Dynamical substantivalism

Tushar Menon

Dianoia Institute of Philosophy,

Australian Catholic University

Abstract

This paper outlines and motivates a novel view in the metaphysics of spacetime: dynamical substantivalism. Dynamical substantivalism results from treating as orthogonal two often conflated questions. First, the question of whether there are immaterial spacetime points. Second, the question of what grounds spatiotemporal structure. Dynamical substantivalism sides with the substantivalist on the former question, and with the dynamical theorist on the latter. This paper (i) offers a diagnosis of the widespread conflation of these two questions and (ii) argues that dynamical substantivalism should be treated as a serious contender in the metaphysics of space and time.

1 Introduction

The project of this paper is to outline and motivate a novel view in the metaphysics of spacetime: dynamical substantivalism. Dynamical substantivalism provides a joint answer to two separate questions. The first is about what sorts of things there are:

`Ontological question`: Are there immaterial (i.e. materially unoccupied) spacetime points?

Substantivalism says there are; relationalism says there are not. Dynamical substantivalism sides with the substantivalists.

The second question is about explanation:

`Explanatory question`: Does spacetime structure ground a law's dynamical symmetry structure or vice versa? (Equivalently: in virtue of what do spacetime points, whether materially occupied or not, stand in the spatiotemporal relations that they do?)

Roughly speaking, according to the dynamical approach, spacetime points stand in the relations that they do in virtue of the primitive symmetry properties of the laws that govern the dynamics of matter (I spell out what precisely this means in §1.1) ; according to the proponent of the geometrical approach, the symmetry properties of the laws that govern the dynamics of matter are what they are in virtue of the primitive relations in which spacetime points stand. Dynamical substantivalism sides with the proponents of the dynamical approach.

Dynamical substantivalism is a novel view whose visibility depends on our treating the ontological and the explanatory question as orthogonal. Historically, they have tended not to be so treated; rather, substantivalism, it has generally been assumed, entails a geometrical answer to the explanatory question (see e.g. [1, 10]). And the dynamical approach has been treated as a form of relationalism (see [1, 27]).[1] This, I argue, should be resisted: substantivalists can go dynamical and geometrical theorists can go relational. Logical space is more variegated and modular than has previously been assumed. The shape of the extant literature masks key choice-points for the metaphysician of spacetime. The matrix of possible views has four acceptable entries, rather than just two; call this the *orthogonality thesis*. I will argue for this thesis in §2, and situate the dispute within the broader discussion of scientific realism. In §3, I discuss what it would take to answer the ontological question. In §4, I engage in a similar discussion over the explanatory question. In that section, I also demonstrate that the recarving of logical space for which I argue is dialectically consequential. In particular, the dynamical version of substantivalism that I discuss in §5 has various dialectical advantages over its traditional geometrical counterpart. Hence, the recarving of logical space is no mere redescription of familiar territory. It substantially changes the shape of the dialectic.

1.1 Dynamical and spacetime symmetries

In this short subsection, I define dynamical and spacetime symmetries. Readers who are familiar with these concepts can safely skip this subsection, and pick up the main thread of the argument in §2.

Consider the following claim: (the extensions of) p and q are spacetime points that are five units of distance apart, and all distances in the theory behave in a manner consistent with the axioms of Euclidean geometry. The standard way to codify geometric facts about a space like this, is via

[1] There are exceptions, of course. Brown [5], Janssen [18] and Read [30] have all, at various points, and to different ends, argued that the two questions should be treated as separate. Friedman [15] discusses two positions, that he dubs 'Leibnizian' and 'Reichenbachian' relationalism, which can be understood to provide answers to the ontological and explanatory questions, respectively. However none of these authors have diagnosed the reasons for the questions having been traditionally conflated, or argued for the viability of a specific view made available by prising apart the two questions, both of which I offer in this paper.

set-theoretic tuples of the form $\langle \mathscr{M}, \mathcal{A} \rangle$, where $\mathscr{M}$ is a smooth manifold (which is just a domain of points structured in a way as to be able to differentiate functions on $\mathscr{M}$ arbitrarily many times) and $\mathcal{A}$ is a geometric object: a mathematical object that codifies the geometric relations between manifold points. In the case of our simple Euclidean geometry, $\mathcal{A}$ is a Euclidean metric tensor. Consider a point r that is 5 units from both p and q. From the fact that the geometric object is a Euclidean metric tensor, we can infer that the three points form an equilateral triangle in which each angle is exactly 60 degrees.

Define a function $\phi : \mathscr{M} \to \mathscr{M}$ such that $p' = \phi(p)$, $q' = \phi(q)$ and $r' = \phi(r)$ respectively. If the triangle $p'q'r'$ is still equilateral with side length 5 and angle 60 degrees, that map is known as a *symmetry* of the triangle structure. If ϕ has the effect of preserving *all* the geometrical relations between all points before and after the mapping, then ϕ is known as a *geometric symmetry*. We can use a function like ϕ, known more generally as a *diffeomorphism* (i.e. an isomorphism of smooth manifolds, whose inverse is also smooth) to construct a new model of the geometric space by dragging along[2] the absolute geometric objects to their 'new' locations (this action is represented as $\phi * \mathcal{A}$). If the model $\langle \mathscr{M}, \mathcal{A} \rangle = \langle \mathscr{M}, \phi * \mathcal{A} \rangle$, then ϕ is a geometric symmetry. And if the model is a model of a *spacetime* geometry, then ϕ is known as a *spacetime symmetry*.

