Grounded Shadows, Groundless Ghosts

E. Rubenstein

ABSTRACT

According to a radical account of quantum metaphysics which I label 'high-dimensionalism', ordinary objects are the 'shadows' of high-dimensional fundamental ontology (for example, (Albert [2013], [unpublished-b]; Ney [2015])). Critics – especially Maudlin ([2007], [2010], [2019]) – allege that high-dimensionalism cannot provide a satisfactory explanation of the manifest image. In this paper, I examine the two main ideas behind these criticisms: that high-dimensionalist connections between fundamental and non-fundamental are 1) inscrutable, and 2) arbitrary. In response to the first, I argue that there is no metaphysically significant contrast regarding the scrutability of low- and high-dimensionalist connections. In response to the second, I argue that the arbitrariness of high-dimensionalist connections has been overstated, and what arbitrariness there is afflicts low-dimensionalist connections too. Thus, the debate should not be focused on whether high-dimensionalism can provide a satisfactory explanation of the manifest image – as it has been in recent literature – but rather on the broader question of whether there is good all-things-considered reason to prefer low-dimensionalist theories.

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 It is what David Albert ([unpublished-a], p.7) describes as 'the totality of opportunities for things to be one way or another'.[1] Specifying the fundamental facts is a matter of specifying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According to Albert ([1996], [2013], [2015], [unpublished-b]) and Alyssa Ney ([2012], [2013], [2015]), however, the success of quantum mechanics suggests a radical rejection of low-dimensionalism.[2] They propose that the fundamental arena corresponds to what is, for the low-dimensionalist, the universe's 'configuration space' (together with a temporal dimension). This arena has $3N+1$ dimensions, where N is the number of particles.[3] 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.)[4]

High-dimensionalism faces an obvious challenge: accounting for the low-dimensional world of our everyday experience. 'The particularly urgent question', as Albert ([2013], p.54) puts it, 'is where, in this picture, all the tables, and chairs, and buildings, and people are'. As Chen ([2019], p.6) emphasizes, at stake in this question is not just high-dimensionalist theories' ability to accommodate our common-sense conception of the

[1] 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').
[2] See also (Loewer [1996]; North [2013]; Ismael [unpublished]).
[3] This is merely a heuristic: the number of dimensions is supposed to determine the number of particles, not *vice versa*.
[4] There is also a hybrid view which envisions both a low-dimensional and a high-dimensional fundamental arena – see (Dorr [unpublished]).

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-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.[5] 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.[6] 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'[7] – fundamental spatiotemporal entities such as particles or fields – determines ordinary truths involving macroscopic objects like tables. As Allori ([2013-a], p.69) puts it:

> [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 §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

[5] See (Sider [2011], §7.6) on 'toy' metaphysical truth-conditions.
[6] See also (Allori [2013-a], [2013-b]; Chen [2017]; Hawthorne [2010]; Monton [2002], [2006]).
[7] For this terminology, see (Allori [2013-a]).

ontology to the manifest image. However, critics have deemed this approach unsatisfactory. Hawthorne ([2010], p.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], p.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], p.351) echoes these concerns, concluding that 'we are right to doubt whether there can be any principled way to close the apparent explanatory gap'.[8]

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-

[8] Hawthorne ([2010], p.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.

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 §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 §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.[9] 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.

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

[9] See, for example, (Lewis [2004]; Allori [2013-b]; Maudlin [2013]; Chen [2017]; Gao [2017]).

location.[10] Thus, they support metaphysical explanations of the manifest image in terms of the fundamentals.

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], p.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.[11]

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

[10] 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.

[11] 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.)

