/teleological) terms: the resulting view connects the phenomenal and broadly physical realms by identity rather than ground. On an alternative approach, phenomenal truths are identified with representational truths that are themselves stated in phenomenal terms (because they involve phenomenal contents and/or modes of representation;

Inspired by Chalmers's (1996: ch.8) speculation, some may be inclined to view the required intermediary truths as 'proto-phenomenal', or even 'informational'. However, these truths are best understood as themselves being either broadly physical or phenomenal.¹⁹⁷ Otherwise, we have no positive conception of their nature: we may as well call them 'some-intermediary-truths-we-know-not-what'. I conclude that the most plausible way of bridging the logical gap between physical and phenomenal truths is to embrace physical-phenomenal identifications.

5.4 Responses

How might grounding physicalists respond to my argument? The first two responses I consider seek to deny that grounding physicalism requires compact physical-phenomenal laws. The other two responses seek to make it plausible that there are some compact physical-phenomenal laws to be found.

5.4.1 Derivative laws

Recall that my argument is roughly parallel in structure to Davidson's (1970): where Davidson argued for identity from the nomological character of causation, I have argued for identity from the nomological character of ground. An important criticism of

see Chalmers 2010: ch.11). These identifications do not take us beyond the phenomenal realm, leaving the gap to the physical realm unbridged. As described below (§5.4.4), representational truths might instead be used to compactify the physical-phenomenal laws.

¹⁹⁷ As discussed below (§5.4.3), proto-phenomenal truths are more plausibly used as the 'gateways' of the phenomenal realm, with compact connections to the physical realm. In my view, 'informational' truths are most plausibly used in the service of a form of idealism, on which physical truths arise from the 'informational structure' realized by a pattern of fundamental phenomenal properties.

Davidson's argument alleges that it relies on an overly strong connection between laws and causation. In particular, Davidson required that causal relations be covered by strict and deterministic laws. However, many have thought that causal relations are covered instead by derivative 'high-level' laws — which may be *ceteris paribus* and/or probabilistic — of the kind that we expect to find in special sciences (including psychology).

A parallel response may be given to my argument. According to this response, just as we should recognize the existence of derivative causal laws, we should also recognize the existence of derivative metaphysical laws. Here are some plausible candidates:

For any x , if x is scarlet, then x is red.

For any region R , if there are some atoms arranged table-wise in R , then there is a table in R .

For any action a , if a involves torturing someone for pleasure, then a is wrong.

These principles are not part of the ultimate roots of the universe's explanatory tree —they operate somewhere higher up, within their own specialized branches. Nonetheless, they seem well-suited to cover grounding relations: indeed, the arguments adapted from Schaffer (2017b), which appealed to the explanatory nature of ground, seem only to require

some kind of general connecting principles, and so do not exclude covering by derivative laws.

Moreover, the response continues, since these derivative laws do not constitute ‘basic structure’, we should not expect them to comprise the kind of compact system that can be written on the front of a T-shirt. Indeed, the incompressible complexity of psychophysical correlations is exactly what we should expect from derivative laws: it is what we find in other specialized branches of the explanatory tree, such as those concerning colours, ordinary objects, or the normative status of actions. Thus, the overall response is that if ‘law’ is relaxed to include these derivative laws — as it must be to justify premise i) — then premise ii) loses its justification.

To avoid name-calling and confusion, I will adopt the terminology suggested by this response: henceforth, I use the term ‘fundamental laws’ to mean what I have hitherto been calling ‘laws’, and use ‘laws’ to include both fundamental and derivative laws. The proposed response is then that the connection between grounding and laws does not warrant the claim I need — that grounding physicalism entails:

Fundamentally Assisted Scrutability: There is some phenomenal truth which is logically entailed by some physical truths together with some *fundamental* laws and/or identifications.

Instead, it only warrants the weaker claim that grounding physicalism entails:

Permissively Assisted Scrutability: There is some phenomenal truth which is logically entailed by some physical truths together with some fundamental *or derivative* laws and/or identifications.

The problem with this reply is that the derivative laws must themselves be metaphysically explained. (It is only reasonable to relax the compactness requirement if these laws are not simply brute posits.) So the reply only pushes the issue a step back: how are these laws themselves to be explained systematically? Effectively, the gap across to phenomenal truths has been bridged at the expense of creating a new gap across to the derivative bridging laws.¹⁹⁸

In particular, I assume that derivative laws must be explained in terms of the fundamental laws and/or the physical truths, and that this explanation must proceed via relations of grounding and/or identity. But then, as I will now argue, these derivative laws must themselves be entailed by some physical truths together with some fundamental laws and/or identifications. Thus, the argument that grounding physicalism entails Fundamentally Assisted Scrutability may proceed via Permissively Assisted Scrutability, from the claim that grounding relations are covered by either fundamental or derivative laws.

¹⁹⁸ There is no analogous difficulty with the response to Davidson's argument, since high-level causal laws are not usually thought to themselves require a causal explanation. However, perhaps a successor of Davidson's argument can be formulated in the light of this response, from the weaker premise that all causal relations must be covered by laws which ultimately rest on strict, deterministic laws, to the weaker conclusion that mental events must either be identical to or grounded in physical events.

The intuitive idea behind my counter-response is that if the explanatory connections between physical and phenomenal truths ultimately rest on fundamental laws which do not bridge the logical gap between physical and phenomenal truths, then these connections must flow through some auxiliary bridging identities. The key lemma is:

All derivative laws are logically entailed by the fundamental laws (LL_{Fun}) together with the physical truths (pp) and identifications.

To see why this lemma is plausible, consider some derivative law L . I assume that it is metaphysically explained by some fundamental laws among LL_{Fun} and/or some physical truths among pp . There are two cases to consider.

Case 1: L is metaphysically explained purely via some identification. Then this identification, together with LL_{Fun} and pp , must logically entail L , so we are done.

Case 2: L is metaphysically explained at least in part via some grounding relation. Then let LL_{Cover} be the laws which cover this grounding relation. Then L must be logically entailed by LL_{Fun} and pp together with some identifications and LL_{Cover} . There are now two sub-cases to consider:

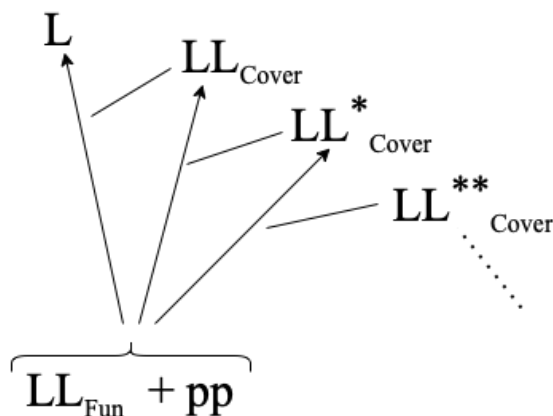
Case 2.1: All the laws in LL_{Cover} are themselves fundamental; we are done.

Case 2.2: Some of the laws in LL_{Cover} are derivative; we repeat the argument for these laws.

The only escape, therefore, is via an infinite sequence of laws belonging to Case 2.2. This amounts to a regress (or perhaps a cycle) of laws, each of which is grounded via covering laws which are themselves derivative. This is not exactly an infinite regress of grounds, but it is an infinite regress of explanatory involvement: each law in the sequence is involved in explaining the previous law in the sequence (by covering its grounding).

Effectively, the complete bridging of the gap to the phenomenal is perpetually postponed on this picture: from the perspective of the fundamental laws and the physical truths, there will always remain a gap that hasn't been closed. I agree with Schaffer's (2017b:316) verdict that such a regress is vicious: 'there cannot be limitlessly descending chains of ever deeper laws'. We should expect that, eventually, the complete explanatory story for each law bottoms out in some ultimate foundation: the laws on which the whole explanatory scaffolding rests.¹⁹⁹

The regress: arrows represent grounding relations; lines connect grounding relations to the laws that cover them.



¹⁹⁹ For further discussion, see Frugé (ms).

It is a short step from the lemma to my conclusion that Permissively Assisted Scrutability entails Fundamentally Assisted Scrutability.²⁰⁰ Suppose that some phenomenal truth q is entailed by some physical truths pp together with some laws LL and/or identifications, where some of these laws, LL^* , are derivative. By the lemma, LL^* are themselves entailed by the fundamental laws and/or identifications together with some physical truths pp^* . Hence, the fundamental laws and/or identifications together with some physical truths, pp and pp^* , entail the phenomenal truth q .