Consider a distinct claim: (the extensions of) p and q are *measured* to be 5 units apart using good measuring devices like rulers (or clocks, if we are measuring spacetime geometries). Rulers and clocks are, of course, material entities subject to dynamical laws and constraints. In order for a ruler to qualify as a good measuring device, intuitively speaking, it should be the case that, if it correctly measures that p and q are 5 units apart, then it should also measure $p' = \phi(p)$ and $q = \phi(q)'$ as being 5 units apart, if ϕ is a spacetime symmetry. But p' and q' name different locations from p and q. So the requirement of veridical measurement translates to a constraint on the dynamical laws: the dynamical laws should not change under a geometric symmetry, since whatever aspect of the dynamics led to the correct measurement before the transformation must be preserved under the transformation. Let P be a mathematical object that represents *dynamical objects* like rulers by associating (tensorial) properties with points in M. A model of a spacetime theory including matter is now of the form $\langle \mathscr{M}, \mathcal{A}, P \rangle$. Define a *dynamical symmetry* as a diffeomorphism ϕ such that if $\langle \mathscr{M}, \mathcal{A}, P \rangle$ is a model of the theory, then so too is $\langle \mathscr{M}, \mathcal{A}, \phi * P \rangle$.[3]

[2] I use the term 'drag along' as a shorthand for the mathematical operations of pullback and pushforward of tensors along diffeomorphisms. For definitions of all these terms, see, e.g. [17].

[3] It can be shown that this model-theoretic characterisation of dynamical symmetries is equivalent to the more common characterisation in terms of satisfaction of a relativity principle. For details, see [13, Ch. 3]

2 Scientific realism and spacetime metaphysics

The ontological and explanatory questions seem to be quite straightforwardly orthogonal. So it is helpful to start by understanding why the two questions might nonetheless have appeared resistant to the kind of decoupling that I propose. To that end, let us contrast the explanatory question with a distinct but nearby question. Call it the abductive question:

`Abductive question`: Are immaterial spacetime points explanatorily indispensable according to our best scientific theory of spacetime?

Suppose you're a scientific realist. Your ontological commitments are just those tacit in our best scientific theories. This makes some ontological decisions easy. For example, it is easy to decide whether to be a realist about things like electrons (you should) and unicorns (you shouldn't). But not all decisions are this straightforward. Consider the following decision: should you be a realist about spacetime points that are not materially occupied? It is not immediately obvious what the scientific realist should say.

Part of what makes the question tricky is the difficulty in pinning down a sufficiently precise formulation of the intuitively plausible norm of scientific realism described above. Different traditions in the philosophy of science have pinned down the core ontological commitment of scientific realism in different ways. We can trace two traditions here. The first, exemplified by Boyd [4], is what we might call *abductive realism*:

`Abductive realism`: Be ontologically committed to x if and only if x is explanatorily indispensable with respect to our best scientific theory or theories.

The second, which follows a broadly Quinean impulse, is what we might call *austere realism*:

`Austere realism`: Be ontologically committed to x if and only if x appears in the domain of quantification of a model of our best scientific theory or theories.

Both of these ways of cashing out scientific realism look respectable. And both have typically been understood in such a way as to rule out dynamical substantivalism. Earman [13] can be read as the paradigm austere realist about spacetime; North [26], Dasgupta [10] and Rosen [32] can be read as abductive realists about spacetime. Let us examine these positions in turn.

The abductive realist eschews the purely model-theoretic construal of scientific realism on the grounds that it drains scientific realism of any real content. Instead of talking about models, says the

abductive realist, we should be talking about the world itself. Rather than testing our ontological commitments indirectly by testing them against the commitments of a model, we should be testing our ontological commitments directly by testing whether they are indispensable with respect to explaining various features of the world itself.

For the abductive realist, not only is the abductive question intelligible, but it is also deeply entangled with the ontological question: the question of whether substantivalism is true, that is, of whether immaterial spacetime points exist, is just the question of whether they are indispensable with respect to some key explanatory work. While the precise nature of the requisite explanatory work varies somewhat between commentators, one fairly constant theme is the explanation of the dynamical symmetry properties of the laws.[4] Therefore, for the abductive realist, the abductive question is also deeply entangled with the explanatory question: an answer to the abductive question entails an answer to the explanatory question. But the explanatory question is just the question of whether or not we should be geometrical theorists. After all, if we need immaterial spacetime points to explain dynamical symmetries, then it is because without such points, the geometrical structure of spacetime cannot explain the dynamical symmetries of the laws.