'grounding connections', and I will talk of fundamental ontology/facts grounding derivative ontology/facts.)[12] 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.[13] 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.[14]

This is worth emphasizing because much of the debate over high-dimensionalism has followed Albert's ([1996], p.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], p.784) objects that high-dimensionalism is 'even more radical than the brain-in-the-vat scenario', Maudlin ([2007], p.3166) objects that 'we cannot appeal to mere fictions' to explain our empirical evidence, and Allori ([2013-b], §7.6) objects that high-dimensionalism – since it makes essential appeal to experience – requires a solution to the mind-body problem. Understanding high-dimensionalism as the claim that spacetime and its occupants are grounded in high-dimensional fundamentals avoids such objections; high-dimensionalism is compatible with the manifest image being just as real as it is for the low-dimensionalist.[15]

The nature of the required grounding connections is central to the debate between high- and low-dimensionalists. The key question is: by what criteria ought we evaluate this part

[12] 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]).

[13] 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.*

[14] 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 (§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.

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

of theorizing, and do high-dimensionalist connections do worse than low-dimensionalist 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-dimensionalist versions of quantum theories replicate the predictions of their low-dimensionalist 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-dimensionalism for thinking that high-dimensionalist connections do not provide satisfactory explanations of the manifest image. But first, I should briefly present these connections.

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.[16]

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 $\langle x_1, x_2, x_3, t\rangle, \ldots, \langle x_{3i-2}, x_{3i-1}, x_{3i}, t\rangle, \ldots, \langle x_{3N-2}, x_{3N-1}, x_{3N}, t\rangle$ is coordinatized as $\langle x_1, \ldots, x_{3N}, t\rangle$.[17]

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 (…

[16] This becomes important in the discussion of arbitrariness below (especially §5.2).

[17] 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.

x_{3i-2} = a, x_{3i-1} = b, x_{3i} = c,…, 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])).[18] 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], p.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.

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.

[18] 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)).

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:

[table at spacetime region R]

[local beables arranged table-wise at spacetime region R]

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

[local beables arranged table-wise at spacetime region R]

[high-dimensional ontology arranged table-wise* at region of fundamental arena $G(R)$]

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 precifisation 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.

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], p.65) alleges that 'once the primitive ontology and its temporal evolution are given, everything else follows', and Maudlin ([2007], p.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.

4.1 Triviality

Maudlin ([2010]) argues that the emergence of pointer positions from Bohmian particle configurations is easy to understand. He ([2010], p.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], p.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], p.20) or 'can be appropriately "read off" from' ([2013-b], p.15) the history of the primitive ontology, allowing us to 'directly compare' ([2013-a], p.66) its macroscopic behaviour to the manifest image. In Maudlin's ([2007], p.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.[19] 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.[20]

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 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

[19] As Maudlin ([2010], p.124) acknowledges, merely having the configurations is insufficient: 'we need at least some schematic reason to think that we could *see* these configurations'.
[20] Thanks to a referee for pointing this out.

us as referring to tiny bits of tables and chairs.[21] 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.

Using such discrepancies in our access to the relevant concepts as a basis for fundamental theorizing seems unduly anthropocentric.[22] 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?

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 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-

[21] I denote concepts in *CAPS*.

[22] 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.

dimensionalism: ordinary ontology could not be grounded in high-dimensional ontology because such connections are 'opaque' (non-transparent).

In fact, following Schaffer ([2017], p.4), 'three not-obviously-equivalent' ways of unpacking opacity are familiar from the philosophy of mind literature, involving conceivability, logical possibility, and epistemic a priority. 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.[23]

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, 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

[23] See Chalmers ([2012]) on 'a priori scrutability'.

familiar and apparently unproblematic cases like mereology. These gaps are bridged by substantive, opaque connections.

As Schaffer ([2017], p.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.[24] 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.

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*.

[24] For example, Thomasson ([2007]). See (Sider [2011], §§9.7-9.10) for discussion.

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 (§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 fn22). 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).[25]

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], p.3166) characterizes these connections as a 'choice of one out of an infinitude' of possible alternatives, and Allori ([2013-b], p.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."[26]

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

---

[25] 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.)