This counter-response assumes that derivative laws must themselves be explained in terms of the fundamental laws and/or the physical truths, without inducing an infinite regress (or cycle) of explanatory involvement. For this reason, I take it to be strongly suggestive rather than demonstrative. The challenge for the grounding physicalist who would employ derivative covering laws is to articulate a conception of such laws which is both acceptably physicalist and escapes the compactness constraint.^{201,202}

²⁰⁰ Note that this argument doesn't quite recover the principle that all grounding relations are covered by fundamental laws, since it may be that the derivative covering laws are themselves only explained by fundamental laws together with some further fundamental truths. But the claim that grounding physicalism entails Fundamentally Assisted Scrutability is all that is needed for the rest of the argument.

²⁰¹ For example, perhaps these derivative physical-phenomenal laws are simply conjunctions of physical-phenomenal counterfactuals. But it is hard to see how such counterfactuals could themselves be explained without somehow invoking some phenomenal truths. Hence, rather than the entire explanatory scaffolding bottoming out in an acceptably physicalist foundation, the derivative physical-phenomenal laws would be involved in explaining themselves. (Similar remarks apply to the proposal that derivative laws be understood as generalizations which are grounded in their instances.) Thanks to Jonathan Schaffer and Verónica Gómez for discussion.

²⁰² Another challenge to Fundamentally Assisted Scrutability leads to a parallel dialectic. Some have suggested that, as with causation, grounding sometimes requires 'enabling background conditions' (e.g. Epstein 2015: ch.9, Skow 2016:111 and Baron-Schmitt 2021). Consider, for example:

5.4.2 Essences

A popular idea amongst fans of ground is that grounding relations derive in some way from essences (Rosen 2010, Fine 2012, Dasgupta 2014b). For example, the grounding relation between a thing's scarletness and its redness is itself partly grounded in a truth about the essence of redness: namely, that it is essential to redness that anything which is scarlet is also red. It is natural on this view to regard essentialist truths as taking up at least some of the slack when it comes to covering grounding relations. Hence, The Principle of the Nomological Character of Ground must be replaced by an even grander principle — The Principle of the Essentialist-Nomological Character of Ground:

All grounding relations are covered by some metaphysical laws and/or essentialist truths.

This amendment no longer guarantees that physicalists are committed to Assisted Scrutability. Instead, given this principle, physicalists need only be committed to:

Ivy is cheeky < All my nieces are cheeky.

This might be thought to be enabled by the background condition that Ivy is my only niece. Hence, one might hold that grounded truths will only be scrutable from grounding truths together with laws, identifications *and background conditions*. However, a view on which the physical truths only ground the phenomenal truths in the presence of some background conditions which are not themselves physical or grounded in the physical would not seem acceptably physicalist — for example, it would allow for the possibility of zombie-worlds (cf. Hawthorne's (2002) verdict on the possibility of 'blockers').

Essentially Assisted Scrutability: There is some phenomenal truth which is logically entailed by some physical truths together with some laws and/or essentialist truths and/or identifications.

To illustrate, take the truth that Joe is in pain. Essentially Assisted Scrutability no longer requires that there be some laws and/or identifications which bridge the logical gap between the physical truths and this phenomenal truth. Instead, the gap may be bridged by a truth about the essence of pain: for example, that it is essential to pain that anyone who has firing C-fibers is in pain. In this way, grounding physicalists may avoid bridging identities by using essences instead.

There are two problems for this proposal, however. The first — and perhaps less serious — is that essentialist connections between the physical and the phenomenal seem to run counter to the non-reductive motivation for grounding physicalism. Part of the attraction of grounding phenomenal truths in physical truths is the idea that it allows them to be genuinely non-physical in nature. Yet the essentialist proposal seems to simply replace bare and incredible assertions of physical identity with equally bare and incredible assertions of physical essence. (Indeed, many non-reductionists seem to be motivated by the intuition that the essence of a phenomenal property is exhausted by the way it feels.)²⁰³ Moreover, these surprising essence facts would be naturally explained by identifications: a natural explanation for its being essential to pain that anyone who's C-fibers are firing is in pain

²⁰³ In this vein, Rosen (2010: §13) observes that physical-phenomenal bridging laws are plausibly viewed as 'Moorean connections': exceptions to the generalization that metaphysical laws derive from essences.

is simply that to be in pain just is to have firing C-fibers (or just is to be in some disjunctive physical state of which firing C-fibers is a disjunct).²⁰⁴

The second problem is a dilemma. On the first horn, the invoked essentialist truths are simply disguised physical-phenomenal identifications, or are explained in terms of such identifications.²⁰⁵ On the second horn, these essentialist truths are nothing to do with identifications. But then the T-shirt problem hasn't been solved: how are the essentialist truths to be systematically explained, if not via identifications? And if unexplained, why are they exempt from the compactness constraint?²⁰⁶

Effectively, this proposal takes the mess of physical-phenomenal bridging laws and merely adds that they are essential to the phenomenal properties in question. It is hard to see why this rebranding should make the overall theory any more plausible. Compare the analogous case of nomic essentialism: the view that the laws of nature are essential to the properties they involve. If true, this view does not exempt physicists from searching for a compact system of fundamental laws. Intuitively, if they are not tied to identifications, essentialist

²⁰⁴ See Correia & Skiles (2017) for a related proposal connecting both essence and ground to identification.

²⁰⁵ The identification of a phenomenal property with a physical property may be thought of as a 'real definition', which may in turn be thought of as spelling out the property's 'full essence'.

²⁰⁶ I argue below that the compactness constraint does not apply to identifications, which are to be systematized by a meta-semantic theory. It is hard to see how any claim about essence could be given a similarly meta-semantic justification, unless this justification went via some identification.

truths still add to the world's complexity; just like fundamental laws, they would constitute brute, worldly, explanation-backing structure. Hence, they should fit on a T-shirt.²⁰⁷

5.4.3 Physicalist panpsychism

In light of the difficulties with invoking non-compact derivative laws and/or essentialist truths, I believe that the most promising strategy for grounding physicalists is to address the argument that physical-phenomenal laws do not belong in a compact system head-on. This would mirror the most popular dualist response to the 'T-shirt problem': rather than being viewed as an insurmountable obstacle, it is seen as motivating the search for compact physical-phenomenal laws. Indeed, this has become a lively research program — one which grounding physicalists, as well as dualists, ought to be engaged in. In this section and the next, I discuss what I take to be the two most promising lines of inquiry.

One prominent idea is that there are some 'proto-phenomenal' properties — resembling in some important way the phenomenal properties we are familiar with — which are connected to fundamental physical properties such as charge and mass via a compact system of laws, and which themselves combine to yield familiar phenomenal properties.²⁰⁸

If this could be made to work, it would avoid both the obstacles to compactification described above: unlike the space of phenomenal properties, the space of proto-

²⁰⁷ Again, I am not making Sider's complaint that, since these truths involve non-fundamental notions, they must be explained. Rather, I am allowing that these truths may be fundamental, or even 'autonomous' (Dasgupta 2014b), and arguing that, if left unexplained, they ought to at least be compactly systematizable.

²⁰⁸ See e.g. Chalmers 1996: §8.4, Collins 2011: §IV; Coleman 2012, Roelofs 2019.

phenomenal properties would only have a few dimensions, and unlike the physical correlates of phenomenal properties, the correlates of proto-phenomenal properties would be simple to state in fundamental terms.²⁰⁹

The crucial observation is that, were some such system to be discovered, it could equally serve the grounding physicalist's purpose: the laws in question could be regarded as metaphysical. On the resulting 'physicalist panpsychism', phenomenal truths would be grounded (at least partially) in proto-phenomenal truths, and the latter would themselves be grounded (via a compact system of fundamental bridging laws) in physical truths.

Indeed, physicalist panpsychism has a significant dialectical advantage over dualist panpsychism. The latter faces the notorious 'combination problem': generating phenomenal truths from proto-phenomenal truths seems to involve the very same challenges that led the dualist to reject physicalism in the first place. In particular, the existence and phenomenal properties of 'higher-level subjects' like ourselves do not seem scrutable on the basis of the proto-phenomenal properties of micro-physical constituents. Embracing further laws to connect the higher-level phenomenal realm to the proto-phenomenal realm lands us back with the T-shirt problem: the higher-level phenomenal realm remains multi-dimensional, and proto-phenomenal correlates are presumably just as complex as physical correlates.