To this, the austere realist has the following rejoinder: we cannot make sense of what the world is like according to some theory of physics without first talking about how we talk about the objects. In particular, we need to be clear about how to establish what Lewis calls the 'semantic glue' [22] between our words and entities in the world. To think otherwise—to think that we can simply bypass semantics—is to think wishfully. Hence the austere realist's suggestion that we use the semantic machinery of model-theory, rather than the metaphysically suspect or obscure ideology of explanation, to tease out a theory's ontological commitments.

For austere realists, the ideology of explanation is perhaps best avoided *tout court*. The concept of explanation lacks respectability, and cannot bear the inferential burden placed upon it. For the austere realist, the dispute between the proponent of the geometrical approach and of the dynamical approach is no more respectable than a dispute between Hindu and Babylonian astrologers. And since the austere realist's ontological view does not rely on explanation, to the extent that they can make sense of the ontological question at all, they can do so independently of both the explanatory and abductive questions.

Interestingly, although they disagree over the extent to which the ontological and explanatory questions are unrelated, both the austere and the abductive realist are naturally inclined to carve out logical space into two camps, rendering invisible positions like dynamical substantivalism

[4] For example, [2, 15, 26, 34].

(and geometrical relationalism). The austere realist denies that there is an intelligible explanatory question, while the abductive realist accepts its intelligibility but denies the orthogonality thesis. The austere realist's position, at least when contrasted with the abductive realist's, consists of two claims, the first a positive semantic one, the second a negative metaphysical one. The semantic claim is simply that model-theory must mediate the link between words and worlds. The metaphysical claim is that explanation plays no role in adjudicating ontological commitments. The abductive realist denies both of these claims.

I propose, as an alternative, a hybrid view I call luxury realism. Luxury realism accepts the orthogonality thesis. It adopts a model-theoretic construal of scientific realism, which reads its ontological commitments directly off those models, but is nonetheless friendly to the ideology of explanation. Explanation plays no role, for the luxury realist, in teasing out ontological commitments, but is still a useful tool that can help articulate a sophisticated and nuanced metaphysical position. Luxury realism, in particular, accepts the intelligibility and importance of the explanatory question, but divorces it from the abductive question. Luxury realism is a hybrid position insofar as it sides with the austere realist in the latter's construal of the ontological question, and sides with the abductive realist in accepting the intelligibility of the explanatory question. But in severing the links between the abductive and explanatory questions, on the one hand, and the abductive and ontological questions, on the other, the luxury realist can maintain the orthogonality of the ontological and the explanatory questions. The reason that dynamical substantivalism has been largely invisible as a option is that most commentators can be read as either austere realists (e.g. [13]) or abductive realists (e.g. [10]).

To get luxury realism off the ground, we need to establish the following two claims. First, that the austere realist's model-theoretic methodology for identifying ontological commitments is in good standing. §3 identifies a problem for the straightforward construal of this methodology, but suggests a fix that makes the methodology viable. Second, that the ideology of explanation is available to the austere realist without thereby committing them to abductive realism. §4 spells out a proposal for how to understand explanation in exactly this sense.

3 The ontological question

Should we believe in the matter-independent existence of spacetime relations? Most of us, I venture, have some sense of what this question asks. Granted, we need to set out the appropriate standards of independence, and we need to agree on a way of distinguishing between matter and spacetime.

But it would be perverse to begin to answer this question by first analysing how we use the words 'spacetime', 'matter' and so on, wouldn't it? Surely a metaphysician would be justified in complaining 'I don't want to talk about how we talk about objects; I want to talk about the objects!'[8]. Against this impulse, as noted above, the austere realist argues that we cannot simply bypass semantics. And in particular, we ought to use model-theory to establish a word-world connection.

Promising though the austere realist's suggestion is, it immediately hits a snag in the context of the debate over realism regarding spacetime points: the standard model-theoretic characterisation of the substantivalist–relationalist debate made popular by Friedman [15] and Earman [13] is rendered incoherent, in a very specific way, by the nature of the model-theoretic apparatus, and specifically by considerations of the indeterminacy of model-theoretic reference made famous by Putnam [28]. I discuss this concern in §3.1.

But all is not lost for the austere realist. In §3.2, I demonstrate how the kernel of substantivalism–relationalism debate can be understood at a purely semantic level. The upshot of the discussion will be that, although the substantivalist–relationalist debate should be understood as a metaphysical dispute, one key aspect of the disagreement can be stated in a way that is entirely neutral with regard to one's metaphysics. I end this section with a novel, metaphysically-agnostic statement of the core of the debate, and demonstrate how it relates to previous commentators' characterisations. This renders the austere realist's construal of the ontological question coherent and intelligible, and makes it available to the luxury realist to purloin.