[26] For worries about arbitrariness, see also (Gao [2017]; Hawthorne [2010]; Lewis [2004]; Maudlin [2010], [2019]; Monton [2002], [2006]).

[27] I borrow this useful term from Albert ([2015], p.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], p.135) objections. As Albert ([2015], pp.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.

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.

5.1 Other-dimensional ghosts

Recall (§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?[28]

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.[29]

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'.[30] 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 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

[28] Hawthorne ([2010], pp.152-3) raises this issue in the context of Many Worlds.

[29] Although, 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 (§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.)

[30] 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], p.282), an 'emergent geometrical space' ([unpublished-a]), and 'the space of ordinary material bodies' ([unpublished-b]). Similarly, Lewis ([2013], p.123) refers to 'the arena in which *spatial* phenomena play out', noting that 'the term *spatial* is intimately connected to the dynamical laws'.

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], fn14) 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').[31] 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.[32] The privilege in question, recall, is that of being the *material objects*. And

[31] 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.

[32] This is contrary to the apparently widespread view that, as Allori ([2013-b], p.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], p.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.

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.[33]

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.[34]

5.2 Displaced ghosts

Maudlin ([2019], p.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 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.

[33] For vivid thought-experiments along these lines see (Albert [unpublished-b]).
[34] See (Ismael [unpublished]; Albert [unpublished-b]).

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.[35] 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-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 (§3). The region $G(p)$ grounds the spacetime point p; the topological and geometric relations holding between spacetime points $p_1,\ldots, p_n$ are determined by the relations between their corresponding regions $G(p_1),\ldots, G(p_n)$. According to high-dimensionalism, then, spacetime points themselves are derivative entities grounded in certain regions of the fundamental arena – namely, corresponding sets of hyperplanes.

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

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:

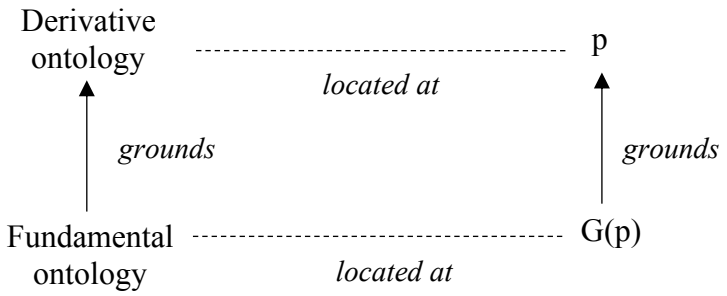

The straightforward thing to say, that is, about how local beables inherit their location from the fundamental ontology which grounds them is just that their locations are grounded in the locations of that fundamental ontology. For example, the undulations of the wavefunction across hyperplanes of the 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 <a, b, c, t>. 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.[36] 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

[36] See §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.

5.3 Alternative local beable ghosts

Maudlin ([2007], pp.3161-2; [2019], p.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.[37] 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.[38] So not only do these alternative local beables both enact material objects, but they enact objects which behave identically in all circumstances.

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

[37] See (Bell [1987], ch.22).
[38] See (Albert [2015], ch.4) for an illuminating discussion of experimental distinguishability.

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.

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 (§4), and secondly that there is no asymmetry in arbitrariness – high- and low-dimensionalism face the same kinds of choices between alternative connections (§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]. p.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], p.63), and the manifest image ought to be 'implicit' within it, falling out more or less straightforwardly 'as a purely *analytical* consequence' (Maudlin [2007], p.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.[39] 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.[40]

'Accommodate' is deliberately vague: they must either explain the truth of these platitudes, or at least their appearance-as-if-true.[41] 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.

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

[39] On this point, see (North [2013], p.200).
[40] 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], p.784)).
[41] Ney ([2013], p.173) makes a similar distinction between 'eliminative' and 'retentive' reductions.

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?[42]

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.