²⁰⁹ This view is 'panpsychist' in that the proto-phenomenal properties are usually assumed to be widespread, and perhaps even ubiquitous. The assumption is natural, given that these proto-phenomenal properties must be had by our own constituents (in order to yield our own phenomenal properties), and they must be simply related to the fundamental physical properties. It is hard to see what simple distinction between our own constituents and the rest could be made in fundamental physical terms.

By contrast, the grounding physicalist can posit inscrutable metaphysical laws connecting the proto-phenomenal to the phenomenal. Since these laws presumably wouldn't comprise a compact system, they would have to be derivative. However, they could be explained via inscrutable auxiliary identities between phenomenal truths and (presumably logically complex) proto-phenomenal truths: whilst grounding physicalists deny that there are identities bridging the physical and phenomenal realms, it is no part of their position that there cannot be any identities bridging the proto-phenomenal and higher-level phenomenal realms. On this view, grounding connections would still play the crucial role in bridging the gap between the physical and phenomenal, albeit with the help of some identities within the phenomenal realm.²¹⁰

Of course, there is a clear cost associated with panpsychism, which many will find unacceptable: the idea that micro-physical entities bear genuinely proto-phenomenal properties remains highly speculative. The more closely these properties resemble phenomenal properties, the more absurd the view appears to be; the less closely they resemble phenomenal properties, the more it appears to collapse into a terminological variant of the standard physicalist picture. But it may be that — like dualists — grounding physicalists can only hope for a suitably compact system of fundamental laws by embracing panpsychism.

²¹⁰ This does not help with the combination problem construed as a problem for the very notion of 'phenomenal combination': see Coleman 2012 for discussion of what this problem might be. But it does avoid the 'tu quoque' problem: the dialectical point that an important argument against physicalism is undermined.

5.4.4 Two-factor physicalism

There may be an alternative way of compactifying the physical-phenomenal laws, however.²¹¹ The major theories of consciousness involve a certain division of labor. One part of the theory concerns what it takes for a mental state (typically thought of as a mental representation) to be conscious. This part of the theory might appeal to functional role and/or neurophysiological realization. For example, it might appeal to monitoring by some higher-order state (Rosenthal 1993, Lycan 1996), broadcasting in the ‘global workspace’ for widespread access and use (Baars 1988, Dehaene & Naccache 2000), or local recurrent processing within the sensory cortex (Lamme 2006, Block 2007).

Another part of the theory tells us how, given a conscious mental state, its phenomenal character is determined. Thus, this second part targets truths such as ‘S is a burning sensation’, or ‘S is a reddish sensation’. To explain phenomenal character, representationalist theories draw on the state’s content (Dretske 1995, Tye 1995) or the content of a corresponding higher-order state (Carruthers 2000, Rosenthal 2002). We may also appeal to the ‘format’ of the state in question, its ‘syntactic’ features, causal role, or intrinsic neuro-physiological makeup.

This division of labor provides a potential route to compactification for grounding physicalists. In particular, they might maintain that a law is only needed to get us from a state’s intrinsic and/or functional properties to its being conscious. This consciousness-generating law would take the form: for any state S, if S has physical property *X* then S is

²¹¹ Thanks to Ned Block for suggesting this response and to Verónica Gómez for discussion.

conscious. Truths about phenomenal character would then be derivable from this law together with identifications of the following form:

$$\text{'S has } Q \text{'} \equiv \text{'S is conscious and has } Y \text{'}$$

where Q is a phenomenal property and Y is a physical property. For example, the complete explanation for why some state feels painful might proceed as follows: first, its consciousness is grounded in its being globally broadcast; second, its painfulness is identified with its being conscious and representing tissue damage.

The general strategy for compactification is to reduce the richly various phenomenal truths via identifications to phenomenally homogenous truths involving the single phenomenal notion of state-consciousness. By 'factorizing out' state-consciousness from phenomenal character in this way, all the rich variation may be quarantined in the remaining physical factor, thus avoiding the 'many phenomenal parameters' aspect of the T-shirt problem. The phenomenally homogenous truths may then be targeted by a single physical-phenomenal law, with the crucial gap between physical and phenomenal truths bridged by grounding. It might be hoped (although, it is far from clearly being the case) that this law will be suitably concise; in particular, it might be hoped that the functional/neural correlate of state-consciousness (as opposed to all the particular kinds of conscious experience) is reasonably simple.

There's another way of turning the trick. Instead of factorizing phenomenal truths conjunctively, we might factorize them into an irreducible phenomenal relation and richly varying physical relata. That is, we might endorse intra-phenomenal identifications of the following form:

$$\text{'S has } Q\text{'} \equiv \text{'S is conscious-of } Y\text{'}$$

where again Q is a phenomenal property and Y is a physical property.²¹² For example, a state's feeling painful might be identified with its being conscious-of tissue damage, and a state's being a reddish sensation might be identified with its being conscious-of a certain surface reflectance property. Phenomenal truths involving the consciousness-of relation might then be grounded in broadly physical truths via a metaphysical law. Again, if the

²¹² Following Russell (1912), being conscious-of a property might be understood in terms of a relation to an object which instantiates that property. Clearly, for the view to be physicalist, the objects in question better not be Russellian sense-data; internal brain states are the most natural physical candidates. To allow for illusion/hallucination, the consciousness-of relation might be better understood in terms of a 'more direct' relation to the property in question. Following Dretske (1995) and Tye (1995), these properties are most naturally understood as ordinary physical properties of external objects which the representation 'tracks' (or functions to track).

Pautz (2017:§5) advocates a non-reductive version of this relational model, but on his view the properties in question are non-physical 'sensible properties'/'secondary qualities' (which may never be instantiated by anything; cf. Chalmers 2006). Since these sensible properties do not themselves have any physicalist reduction, this view does not avoid the 'many parameters' problem.

There is also the possibility of a 'three-factor' view, which factorizes the conscious-of relation itself as follows:

$$\text{'S is conscious-of } Y\text{'} \equiv \text{'S is conscious and S represents } Y\text{'}$$

The idea is then that the relevant representation relation is itself irreducible, but grounded according to some further, suitably concise metaphysical law.

Whether two-factor or three-factor, these relational versions of grounding physicalism would appear to inherit the difficulties of representationalism: for discussion, see Block 1996, Chalmers 2010: ch.11, Papineau 2016.

physical correlate of this relation is reasonably simple (perhaps involving some special kind of causal/nomic connection), then the law in question might be suitably concise.

The question for these ‘two-factor’ versions of grounding physicalism is whether the non-reductive advantages of treating the first component of the theory as a law rather than an identification outweigh the cost of complicating the overall lawbook. Recall the three initial motivations for denying physical-phenomenal identifications: they are counterintuitive and mysterious; they ignore the special significance of consciousness; and they face a trilemma resulting from multiple realizability. The two-factor theories considered above might avoid these issues as applied to state-consciousness/the conscious-of relation, but they do nothing to address them as applied to phenomenal properties themselves.

First, by making all phenomenal variation physical variation, they seem to retain much of what is mysterious and counter-intuitive about assimilating phenomenal truths to physical truths. (If anything, the denial that consciousness simpliciter is physical seems to pile mystery upon mystery: the idea that it is a physical matter what kind of conscious experience a state is seems even stranger when consciousness itself is not physical!)

Second, two-factor views may be able to accommodate the idea that consciousness and/or the conscious-of relation have special significance, but they do nothing to accommodate the special significance of particular phenomenal characters, such as pleasure and pain (since these phenomenal characters involve some complex physical property, which presumably has many intrinsically similar candidates in its vicinity).

Third, by identifying truths involving phenomenal properties with truths involving some corresponding broadly physical properties (together with an irreducible phenomenal component), they face the challenge of further explaining the truths about these physical properties in a way which avoids extensional inadequacy, disunification, and causal ineligibility. The idea that all phenomenal properties share some irreducible phenomenal component might help to accommodate the multiple realizability of this shared component, but presumably the multiple realizability of phenomenal properties goes beyond this: even fixing the way that the shared component is realized, there are many ways of realizing the specific phenomenal property.