3.1 Sets, Models and Worlds

On a standard model-theoretic semantics, we begin with a *structure* $\mathfrak{A} := \langle D, \mathcal{I} \rangle$, which is a set-theoretic pair that comprises a domain of discourse, D, and an interpretation function $\mathcal{I}$. The interpretation function maps constant symbols to elements of D, $n-$place predicate symbols to $n-$tuples of elements of D and $n-$place function symbols to $n+1$-tuples on D. The domain of the interpretation function is called the *signature*, and elements of the signature are called *parameters*. Images of the signature under the interpretation function are called *extensions*: extensions of constants are domain elements, extensions of predicates are relations (of which properties are a special case) and extensions of function symbols are functions. Finally, we introduce a set of variables, V that are mapped to elements of D under a variable assignment g. Sentences are built out of variables, quantifiers, logical connectives and parameters by applying the appropriately defined syntactic rules. If, under the appropriate action of a variable assignment and an interpretation function, a sentence

is rendered true in a structure $\mathfrak{A}$, then that structure is called a model of that sentence relative to that variable assignment. Call the collection of sentences that are intended to be about a particular world, and that are modelled by some $\mathfrak{A}$ a *linguistic description* of that world.

One of Tarski's great contributions to semantics was to operationalise the concept of truth by relativising it to a model. But the operational gains come at a cost for metaphysical realists, i.e. people committed minimally to the mind- and model-independent truth of certain claims about the world. The problem is this that model-theoretic truth is criminally easy to come by, and as such, says almost nothing about the world that models are intended to model. The most forceful articulation of this fact is associated with Putnam [28], although precursors include Newman [25] and Quine [29]. I present the argument here as a dilemma.

Consider a model that purports to represent a world, in the sense that its domain consists of objects in the world, and its interpretation function picks out the properties and relations that the domain elements actually instantiate in that world. The (model-theoretic) truth of a linguistic description is then tested by checking whether its sentences are rendered true under that model. As Putnam famously pointed out, there is no semantic glue to bind a linguistic description to that domain. This leads to a serious problem when combined with the so-called 'homomorphism theorem', according to which if $\mathfrak{A}_1$ models a linguistic description, and $\mathfrak{A}_2$ is isomorphic as a structure to $\mathfrak{A}_1$, then $\mathfrak{A}_2$ also models that linguistic description (for a proof, see e.g. [9, p. 52-54]).

To see why this is a problem, consider the theory's 'intended model': the model according to which if the linguistic description is true about the world, it is true in that model. By the homomorphism theorem if $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are isomorphic, then they render true and false exactly the same sentences. Now consider some sentence S rendered false by the intended model $\mathfrak{A}_1$. By judiciously permuting the labels (i.e. by switching to a new interpretation function, and variable assignment, leaving unchanged their domains and target), we can create a new model $\mathfrak{A}_2$ on which S is rendered true. But, since all we did was permute the elements of the range of the original interpretation function (this is sometimes known as the 'push-through construction'), the new model $\mathfrak{A}_2$ is isomorphic as a model to $\mathfrak{A}_1$. So we end up with a dilemma: either every one of the isomorphic models is also intended, in which case, on pain of violating the homomorphism theorem, none of them can model S, or not all models isomorphic to an intended model are intended. In the latter case, we need some way to identify some proper subclass of isomorphic models as intended, and disallow unintended models. This is Putnam's challenge. It leads to a serious problem with the standard model-theoretic formulations of substantivalism and relationalism, as I discuss below.

3.2 Substantivalism and relationalism

On the implicitly reductionist picture employed by most commentators, the only irreducible physical entities in the domain of a model of a spacetime theory are points: everyday objects are built out of matter, conceived of either as particles, which are individually material points, or fields, which are regions of material points. Without loss of generality then, we can fully characterise the models of a spacetime theory by introducing two primitive properties (i.e. one-place relations) on the domain of the models: M: '...is material' and I: '...is immaterial'.[5]

According to Earman, and updating his terminology for consistency with this paper, the substantivalist is someone for whom '[material points] may or may not appear in the domain of every intended model, but [immaterial] points and regions do'. And according to the relationalist, '[material points] exhaust the domains of the intended models of the [nomically possible] worlds' [13, p. 114]. In light of Putnam's challenge, the problem with Earman's characterisation is clear: if every isomorphic model is intended, then there is no way to set up the substantivalist position (*mutatis mutandis* the relationalist). This is because every model that has a non-empty extension for 'immaterial' is isomorphic to a model that has a non-empty extension for 'material'. In other words, there exists a push-though construction that commits Earman's substantivalist to accepting some intended models with material points but no immaterial points or regions. The only way out, it seems, is to privilege a proper subclass of isomorphic models as intended. It seems, then, as if we need to resolve Putnam's challenge in order to state the substantivalism–relationalism debate.

However, there is an important sense in which we do not. We can accept the force of Putnam's challenge, and still state the core of the debate once we realise that the worry about whether or not M *really* picks out the *material* points is dialectically downstream. The more basic dispute is over whether or not there are *two* types of spacetime points (i.e. points that fall under extensions of the two predicates I and M) or just *one* type (i.e. points that fall under the extension of only one of the predicates; it does not matter which).

I suggest the following refinement of Earman's characterisation of the debate, which eschews intended models:[6]

`Substantivalism`$_1$: On the subdomain of relata of spatiotemporal relations, I has non-empty

[5] Not a great deal turns on precisely how to understand 'immaterial'. It can just as well be read as '...is materially unoccupied' or even '...is the kind of point whose existence is denied by the relationalist'.