Acknowledgements

Thanks to David Albert, Verónica Gómez Sánchez, Barry Loewer, Jill North, Jonathan Schaffer, Ted Sider, Isaac Wilhelm, and the participants of the Rutgers Third Year Seminar for discussion which has greatly benefitted this paper. Thanks also to three anonymous referees for this journal for their helpful comments and suggestions.

[42] Schaffer ([2017], p.14) makes a similar comparison.

Ezra Rubenstein
Rutgers University
New Brunswick, NJ, USA
ezra.rubenstein@rutgers.edu

References

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. 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. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.

—— [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/.

Bell, J. S. [1987]: *Speakable and Unspeakable in Quantum Mechanics*, Cambridge, UK: Cambridge University Press.

Chalmers, D. J. [1996]: *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.

—— [2012]: *Constructing the World*, Oxford: Oxford University Press.

Chen, E. K. [2017]: 'Our Fundamental Physical Space: An Essay on the Metaphysics of the Wave Function', *Journal of Philosophy* **114 (7)**, pp. 333–65.

—— [2019]: Realism About the Wave Function, *Philosophy Compass* **14 (7)**.

Dorr, C. [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*, Oxford: Blackwell.

Fine, K. [2012]: ‘Guide to Ground’, in F. Correia and B. Schnieder (eds.), *Metaphysical Grounding*, Cambridge: Cambridge University Press.

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* **D34 (D)**, pp. 470–491.

Hawthorne, J. [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.

Hill, C. S. [1997]: ‘Imaginability, Conceivability, Possibility, and the Mind-Body Problem’, *Philosophical Studies* **87 (1)**, pp. 61–85.

Ismael, J. [unpublished]: ‘What Entanglement Might Be Telling Us’.

Lewis, P. J. [2004]: ‘Life in Configuration Space’, *British Journal for the Philosophy of Science* **55 (4)**, pp. 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.

Levine, J. [1983]: ‘Materialism and Qualia: The Explanatory Gap’, *Pacific Philosophical Quarterly* **64**, pp. 354–61.

Loar, B. [1990]: ‘Phenomenal States’, *Philosophical Perspectives* **4**, pp. 81–108.

Loewer, B. [1996]: ‘Humean Supervenience’, *Philosophical Topics* **24 (1)**, pp. 101–127.

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 and D. Wallace (eds.), *Many Worlds?: Everett, Quantum Theory & Reality*, Oxford: Oxford University Press.

—— [2013]: ‘The Nature of the Quantum State’, in A. Ney and 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.

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.

Ney, A. [2012]: ‘The Status of Our Ordinary Three Dimensions in a Quantum Universe’, *Noûs* **46 (3)**, pp. 525–60.

—— [2013]: ‘Ontological Reduction and the Wave Function Ontology’, in A. Ney and 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)**, pp.3105–24.

North, J. [2013]: ‘The Structure of the Quantum World’, in A. Ney and D. Albert (eds.), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Oxford: Oxford University Press.

Rosen, G. [2010]: ‘Metaphysical Dependence: Grounding and Reduction’, in B. Hale and A. Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*, Oxford: Oxford University Press.

Saunders, S., J. Barrett, A. Kent and D. Wallace (eds.) [2010]: *Many Worlds?: Everett, Quantum Theory & Reality*, Oxford: Oxford University Press.

Schaffer, J. [2017]: ‘The Ground Between the Gaps’, *Philosophers' Imprint* **17(11)**, pp.1–26.

Sider, T. [2011]: *Writing the Book of the World*, Oxford: Oxford University Press.

—— [2013]: ‘Against Parthood’, *Oxford Studies in Metaphysics* **8**, pp. 237–93.

Thomasson, A. [2007]: *Ordinary Objects*, Oxford: Oxford University Press.

Wilson, J. [2014]: ‘No Work for a Theory of Grounding’, *Inquiry* **57:5-6**, pp. 535–79.