All in all, the non-reductive advantages of two-factor grounding physicalism look too flimsy to justify the addition of a special consciousness-generating metaphysical law.²¹³

5.5 Hybrid Physicalism

5.5.1 Motivation

Some physicalists will view the problem that I have presented for grounding physicalism as vindicating a return to extremism: a ‘pure’ identity physicalism, on which all metaphysical explanation is backed by identification. However, this won’t be attractive for

²¹³ A further issue concerns the integration of the consciousness-generating law into the lawbook. In particular, there is the possibility of mismatch between the truths that the other laws output on the basis of the fundamental physical truths and the broadly physical truths that the consciousness-generating law needs as inputs. It is an open question whether it would be legitimate for grounding physicalists to rely on identifications to bridge this gap, or whether instead all metaphysical laws must operate within a ‘common language’. In the latter case, the consciousness-generating law may turn out to be far more complex than initially hoped.

those who found the initial motivation for grounding physicalism compelling: the need to reconcile the unity, significance, causal efficacy, and apparent irreducibility of phenomenal properties with their non-fundamentality.

My argument suggests a natural way of navigating between the horns of this dilemma, which I call ‘hybrid physicalism’. According to hybrid physicalism, the metaphysical explanation of consciousness proceeds in two stages. First, phenomenal properties are identified with some high-level but broadly physical properties: perhaps these properties are neuro-biological, or perhaps they are functional. Secondly, the truths about these high-level but broadly physical properties are grounded in some fundamental physical truths. (Hence, this view preserves the positive component of grounding physicalism at the expense of its negative component.)

The first stage allows this hybrid approach to avoid the problem with grounding physicalism presented above: by identifying phenomenal properties with broadly physical properties, we avoid the need for fundamental physical-phenomenal laws, so their generation from fundamental physics can be subsumed by a suitably compact system. For example, if phenomenal properties are identified with functional properties, then their generation can be covered by a very general law governing the generation of functional roles from realizers and laws of nature. This law isn’t specific to the phenomenal realm: it covers the plethora of notions with functional analyses, including, presumably, ordinary object predicates such as ‘is a table’, and perhaps even extending to spatiotemporal notions such as ‘is five feet from’.

Meanwhile, the second stage promises to alleviate (at least some of) the problems for the ‘pure’ version of identity physicalism which motivated the shift towards grounding physicalism in the first place. By denying that phenomenal properties can be given any analysis in fundamental physical terms, we avoid turning them into complex, arbitrary-looking and explanatorily irrelevant disjunctions of actual and/or possible realizers. For example, we can identify phenomenal properties with high-level functional properties, without needing to provide any identification of truths involving these properties with the truths of fundamental physics.²¹⁴

Some problems do remain, however. On the cruder visceral side, the concerns about ‘incredible identity assertions’ and radical eliminativism loom large (presumably, restricting the reduction to higher-level physical properties doesn’t make it any more intuitive). On the nuanced theoretical side, the putative dilemma between extensionally inadequate realizer properties and causally inefficacious role properties must be addressed. Moreover, it is unclear whether identifying phenomenal properties with high-level physical properties helps to accommodate their ‘significance’. On the one hand, there would still presumably be many intrinsically similar properties in their vicinity. On the other, by denying that these high-level properties are reducible, we allow that they are joint-carving

²¹⁴ Compare the idea that sets are grounded in their members, which avoids the need to reduce the rich truths of set theory to truths which are free of set-theoretic notions. Whether or not the hybrid approach has a genuine advantage here remains to be seen. First, how exactly is the hybrid approach to be implemented: in particular, how are the compact laws in question to be spelled out? Second, what resources does the pure identity physicalist have to accommodate multiple realizability: in particular, could there be analyses of high-level functional properties which accommodate their naturalness? For a step in the first direction, see Schaffer’s (2021) ‘mind enactment machine’. For a step in the second direction, see Gómez Sánchez’s (forthcoming) ‘iterated lawbook’ construction.

(and hence that they may feature in reasonably simple normative, epistemic and meta-semantic principles).

The availability of hybrid physicalism invites a comparative construal of my argument. Rather than wielding compactness as an absolute constraint on fundamental metaphysical laws, we might argue instead that, *ceteris paribus*, physicalist theories are attractive to the extent that they posit compact fundamental metaphysical laws. Since hybrid physicalism's laws are much more compact than the physical-phenomenal laws that grounding physicalism needs, hybrid physicalism is (*ceteris paribus*) more attractive. (Whether the *ceteris* are indeed *paribus* in this case will depend on the plausibility of the required physical-phenomenal identifications.)

5.5.2 Systematicity

It might be worried, however, that hybrid physicalists (and indeed, pure identity physicalists) face a difficulty analogous to the grounding physicalists' need for compact laws. Don't they simply replace grounding physicalism's compendious system of physical-phenomenal laws with their own compendious system of physical-phenomenal identities? In sartorial terms, hasn't the hybrid physicalist simply snuck all the small print from the front of the T-shirt (where the laws are) around to the back of the T-shirt (where the identifications are lurking)?!

To alleviate this concern, consider a case in which identifications are much less controversial: the explanation of truths about ordinary objects like tables and chairs. Here,

I take it, it is very plausible that our talk of these objects may be reduced in some sense to truths about their microphysical constituents. For example, it is plausible that for something to be a table just is for it to be the mereological fusion of some particles arranged into a certain shape in such a way that they exhibit certain behavioral dispositions, and that the property of being a chair, or a lamp, or a window, or a soccer ball, can be given analyses along similar lines.²¹⁵

Can this list of analyses be compactified into a simple system of basic principles? Surely not: even if the analyses are similar in the sense that they all involve, say, being a fusion of particles with certain shape and disposition-properties, there is no overarching principle from which each analysis may be derived. There is no concisely expressible function which takes in predicates such as ‘is a table’ and ‘is a chair’ and outputs their analyses: a machine which learnt the analyses of many such predicates would have no means of extrapolating the analysis of some entirely new predicate. Instead, as when we are learning a new language, we must proceed case-by-case: each primitive predicate must be taken on its own merits. The resulting compendious list of analyses does not seem to impugn identity physicalism about ordinary objects in the least.²¹⁶

The analogy with ordinary objects is a purely defensive manoeuvre: it suggests that identifications should not be required to be compact, but does not indicate what the relevant

²¹⁵ Setting aside complexities arising from the role that intentions and other mental phenomena may play in the correct application of these predicates — if necessary, they can be replaced by artificial predicates whose application ignores such factors.

²¹⁶ Although, as I discuss shortly, there should be a kind of ‘meta-semantic’ systematicity.

difference from laws is. Indeed, the initial motivation for the T-shirt problem was that the universe's 'basic structure' should be simple. But physical-phenomenal identifications seem aptly characterized as 'basic structure': they appear to be brute, unexplained posits which provide the physicalist's explanatory 'scaffolding'. They support physical-phenomenal explanations in the same way that physical-phenomenal laws do for grounding physicalists. As Chalmers and Jackson (2001:353) put it:

Ontologically, these identities may differ from laws. But epistemically, they are just like laws. They are epistemically primitive psychophysical "bridging" principles that are not themselves explained, but that combine with physical truths to explain phenomenal truths.²¹⁷

In response, it might be maintained that identifications do not themselves need to be explained: as Dorr (2016:41) observes, they are 'excellent stopping places for explanation'. However, I don't think this is the source of the crucial asymmetry, since it is unclear that fundamental laws can't be equally excellent stopping places: why should the ultimate principles by which worldly generation works call for any deeper explanation? We expect to rest content with the fundamental dynamical laws of some final physical theory; why would we not do the same with the fundamental metaphysical laws?