[6] On this view, supersubstantivalism, is indistinguishable from relationalism at the semantic level. Supersubstantivalism is the view that only spacetime, but not material bodies, exists. It qualifies as a tertium quid only relative to a specific class of resolutions to Putnam's challenge; Lehmkuhl [20] discusses the extent to which supersubstantivalism should even be considered a tertium quid at all.

extension in every model and M has non-empty extension in some model(s).

`Relationalism`$_1$: On the subdomain of relata of spatiotemporal relations, M has non-empty extension in every model, and I has no non-empty extension in any model of the theory in question.

One might object: it's incredibly easy to divide the world into more than one type. For example, a self-professed relationalist might (indeed, should) be committed to matter fields of different types, for example, electromagnetic fields and electron fields. So even for the relationalist, points of spacetime do fall under the extensions of two disjoint predicates, '...is occupied by an electron field' and '...is occupied by an electromagnetic field'. Surely, then, this characterisation falls at the first hurdle.

The response to this objection begins with the following observation: although the characterisation of a relation as spatiotemporal is doing some important work here, that work is the same for the substantivalist and the relationalist, both of whom agree that certain relations are irreducibly spatiotemporal (for example, time intervals, distances, shapes, topological connectedness). By 'irreducible', I simply mean that their status as spatiotemporal relations as opposed to non-spatiotemporal relations is not subject to further reductive explanation from within the resources of the theories under consideration. As an analogy: that green is a colour is an irreducible fact about greenness, even though greenness is not necessarily itself an irreducible property of objects. Once we have established that the sorts of relations that obtain between two points are irreducibly spatiotemporal, as is assumed by both the substantivalist and relationalist, then any further non-spatiotemporal relations in which they may stand (such as '...is of the same material field type as...') are irrelevant to the discussion about substantivalism and relationalism. In other words, the level at which the points are typed for the purposes of setting up the substantivalist–relationalist debate is at the level at which the relevant distinction is between spatiotemporal and non-spatiotemporal; that the non-spatiotemporal relata may be further typed is irrelevant. This suggests the following refinement:

`Substantivalism`$_2$: On the subdomain of relata of spatiotemporal relations, *qua* spatiotemporal relata, the extension of I is non-empty in every model and the extension of M is non-empty in some model(s).

`Relationalism`$_2$: On the subdomain of relata of spatiotemporal relations, *qua* spatiotemporal relata, the extension of M is non-empty in every model, and the extension of I is empty in every model of the theory in question.

Importantly, there is no push-through construction that renders a substantivalist linguistic description

equivalent to a relationalist one. Even if we switch our intuitive understanding of what I and M are intended to pick out (i.e. if I picks out material points and M picks out immaterial ones), we have two inequivalent positions. This was not true of Earman's characterisation. So this construal of the positions is neutral with respect to resolutions to Putnam's dilemma. To be a relationalist is simply to be committed to there being only one type of spatiotemporal relatum, while to be a substantivalist is to be committed to there being two. That is all.

For the austere realist, the dispute is not over whether we should be realists about spacetime (whatever that means), rather it is over whether our best theories of spacetime, on whatever realist construal sidesteps Putnam-style worries, commit us to one or to two types of spatiotemporal relata. How one further understands this ontological commitment, in particular the properties and characteristics of the elements to which one is ontologically committed, will depend greatly on the further cocktail of semantic, metaphysical and epistemological considerations that characterise positions in the broader debates over scientific and metaphysical realism.

The austere realist is best-off reformulating the ontological question as:

`Ontological question*`: On the subdomain of relata of spatiotemporal relations, how many types of spatiotemporal relata are there, across models?

4 The explanatory question

The austere realist's methodology for arriving at ontological commitments regarding spacetime is now in good standing. And it makes no reference to explanatory dispensability. But just because we were able to answer the ontological question without invoking explanation, doesn't mean that the ideology of explanation should play no further role in articulating our broader metaphysical commitments. In this section, I defend the idea that the ideology of explanation is in good enough standing that the austere realist can put it to work in answering the explanatory question without thereby being committed to an affirmative answer to the abductive question.

Earman proposes, as a methodological principle reflecting good philosophical hygiene, that we always employ a formalism in which spacetime symmetries are co-extensive with dynamical symmetries and vice versa [13]. As a methodological suggestion, it is utterly plausible, almost mundane. After all, if the dynamical symmetries outstrip the spacetime symmetries, this means that spacetime has more structure than is required to articulate the laws (since fewer symmetries = more structure). So on a very basic Occamist norm regarding metaphysical theorising, we should excise this extra structure. Conversely, if spacetime symmetries outstrip the dynamical symmetries, the

spacetime structure proposed explicitly does not have enough structure to articulate the laws. So either the laws are improperly articulated, or they secretly invoke some extra spacetime structure that is not officially sanctioned by our choice of spacetime symmetries. Either way, bad methodological practice results in a mismatch.