As I see it, the contrast derives instead from the constraints which apply when we are positing these two forms of brute explanation-backing structure. Grounding-backed explanation is importantly different from identification-backed explanation, and so we

²¹⁷ Of course, they have in mind the laws of naturalistic dualism, but the point is arguably even clearer with respect to the laws of grounding physicalism, since these are also used to support *metaphysical* explanation.

should not necessarily expect the constraints on the two kinds of explanation to be the same. In the first case, we are positing a certain kind of worldly structure: a principle, or mechanism, by which derivative portions of reality are generated from more basic portions. But in the second case, we are reducing — or collapsing — worldly structure: we are positing that there is only a single portion of reality where there might have seemed to be two. For this reason, it should not be surprising if the plausibility of these two kinds of explanation ought to be evaluated in quite different ways.²¹⁸

One way in which the two kinds of explanation plausibly differ concerns the sense in which they ought to be systematic. Our inability to compactify physical-phenomenal laws makes them unsuited to be fundamental — and hence, makes them ‘cry out for’ deeper explanation, just as the dynamical laws of a physical theory would if they were similarly complex and arbitrary-looking. But it seems unreasonable to expect identification-backed explanation to be systematic in the same way: rather than being ‘nominally systematic’ — subsumed by a compact system of fundamental laws — we should only expect it to be ‘meta-semantically systematic’.²¹⁹

Meta-semantic systematicity is not motivated by the idea that worldly generation is systematic but rather by the idea that our language/thought latches onto reality in a

²¹⁸ Cf. Pautz (2017:355): ‘intuitively identifications have a unique feature: unlike certain kinds of brute “grounding laws”, *they don’t add to the complexity of our theory.*’

²¹⁹ My ‘Two Approaches to Metaphysical Explanation’ (ms) situates this contrast within the broader distinction between grounding (or ‘generation’) and identification (or ‘reduction’).

systematic way. A meta-semantically systematic theory involves a dictionary of reductive analyses rather than a lawbook of generative principles. A dictionary is not merely an incompressible compendium, however: we should expect it to be subsumed under some reasonably systematic meta-semantic account. This account will connect the analysis of a given word/concept to its ‘meta-semantic profile’: that is, whichever of its features are relevant to its latching onto the world in the way that it does (such as its history of use, its causal/nomic connections to the environment, its biological function, etc.). It is this meta-semantic theory which provides the general principles describing when two ways of representing latch onto a single portion of reality. Roughly, a meta-semantic theory stands to a system of metaphysical laws as a semantic theory stands to a system of physical laws. Where the latter is foundational and comprehensive, the former is high-level and specialized. For this reason, they should be held to quite different standards.²²⁰

In the case of consciousness, the phenomenal truths seem to defy generation from physical truths via a nomically systematic lawbook. But this does not prevent their being identified with physical truths via a meta-semantically systematic dictionary, where this dictionary is subsumed by a theory which connects the meta-semantic profiles of our phenomenal concepts to their physical analyses. Treating the difference between phenomenal and physical truths as worldly creates a gap that cannot be bridged by a suitably systematic

²²⁰ It might be worried that, given the broadness of my conception of ‘grounding’, this opens the door for special pleading in the case of certain ‘grounding’ relations. Couldn’t the fan of realization, for example, hold that realization-backed explanations are also an exception to the nomic systematicity constraint? (Thanks to Luke Roelofs for raising this objection.)

I think not. The need for nomic systematicity derives from very general features that all grounding relations share: they constitute explanation-backing structure which is *worldly*. It is because grounding-backed explanations add to the world that they ought to be nomically systematic.

lawbook; treating it as representational may allow the gap to be bridged by a suitably systematic meta-semantics.

Thus, a compendious list of identifications may be justifiable in a way that a compendious list of laws cannot be. Those who are impressed by the idea that identifications ought to be scrutable — that is, a priori entailed by the fundamental truths — will view the result as a dilemma for physicalism. Metaphysical laws, being worldly, needn't be scrutable, but they ought to be compact; identifications, being representational, needn't be compact, but they ought to be scrutable.²²¹ Hence, the gap between physical and phenomenal truths is not satisfactorily bridged in either way.

However, there are reasons to dislike scrutability as a constraint on metaphysical explanation which do not rely on the worldliness of the connections. For example, one might argue that paradigm cases of unproblematic identifications — such as those connecting truths about ordinary objects to truths about their mereologically simple parts — are not themselves scrutable. Or one might hold that phenomenal concepts (like indexical concepts, and normative concepts) are special in a way that makes it unreasonable to expect the truths involving them to be scrutable.²²² If correct, these arguments would

²²¹ As Schaffer (2021:185) puts it: 'causal laws and grounding principles are both principles of real generation in nature... From such a realist perspective, it seems simply wrong-headed to demand that the clockwork of nature be transparent to pure reason.' Identifications are not part of 'the clockwork of nature': indeed, this is precisely why they needn't be nomically systematic.

²²² For criticism of scrutability as a constraint on identifications, see Loar (1990), Tye (1999), Block & Stalnaker (2001), Balog (2012) and Mehta (2019). I discuss the issue in 'Scrutability and Metaphysical Explanation' (ms).

show that scrutability is no more appropriate as a constraint on identifications than on metaphysical laws.

Finally, one might wonder about the scope and limits of hybrid physicalism. I have focused on the the question of how phenomenal truths are to be connected to the broadly physical realm because it constitutes a central challenge for the physicalist world-view. But of course, there remains the question of the connections within the broadly physical realm: how are broadly physical higher-level truths — such as those concerning biological organisms, or ordinary macroscopic objects — connected to the truths of fundamental physics? Almost any higher-level realm of truths seems to have the two features that led to the T-shirt problem: many independent dimensions of similarity, and complexity of physical correlates. For this reason, I expect the hybrid approach to apply widely. Many higher-level truths can only be systematically grounded once they have been ‘translated’ into the language of the fundamental metaphysical laws: we should not expect these laws to deal directly in tigers and tables. But the defense of this broader view of metaphysical explanation must be left to future work.²²³

²²³ I also expect the hybrid approach to apply to other forms of explanation, such as causal and mathematical explanation. Perhaps explanation in general has a worldly, law-backed dimension and a representational, identification-backed dimension.

5.6 Concluding Remarks: The Explanatory Gap

I have been presenting a challenge for grounding physicalism's explanation of consciousness. In conclusion, I would like to relate this challenge to the so-called 'explanatory gap' between physics and phenomenology.

It seems likely that a few distinct though inter-connected concerns lie in the vicinity of the explanatory gap as it has been presented in the literature. However, one central strand of the challenge for physicalist explanations (which strikes me as especially central to our 'gappiness' intuitions) has to do with 'arbitrariness'. Here is Levine's (1983:357) classic articulation of the worry:

what is left unexplained by the discovery of C-fiber firing is *why pain should feel the way it does!* For there seems to be nothing about C-fiber firing which makes it naturally "fit" the phenomenal properties of pain, any more than it would fit some other set of phenomenal properties.

In later work, Levine (2006:145) calls this 'sense of arbitrariness' the 'core problem':

The connection between the neurological description and our first-person conception of what it's like seems totally arbitrary. One feels that this neurological configuration could just as easily have gone with a bluish visual experience as a reddish one.

Arbitrariness, as I see it, is a matter of lacking 'privilege' over some relevant alternatives. In the case of explanation, it is the proposed explanatory connection which is alleged to lack the required privilege. An arbitrariness concern is to be answered by justification, not explanation: the choice of the explanatory connection over its alternatives must be justified, but need not be explained — the obtaining of the connection may simply be a brute matter

of fact. Depending on how exactly ‘privilege’ is construed, and what the ‘relevant alternatives’ are, this justification may show that the connection in question is privileged, or it may simply make it plausible despite being somewhat arbitrary. For example, if the privilege of a dynamical law is understood in terms of its relative simplicity, it may be that ‘ $F = ma$ ’ is privileged over ‘ $F = ma^2$ ’ or ‘ $F = ma + k$ ’, but not privileged over ‘ $F = 2ma$ ’ (which is only more complex relative to a choice of units). Nonetheless, the choice of ‘ $F = ma$ ’ over ‘ $F = 2ma$ ’ is justified by our evidence.

From the hybrid physicalist perspective, it is natural to distinguish two quite different aspects of this arbitrariness concern as it applies to physicalist explanations of consciousness. The aspect which is usually focused on concerns the arbitrariness of physical-phenomenal identifications: what privileges these identifications over their myriad conceivable alternatives? This issue has not been my focus here, but I agree with the general consensus that it is best addressed by focusing on ‘phenomenal concepts’: identifications are to be justified (though not explained!) in terms of the meta-semantic profiles of the concepts involved.²²⁴ The idea is that the underlying analyses have a kind of ‘representational privilege’: the corresponding high-level properties are highlighted, at least in part, by our own particular (and perhaps parochial) perspective on the world, as embodied by our concepts. This allows that these analyses are not ‘objectively privileged’ in the sense that aliens might have phenomenal-like concepts which latch onto intrinsically similar high-level properties (cf. Lee 2018).

²²⁴ Ideally, this meta-semantic approach would also provide an error theory for the intuition that the identifications in question are ‘incredible’.

However, a second aspect of the arbitrariness concern is just as important: the idea that fundamental physical-phenomenal laws are also unacceptably arbitrary. Their lack of compactness belies an absence of objective privilege. Compare a physical theory according to which the nomic connection between force, mass and acceleration varies unsystematically across spacetime. We would naturally regard the resulting lawbook as implausibly arbitrary: such complicated and compendious dynamics would not seem to be privileged over the myriad conceivable alternatives.

From this perspective, the arbitrariness concern amounts to a dilemma: physicalists must weigh the apparent arbitrariness of physical-phenomenal identifications against that of physical-phenomenal laws. The point I have been pressing is this: grounding physicalists only avoid the first at the cost of embracing the second.

References

- Adams, R. M. (1979). Primitive thisness and primitive identity. *Journal of Philosophy* 76(1):5-26.
- (1987). Flavors, Colors, and God. *The Virtue of Faith and Other Essays in Philosophical Theology*: 243–62. Oxford: Oxford University Press.
- Albert, D. Z. (1996). Elementary Quantum Metaphysics. In J. T. Cushing, A. Fine, and S. Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, Dordrecht: Kluwer.
- (2013). Wave Function Realism. In A. Ney and D. Z. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.
- (2015). *After Physics*, Cambridge, MA: Harvard University Press.
- (unpublished-a). On the Emergence of Space and Time.
- (unpublished-b). How to Teach Quantum Mechanics. Available at <http://philsci-archive.pitt.edu/15584/>.
- Allori, V. (2013-a). Primitive Ontology and the Structure of Fundamental Physical Theories. In A. Ney and D. Z. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.
- Allori, V. (2013-b). On the Metaphysics of Quantum Mechanics. In S. Le Bihan (ed.) *La Philosophie de la Physique: d'Aujourd'hui a Demain*, Paris: Editions Vuibert. Accessed online at <http://philsci-archive.pitt.edu/9343/>.
- Armstrong, D. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Audi, P. (2012). Grounding: Toward a Theory of the In-Virtue-Of Relation. *Journal of Philosophy* 109(12):685–711.
- Baars, B. (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Bacon, A. (2019). Is reality fundamentally qualitative? *Philosophical Studies* 176 (1):259-295.
- (2020). Logical Combinatorialism. *Philosophical Review* 129 (4):537-589.
- Balog, K. (2012). In Defense of the Phenomenal Concept Strategy. *Philosophy and Phenomenological Research* 84 (1):1-23.
- Barnes, E. (2014). Going Beyond the Fundamental: Feminism in Contemporary Metaphysics. *Proceedings of the Aristotelian Society* 114 (3pt3):335-351.
- Barnes, E. & Williams, J. R. G. (2011). A Theory of Metaphysical Indeterminacy. In K. Bennett & D. W. Zimmerman (eds.), *Oxford Studies in Metaphysics Vol. 6*: 103-148.
- Baron-Schmitt, N. (2021). Contingent grounding. *Synthese* 199 (1-2):4561-4580.

- Baron, S. & Norton, J. (2019). Metaphysical Explanation: The Kitcher Picture. *Erkenntnis*:1–21.
- Bell, J. S. (1987). *Speakable and Unspeakable in Quantum Mechanics*, Cambridge: Cambridge University Press.
- Benacerraf, P. (1965). What Numbers Could Not Be. *The Philosophical Review* 74(1): 47–73.
- Bennett, K. (2004). Spatio-temporal coincidence and the grounding problem. *Philosophical Studies* 118(3):339–371.
- (2009). Composition, colocation, and metaontology. In D. J. Chalmers, D. Manley & R. Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.
- (2011a). Construction area (no hard hat required). *Philosophical Studies* 154(1):79–104.
- (2011b). By Our Bootstraps. *Philosophical Perspectives* 25 (1):27–41.
- (2017). *Making Things Up*. Oxford: Oxford University Press.
- (forthcoming). Why I am Not a Dualist. *Oxford Studies in Philosophy of Mind* 1.
- Black, M. (1952). The identity of indiscernibles. *Mind* 61(242):153–164.
- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in the Philosophy of Science* 9:261–325.
- (2007). Consciousness, accessibility and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences* 30: 481–548.
- Block, N. & Stalnaker, R. (1999). Conceptual analysis, dualism, and the explanatory gap. *Philosophical Review* 108 (1):1–46.
- Builes, D. (2021). Ontology and Arbitrariness. *Australasian Journal of Philosophy*. doi: [10.1080/00048402.2021.1931377](https://doi.org/10.1080/00048402.2021.1931377)
- Cameron, R. P. (2008). Truthmakers and ontological commitment: or how to deal with complex objects and mathematical ontology without getting into trouble. *Philosophical Studies* 140 (1):1–18.
- (2010). How to Have a Radically Minimal Ontology. *Philosophical Studies* 151: 249–64.
- Carruthers, P. (2000). *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- (2006). Perception and the Fall from Eden. In T. Szabo Gendler and J. Hawthorne (eds.) *Perceptual Experience*. Oxford: Oxford University Press.
- (2010). *The Character of Consciousness*. Oxford: Oxford University Press.
- (2012). *Constructing the World*. Oxford: Oxford University Press.

- Chalmers, D. J. & Jackson, F. (2001). Conceptual analysis and reductive explanation. *Philosophical Review* 110 (3):315-61.
- Chen, E. K. (2017). Our Fundamental Physical Space: An Essay on the Metaphysics of the Wave Function. *Journal of Philosophy* 114 (7): 333–65.
- (2019). Realism About the Wave Function, *Philosophy Compass* 14 (7).
- Chisholm, R. M. (1967). Identity through possible worlds: Some questions. *Noûs* 1 (1):1-8.
- Coleman, S. (2012). Mental Chemistry: Combination for Panpsychists. *Dialectica* 66 (1):137-166.
- Collins, R. (2011). A Scientific Case for the Soul. In M. C. Baker & S. Goetz (eds.), *The Soul Hypothesis: Investigations Into the Existence of the Soul*. New York: Continuum Press. pp. 222-246.
- Correia, F. (2010). Grounding and truth-functions. *Logique Et Analyse* 53(211):251–279.
- Correia, F. & Schnieder, B. (eds.) (2012). *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge: Cambridge University Press.
- Correia, F. & Skiles, A. (2017). Grounding, Essence, and Identity. *Philosophy and Phenomenological Research* (3):642–670.
- Dasgupta, S. (2009). Individuals: an essay in revisionary metaphysics. *Philosophical Studies* 145(1):35-67.
- (2014a). On the Plurality of Grounds. *Philosophers Imprint* 14(20): 1-28.
- (2014b). The Possibility of Physicalism. *The Journal of Philosophy* 111:557–92.
- (2015). Quality and Structure. In E. Barnes (ed.), *Current Controversies in Metaphysics*. Routledge.
- (2017). Constitutive Explanation. *Philosophical Issues* 27(1):74–97.
- Davidson, D. (1970). Mental Events. In D. Davidson, *Essays on Actions and Events*, Oxford: Oxford University Press, pp. 207–223.
- Dehaene, S. & Naccache, L. (2000). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79:1–37.
- deRosset, L. (2013). Grounding Explanations. *Philosophers' Imprint* 13: 1–26.
- (2017). Grounding the Unreal. *Philosophy and Phenomenological Research* 95 (3):535-563.
- Dewar, N. (2019). Algebraic structuralism. *Philosophical Studies* 176 (7):1831-1854.
- Dorr, C. (2005). What We Disagree About When We Disagree About Ontology. In M. Kalderon (ed.), *Fictionalist Approaches to Metaphysics*. Oxford: Oxford University Press.
- (2007). There Are No Abstract Objects. In T. Sider, J. Hawthorne and D. Zimmerman (eds.), *Contemporary Debates in Metaphysics*, Malden: Wiley-Blackwell.
- (2014). Quantifier Variance and the Collapse Theorems. *The Monist* 97 (4):503-570.