But bad methodological practices sometimes yield good results. So some commentators have sought a metaphysical principle that might underwrite Earman's methodological principles. This brings us back to the explanatory question:

`Explanatory question`: Does spacetime structure ground a law's dynamical symmetry structure or vice versa?

Broadly speaking, the machinery of grounding is what allows us to formalise and characterise the '...in virtue of...' locutions that are ubiquitous in both natural and scientific languages. In particular, if A grounds B then, among other things, B obtains in virtue of A, which is to say A necessitates B in some nontrivial way. Thus if spacetime structure grounds the dynamical symmetry structure of a theory, in the sense that the theory's dynamical symmetries are what they are *in virtue of* spacetime structure being what it is, then spacetime structure necessitates the dynamical symmetry structure and Earman's principles come out as necessary. This fact remains unchanged even if there are multiple grounds, i.e. if for example spacetime structure only partially grounds dynamical symmetry structure, in the sense that the latter is also grounded in some other facts (in this case, the dynamical structure is necessitated by the combination of spacetime structure and something else).

Approaching the substantivalist–relationalist debate in terms of grounding is, by now, a relatively well-established move. Dasgupta [11] and North [26] construe the debate this way, as do Rosen [32], and arguably (although not explicitly) Belot [3]. The word 'grounding' appears in a variety of contexts. For some this fact is indicative of a deep and hidden feature of the world whose properties are slowly being uncovered by the techniques of analytic metaphysics in much the same way as traces in a bubble chamber are indicative of sub-atomic particles whose properties are slowly uncovered by the techniques of physics. Dasgupta summarises this view appropriately acerbically: '[the contemporary literature on grounding] seemed to take ground to be some *part* of reality, some metaphysical analogue of the Higgs boson that somehow held the world together. The job of a metaphysician, on this new conception, was to peer into reality and discern where these "groundons" were flowing (of course, to see these groundons one needed goggles provided by specialist departments)' [12, p. 74]

This is emphatically not the notion of ground that I rely on in this paper. I opt, instead, for

something closer to Dasgupta's deflationary account, which sits in a well-established tradition of metaphysical analysis that is sensitive to the dangers of taking too seriously the surface grammar of natural language. We have no more prima facie reason to believe that our talk of some fact grounding another fact commits us to some worldly relation between extant entities called facts than that our talk of doing things for other people's sake commits us to worldly entities called 'sakes' [29]. Dasgupta's characterisation of the broad contours of how grounding allows for 'limning many issues of intellectual interest' will suffice for our purposes. His proposal is straighforward: it is sufficient to treat grounding as nothing more than providing a *constitutive explanation* of certain facts (or claims or propositions).

The notion of a constitutive explanation is perhaps best understood with the help of an example that contrasts it with the more intuitive and familiar notion of a casual explanation. Suppose we were to ask the question 'why did India lose the final of the cricket World Test Championship?' Our everyday-use of 'why' suggests we are looking for a causal explanation, something like 'because they were unnecessarily attacking on the final day'. But there is another equally legitimate answer, which construes the question as asking for a constitutive explanation like 'because they lost twenty wickets'. Thus, while their loss was caused by particular choices of tactics, that they lost was in virtue of, or *grounded in*, their having lost twenty wickets. Dasgupta rightly cautions us not to succumb to the temptation to reify more than is required: the fact that one of the teams lost can be grounded in the fact that they lost twenty wickets without it being the case that we are thereby committed to an ontology of facts, truthmakers, propositions or anything beyond two teams who engaged in a cricket match.

The proponent of the geometrical approach to spacetime claims that the spatiotemporal relations that obtain between spacetime points constitutively explain the fact that the laws that describe the goings-on of matter in that spacetime have certain symmetries. All this means is that spacetime structure is, in an appropriate sense, more metaphysically basic than nomological structure in such a way that brute facts about spacetime can underwrite a particular sort of explanation—a constitutive explanation—about certain properties of the laws. This account is, of course, wholly consistent with more inflationary stories about grounding (including ones which embrace Dasgupta's groundons), but crucially, this account does not require any such heavy-duty metaphysical commitments.

The proponent of the dynamical approach simply reverses the direction of the constitutively explanatory arrow. In their original presentation of the dynamical approach [7], Brown and Pooley identify, as what grounds the structure of Minkowski spacetime, the Lorentz-invariance (i.e. the dynamical symmetry properties) of the laws that describe matter fields. They call this the *truncated*

Lorentzian pedagogy: raise to the status of unexplained postulate the fact that all matter fields (whether quantum or not) are Lorentz-invariant. It will follow that measuring devices like rods and clocks (call them 'surveyors') built out of such fields will all read out the same value of proper time along the same trajectory. [7]Consequently, there is no need to posit an ontologically independent metric field to be measured—it would just be an idle wheel since, as Brown famously puts it '[material entities] have no spacetime [geometry] feelers' [5, p. 24]. Therefore, on Occamist grounds, one can excise the Minkowski metric from the primitive structure of SR (hence Brown and Pooley's sometimes-misinterpreted claim, in a later paper, that Minkwoski spacetime is a 'glorious non-entity'; their claim is that the Minkowski metric is not an ontologically independent entity [6]).