- (2016). To Be F Is To Be G. *Philosophical Perspectives* 30 (1):39-134.
- (unpublished). Finding Ordinary Objects in Some Quantum Worlds. Available at <http://www.nyu.edu/projects/dorr/papers/Finding.pdf>.
- Dorr, C. and G. Rosen (2002). Composition as a Fiction. In R. Gale (ed.), *The Blackwell Guide to Metaphysics*, 151-74. Oxford: Blackwell.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, Mass.: MIT Press.
- Elgin, S. Z. (forthcoming). The semantic foundations of philosophical analysis. *Review of Symbolic Logic*:1-25.
- Epstein, B. (2015). *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. Oxford: Oxford University Press.
- Fairchild, M. (2019). The Barest Flutter of the Smallest Leaf: Understanding Material Plenitude. *Philosophical Review* 128: 143–178.
- Fairchild, M. & Hawthorne, J. (2018). Against Conservatism in Metaphysics. *Royal Institute of Philosophy Supplement* 82:45-75.
- Field, H. (1994). Disquotational truth and factually defective discourse. *Philosophical Review* 103 (3):405-452.
- Fine, K. (1994). Senses of Essence. In W. Sinnott-Armstrong, D. Raffman & N. Asher (eds.), *Modality, Morality and Belief. Essays in Honor of Ruth Barcan Marcus*. Cambridge University Press.
- (2001). The question of realism. *Philosophers' Imprint* 1:1–30.
- (2005). Tense and reality. In *Modality and Tense*. Oxford: Oxford University Press.
- (2010). Towards a Theory of Part. *Journal of Philosophy* 107 (11):559-589.
- (2012). Guide to Ground. In Correia & Schnieder (2012).
- (2013). Fundamental Truth and Fundamental Terms. *Philosophy and Phenomenological Research* 87(3):725–732.
- (2015). Unified Foundations for Essence and Ground. *Journal of the American Philosophical Association* 1(2):296–311.
- (2016). Identity criteria and ground. *Philosophical Studies* 173 (1):1-19.
- (2017). A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic* 46 (6):625-674.
- Fodor, J. A. (1974). Special Sciences: Or, The Disunity of Science as a Working Hypothesis. Reprinted in J. Fodor, *Representations*, Cambridge, MA: MIT Press, 1981.
- (1990). *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Frugé, C. (ms). Janus-Faced Grounding.
- Gao, S. (2017). *The Meaning of the Wave Function: In Search of the Ontology of Quantum Mechanics*, Cambridge: Cambridge University Press.

- Ghirardi, G., A. Rimini and T. Weber (1986). Unified Dynamics for Microscopic and Macroscopic Systems. *Physical Review D* 34 (D): 470–491.
- Glazier, M. (2016). Laws and the Completeness of the Fundamental. In M. Jago (ed.), *Reality Making*. Oxford: Oxford University Press. pp. 11-37.
- (2017). Essentialist Explanation. *Philosophical Studies* 174(11):2871–2889.
- Gómez Sánchez, V. (forthcoming). From Nomic Humeanism to Normative Relativism. *Philosophical Perspectives*.
- (forthcoming). Naturalness by Law. *Noûs*.
- Goodman, J. (2017). Reality Is Not Structured. *Analysis* 77 (1): 43–53.
- Hare, C. (2009). *On Myself, and Other, Less Important Subjects*. Princeton: Princeton University Press.
- Hawthorne, J. (2002). Blocking Definitions of Materialism. *Philosophical Studies* 110 (2):103-113.
- (2010). A Metaphysician Looks at the Everett Interpretation. In S. Saunders, J. Barrett, A. Kent and D. Wallace (eds.), *Many Worlds?: Everett, Quantum Theory & Reality*. Oxford: Oxford University Press.
- (O’Leary-)Hawthorne, J. & A. Cortens (1995). Towards ontological nihilism. *Philosophical Studies* 79(2):143–65.
- Heil, J. (2003). Levels of Reality. *Ratio* 16(3): 205–21. doi:10.1111/1467-9329.00218
- Hill, C. S. (1997). Imaginability, Conceivability, Possibility, and the Mind-Body Problem. *Philosophical Studies* 87 (1): 61–85.
- Hirsch, E. (1999). The Vagueness of Identity. *Philosophical Topics* 26: 139–159.
- Ismael, J. (unpublished). What Entanglement Might Be Telling Us.
- Jenkins, C. (2011). Is Metaphysical Dependence Irreflexive? *The Monist* 94(2):267–276.
- Kennedy, C. & Stanley, J. (2009). On ‘Average’. *Mind* 118(471):583–646.
- Kim, J. (1993). Postscripts on Supervenience. *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- (1994). Explanatory Knowledge and Metaphysical Dependence. *Philosophical Issues* 5:51–69.
- (1998). *Mind in a Physical World*, Cambridge: Cambridge University Press.
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science* 48: 507–531.
- Kment, B. (2014). *Modality and Explanatory Reasoning*. Oxford: Oxford University Press.
- Korman, D. Z. (2010). Strange Kinds, Familiar Kinds, and the Charge of Arbitrariness. *Oxford Studies in Metaphysics*, 5: 119–144.
- Krämer, S. & Roski, S. (2015). A Note on the Logic of Worldly Ground. *Thought: A Journal of Philosophy* 4(1):59–68.

- Kriegel, U. (2015). Existence, Fundamentality, and the Scope of Ontology. *Argumenta* 1 (1):97-109.
- Kripke, S. A. (1980). *Naming and Necessity: Lectures Given to the Princeton University Philosophy Colloquium*. Cambridge, MA: Harvard University Press.
- Lamme, V. A. F. (2006). Toward a true neural stance on consciousness. *Trends in Cognitive Science* 10:11, 494–501.
- Latham, N. (2000). Chalmers on the addition of consciousness to the physical world. *Philosophical Studies* 98 (1):71-97.
- Lee, G. (2018). Alien Subjectivity and the Importance of Consciousness. In A. Pautz and D. Stoljar (eds.), *Blockheads: Essays on Ned Block's Philosophy of Mind and Consciousness*. Cambridge, MA: MIT Press.
- Levine, J. (1983). Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64: 354–61.
- (2006). Phenomenal concepts and the materialist constraint. In Torin Alter & Sven Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press.
- Lewis, D. K. (1980). Mad pain and Martian pain. In N. Block (ed.), *Readings in the Philosophy of Psychology*. Harvard University Press. pp. 216-222.
- (1983a). New work for a theory of universals. *Australasian Journal of Philosophy* 61(4):343–377.
- (1983b). Counterpart Theory and Quantified Modal Logic. Reprinted with additional “Postscripts” in *Philosophical Papers, vol. 1*. Oxford: Oxford University Press.
- (1986a). *On the Plurality of Worlds*. Malden: Blackwell.
- (1986b). *Philosophical Papers, vol. 2*. Oxford: Oxford University Press.
- (1991). *Parts of Classes*. Oxford: Blackwell.
- (1994a). Humean Supervenience Debugged. *Mind* 103:473–490.
- (1994b). Reduction of Mind. In S. Guttenplan (ed), *A Companion to the Philosophy of Mind*, Oxford: Blackwell, pp. 412–431.
- (1999). Many but Almost One. In *Papers in Metaphysics and Epistemology: Volume 2*. Cambridge: Cambridge University Press.
- Lewis, P. J. (2004). Life in Configuration Space. *British Journal for the Philosophy of Science* 55 (4): 713–729.
- (2013). Dimension and Illusion. In A. Ney and D. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.
- (2016). *Quantum Ontology: A Guide to the Metaphysics of Quantum Mechanics*. New York, NY: Oxford University Press.
- Lipman, M. A. (2016). Perspectival Variance and Worldly Fragmentation. *Australasian Journal of Philosophy* 94 (1):42-57.

- (2018). Against Fundamentality-Based Metaphysics. *Noûs* 52(3):587–610.
- Loar, B. (1990). Phenomenal States. *Philosophical Perspectives* 4: 81–108.
- Loewer, B. (1996). Humean Supervenience. *Philosophical Topics* 24(1):101-127.
- Lycan, W. (1996). *Consciousness and Experience*. Cambridge, MA: MIT Press.
- Massimi, M. (2017). Laws of Nature, Natural Properties, and the Robustly Best System. *The Monist* 100 (3):406-421.
- Maudlin, T. (2007). Completeness, Supervenience and Ontology. *J. Phys. A: Math. Theor.* (40): pp. 3151–71.
- (2010). Can the World Be Only Wavefunction? In S. Saunders, J. Barrett, A. Kent & D. Wallace (eds.), *Many Worlds?: Everett, Quantum Theory & Reality*, Oxford: Oxford University Press.
- (2013). ‘The Nature of the Quantum State’, in A. Ney & D. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. Oxford: Oxford University Press.
- (2019). *Philosophy of Physics vol. 2: Quantum Theory*. Princeton: Princeton University Press.
- McGee, V. & McLaughlin, B. P. (2000). The Lessons of the Many. *Philosophical Topics* 28:129–152.
- McKenzie, K. (2014). Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis. *British Journal for the Philosophy of Science* 65 (2):353–380.
- Mehta, N. (2019). Phenomenal, Normative, and Other Explanatory Gaps: A General Diagnosis. *Philosophy and Phenomenological Research* 98 (3):567-591.
- Monton, B. (2002). Wave Function Ontology. *Synthese* 130 (2), pp. 265–77.
- (2006). Quantum Mechanics and 3N-Dimensional Space. *Philosophy of Science* 73, pp. 778–89.
- Moore, G. E. (1925). A Defence of Common Sense. In J. H. Muirhead (ed.) *Contemporary British Philosophy* (2nd series). London: Allen and Unwin.
- Murphy, D. (ms). Symmetry-breaking Determination.
- Nagel, T. (1974). What is it Like to Be a Bat? *The Philosophical Review* 83: 435–50.
- Ney, A. (2012). The Status of Our Ordinary Three Dimensions in a Quantum Universe. *Noûs* 46 (3): 525–60.
- (2013). Ontological Reduction and the Wave Function Ontology. In A. Ney & D. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.
- (2015). Fundamental Physical Ontologies and the Constraint of Empirical Coherence: a Defense of Wave Function Realism. *Synthese* 192(10): 3105–24.

- North, J. (2013). The Structure of the Quantum World. In A. Ney & D. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.
- Papineau, D. (2016). Against Representationalism (about Conscious Sensory Experience). *International Journal of Philosophical Studies* 24: 324–347.
- Pautz, A. (2017). The significance argument for the irreducibility of consciousness. *Philosophical Perspectives* 31 (1):349-407.
- (forthcoming). How to Achieve the Physicalist Dream Theory of Consciousness: Identity or Grounding?. In G. O. Rabin (ed.), *Grounding and Consciousness*.
- Pereboom, D. (2011). *Consciousness and the Prospects of Physicalism*. Oxford: Oxford University Press.
- Place, U. T. (1956). Is consciousness a brain process? *British Journal of Psychology* 47 (1):44-50.
- Putnam, H. (1967). Psychological Predicates. In W.H. Capitan and D.D. Merrill (eds.), *Art, Mind, and Religion*, Pittsburgh: University of Pittsburgh Press, 37–48.
- Rabin, G. O. (2019). Fundamentality physicalism. *Inquiry: An Interdisciplinary Journal of Philosophy* (1):77-116.
- Roelofs, L. (2019). *Combining Minds: How to Think about Composite Subjectivity*. New York: Oxford University Press.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In B. Hale & A. Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press.
- (2011). The Reality of Mathematical Objects. In J. Polkinghorne (ed.), *Meaning in Mathematics*. Oxford: Oxford University Press.
- (2017). Ground by Law. *Philosophical Issues* 27(1):279–301.
- Rosenthal, D. M. (1993). Thinking that one thinks. In M. Davies & G. Humphreys, eds. *Consciousness: Psychological and Philosophical Essays*. Oxford: Blackwell.
- (2002). Explaining Consciousness. In D. J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press. pp. 109-131.
- Russell, B. (1912). *The Problems of Philosophy*. London: Williams and Norgate.
- (1918). The Philosophy of Logical Atomism. In R. C. Marsh (ed.), *Bertrand Russell: Logic and Knowledge, Essays 1901–1950*, London: Unwin Hyman.
- (2003). Analytic Realism. In S. Mumford (ed.), *Russell on Metaphysics*. London: Routledge.
- Russell, J. S. (2015). Possible Worlds and the Objective World. *Philosophy and Phenomenological Research* 90(2):389–422.
- (2016). Qualitative Grounds. *Philosophical Perspectives* 30(1):309–348.
- (2018). Quality and Quantifiers. *Australasian Journal of Philosophy* 96(3):562–577.

- Saunders, S., J. Barrett, A. Kent & D. Wallace (eds.) (2010). *Many Worlds?: Everett, Quantum Theory & Reality*, Oxford: Oxford University Press.
- Schaffer, J. (2004). Two conceptions of sparse properties. *Pacific Philosophical Quarterly* 85:92–102.
- (2008). Truthmaker commitments. *Philosophical Studies* 141 (1):7–19.
- (2009). On what grounds what. In D. Manley, D. J. Chalmers & R. Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.
- (2012). Grounding, transitivity, and contrastivity. In Correia & Schnieder (2012).
- (2013). Metaphysical Semantics Meets Multiple Realizability. *Analysis* 73(4):736–751.
- (2015). What Not to Multiply Without Necessity. *Australasian Journal of Philosophy* 93(4):644–664.
- (2016). Grounding in the image of causation. *Philosophical Studies* 173(1):49–100.
- (2017a). The Ground Between the Gaps. *Philosophers' Imprint* 17(11): 1–26.
- (2017b). Laws for Metaphysical Explanation. *Philosophical Issues* 27(1):302–321.
- (2021). Ground Functionalism. *Oxford Studies in the Philosophy of Mind*.
- (forthcoming). Naturalistic Dualism and the Problem of the Physical Correlate. In G. O. Rabin (ed.), *Grounding and Consciousness*.
- Shoemaker, S. (2007). *Physical Realization*, Oxford: Oxford University Press.
- Sider, T. (2001). *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
- (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- (2013a). Symposium on Writing the Book of the World. *Analysis* 73(4):751–770.
- (2013b). Replies to Dorr, Fine, and Hirsch. *Philosophy and Phenomenological Research* 87(3):733–754.
- (2013c). Against Parthood. *Oxford Studies in Metaphysics* 8:237–293.
- (2020). *The Tools of Metaphysics and the Metaphysics of Science*. Oxford: Oxford University Press.
- Skiles, A. (2015). Against Grounding Necessitarianism. *Erkenntnis* 80 (4):717–751.
- Skow, B. (2016). *Reasons Why*. Oxford: Oxford University Press.
- Spencer, J. (2016). Relativity and Degrees of Relationality. *Philosophy and Phenomenological Research* 92 (2):432–459.
- Stoljar, D. (2021). Physicalism. *Stanford Encyclopedia of Philosophy*.
- Thomasson, A. (2007). *Ordinary Objects*, Oxford: Oxford University Press.
- Turner, J. (2010). Ontological Pluralism. *Journal of Philosophy* 107(1):5–34.

- (2011). Ontological Nihilism. *Oxford Studies in Metaphysics* 6:3–54.
- (2017). Can We Do without Fundamental Individuals? In E. Barnes (ed.), *Current Controversies in Metaphysics*. Routledge.
- (forthcoming). On Doing Without Ontology: Feature-Placing on a Global Scale. In J. Cumpa (ed.), *The Question of Ontology*. Oxford University Press.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Mass.: MIT Press.
- (1999). Phenomenal consciousness: The explanatory gap as a cognitive illusion. *Mind* 108 (432):705-25.
- Van Inwagen, P. (1990). *Material Beings*. Ithaca: Cornell University Press.
- Williams, J. R. G. (2012). Requirements on Reality. In Correia & Schnieder (2012).
- Wilsch, T. (2016). The Deductive-Nomological Account of Metaphysical Explanation. *Australasian Journal of Philosophy* 94 (1):1-23.
- Wilson, A. (2018). Metaphysical Causation. *Noûs* 52 (4):723-751.
- Wilson, J. (2009). Determination, Realization, and Mental Causation. *Philosophical Studies* 145(1): 149–169. doi:10.1007/s11098-009-9384-8
- (2011). Non-reductive Realization and the Powers-based sub-set strategy. *The Monist* 94: 121–154.
- (2013). A determinable-based account of metaphysical indeterminacy. *Inquiry: An Interdisciplinary Journal of Philosophy* 56 (4):359-385
- (2014). No Work for a Theory of Grounding. *Inquiry* 57:5-6. 535–79.
- Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. Trans. C. K. Ogden. London: Routledge and Kegan Paul Ltd., 1981.
- Yablo, S. (1992). Mental Causation. *Philosophical Review* 101:245–80.