A quick note to hopefully assuage the misgivings of some metaphysicians, regarding the intensional characterisation of constitutive explanation presented above. There is an intuition shared by many, according to which explanation is a hyperintensional concept: the truth-value of propositions built from relations between explanantia and explananda can change even under substitution of those explanantia and explananda by necessary equivalents. This intuition is pumped by the examples in which the relation of explanation is intuitively more fine-grained than modal distinctions allow: for example, that some mathematical truths explain others, despite all such truths being necessary. Does it follow, then, that every legitimate use of 'explains' needs to invoke hyperintension? No. For Lewis [23], causal explanations are perfectly intelligible and expressed in modal terms (notice how most counterexamples to the sufficiency of modal explanations are non-causal). For Kincaid [19], supervenience explanations are legitimate accounts of some relations between lower- and higher-level scientific theories. For Friedman [15] explanation in physics can be characterised in terms of embeddability of models—a modal concept. The deflationary view of 'grounding-as-constitutive-explanation' that I endorse here simply sees constitutive explanation as another form of explanation. It does not deny that some explanations are hyperintensional, it merely denies that all good explanations are.

[7] One might worry that the dynamical approach can, consequently, not be extended to a theory like General Relativity (GR) because, in that theory, generic spacetimes do not have any non-trivial symmetries. This is not a concern; there are two proposed routes for extending the dynamical approach to GR. The first is from Brown himself [5, Ch. 9], and involves understanding the chronogeometric significance of the metric (i.e. the fact that measuring devices read off intervals of proper time along their trajectories) as grounded in the so-called strong equivalence principle (see [21, 31]), which invokes claims about local dynamical symmetries in GR (see [14, 35] for discussions of the viability of this approach). An alternative suggestion, from Pooley [27, §6], is to extend Huggett's [16] 'regularity approach' to the GR mosaic. In this paper, I bracket these discussions and focus on dynamical substantivalism as it applies to highly symmetric spacetimes like Newtonian or Minkowskian. But the doctrine of dynamical substantivalism can straightforwardly be extended to GR; the precise details will depend on how the dynamical approach is construed in that theory.

5 Luxury realism and dynamical substantivalism

In §2, I introduced luxury realism is a hybrid position constructed from elements of both austere realism and abductive realism. In §3, I demonstrated that the austere realist's methodology for extracting ontological commitments using model-theory was viable. And in §4, I argued that although the ideology of explanation need not play a role in spelling out the realist's ontological commitments, it can nonetheless play an important role in fleshing out an attractive metaphysical position. In this section, I argue that luxury realism is an attractive position: it inherits many of the strengths of austere realism, and adds to it further inferential resources, which together make available the position of dynamical substantivalism.[8] I demonstrate the attractiveness of dynamical substantivalism by discussing how the dynamical substantivalist makes sense of one of the most central thought experiments in the philosophy of space and time: Newton's spheres.

In the final paragraph of the Scholium to the definitions of the Principia, Newton introduces a thought experiment: imagine two rigid spheres of different masses connected by a light inextensible string in an otherwise matter-free possible world that is nomically possible according to Newton's laws. Consider the situation in which the spheres revolve about their common centre of mass. Intuitions about the relationship between angular momentum and force suggest that the tension in the string will increase as a function of the angular velocity. Assume all parties agree that the magnitude of the force in a world supervenes only on quantities in that world.

For the substantivalist, angular velocities can be defined with respect to the immaterial points of spacetime, so different worlds can agree about the laws (or more generally be nomically possible according to the same laws) yet disagree about the magnitude of the tension in the string, since those worlds can disagree about the magnitude of angular velocity of the system. Note that there need not be any commitment to the spatiotemporal structure grounding the laws: it is merely the existence of relations between material and immaterial points that is doing the necessary metaphysical work for the substantivalist. The relationalist does not quantify over immaterial points, so cannot make sense of the quantity of angular velocity (and variations thereof) in an otherwise matter-free world. Therefore, the only way in which the relationalist can make sense of different tensions in the string in different possible worlds is by conceding that those worlds are governed by different, mutually inconsistent laws: the force function in each world maps the same relative distance to different string tensions. Consequently, the relationalist's plurality of nomically possible worlds, for a large class of

[8] It also makes available the position of geometrical relationalism; I am less confident of the plausibility of that position, but am nonetheless happy to have it rendered intelligible by luxury realism.

mechanical laws like Newton's, is severely impoverished; whereas the Newtonian substantivalist can claim a continuously infinite number of nomically possible worlds corresponding to any given relative distance between the spheres, the Newtonian relationalist can claim only one nomically possible world for each relative distance.[9]

Without luxury realism, those who are led by Newton's thought experiment to endorse substantivialism have their hand forced in the dynamical/geometrical debate: they need to go geometrical. Luxury realism gives the substantivalist a new freedom. Their hand is no longer forced—they can choose whether to go dynamical or geometrical. A similarly novel and attractive option becomes available to the proponent of the dynamical approach. And this should be welcomed: it is unfair to saddle the proponent of the dynamical approach, for whom questions of explanatory priority concern the behaviour of matter rather than ontology, with the relationalist's ontologically impoverished set of nomically possible worlds. As an example to help make vivid this worry, consider how a Maudlinian primitivist view of laws of nature [24] might motivate a dynamical view.Nothing of importance turns on this particular view of laws of nature. The choice of primitivism was made for dialectical convenience. Analogous arguments are available to proponents of most mainstream views on laws of nature, such as the Mill-Ramsey-Lewis theorist as well as the Dretske-Armstrong-Tooley theorist.

On a four-dimensionalist picture of the world, dynamical laws provide a way of encoding the patterns of configurations of material in spacetime. Each such configuration is a complete history of all the goings-on in a world. That some universal generalisations over this configuration count as laws is just a primitive fact about a world. On the nomic primitivist's picture, the behaviour (i.e. four-dimensional configuration) of entities in a possible world is fixed by the laws of that world. It might be a contingent fact that certain entities exist in a possible world, but if those entities exist in a world, their behaviour is necessitated by the laws. Consider a substantivalist who is a primitivist about laws in this sense. For a substantivalist of this bent, if there exist material particles in some world (that is to say if some spacetime points are material), their configuration across spacetime is determined entirely by the laws in that world. But a spacetime configuration is nothing more than the collection of spatiotemporal relations in which the material points stand: the spacetime geometry, in other words. Thus, the primitive fact that the laws of Newtonian mechanics is Galilean-invariant, combined with the primitivity of the laws themselves constitutively explains the Galilean geometry of (at least the materially occupied points of) Newtonian worlds.

[9] For relationalists like Mach, for example, this was hardly a demerit: their argument was effectively a Moorean shift against the substantivalist, to argue that the space of nomically possible worlds *should* be impoverished by comparison with the substantivalist's.

So much for the materially occupied points: the laws imbue them with geometric structure. But what of the immaterial points? Consider a model which consists of a smooth manifold of spacetime points, some of which are material. Dynamical equations govern the behaviour (and thus ground the spatiotemporal geometry) of all and only these materially occupied points. Suppose, according to the dynamical laws, the material spacetime point p stands in some determinate spatiotemporal relation to (say is 7 units of proper time away from) another material point q. The substantivalist is committed to the existence of immaterial points as well, so they can ask the question: in what relation(s) does immaterial point r stand with respect to p and q?

For the substantivalist, once a point r has been chosen then, on pain of inconsistency, the relations between r and p, and r and q must compose in such a way as to reproduce the determinate relation between p and q. Since the material points stand in certain metric relations to one another, it cannot be the case that they fail to stand in metric relations to immaterial points. So, in fact, *all* geometrical relations are grounded in the symmetry structure of the laws. Thus, the combination of the dynamical approach and substantivalism is entirely coherent.

In fact, dynamical substantivalism is more than merely coherent, it is arguably rather attractive. Let us return to Newton's spheres, and consider some relation R_1 that obtains between two material spheres, idealised as points named p and q. According to the dynamical substantivalist, this relation can ground the existence of two other relations, call them R_2 which obtains between the p and the immaterial point r and R_3 which obtains between r and q. Nothing prevents the substantivalist from then asserting that there exist distinct nomically possible worlds that agree on R_1 but disagree on R_2 and R_3, and further, that all of these facts are grounded in facts about the laws. On accepting this, the proponent of the dynamical approach can argue that angular velocity is a legitimate quantity to subvene the tension in the string, and thus repopulate their space of nomically possible worlds to bring it back on par with the substantivalist's. Thus luxury realism can allow for a dynamical view on the explanatory question while still leaving room for a substantivalist view on the ontological question.

Of course, you do not need to be a dynamical substantialist in order to be a luxury realist; the other three permutations of positions are still available to the luxury realist. The arguments in this section are intended as a proof-of-concept, to demonstrate the value of luxury realism in making available a metaphysics of spacetime in which both substantivalism and the dynamical approach can be made compatible. So even if you are unmoved by dynamical or substantival considerations, it should be clear that luxury realism allows for recipe to make positions in the metaphysics of spacetime compatible with a wider range of antecedent metaphysical commitments.

6 Conclusion

The goal of this paper was to highlight the distinctness of certain metaphysical options, and to demonstrate that they can be combined in various ways to carve out a unique and attractive position regarding the metaphysics of spacetime. On prising apart the ontological and explanatory questions, I suggested a hybrid form of realism, which incorporated elements of both austere as well as abductive realism, called luxury realism. I argued in favour of the latter on the basis of its ability to make available a more finely-grained metaphysics of spacetime, in particular, its ability to allow for geometrical relationalism and dynamical substantivalism. Finally, I provided an example of a coherent and attractive account of Newton's spheres that is available only to the dynamical substantivalist. With this, I hope to have shown that reports of the death of the substantivalism–relationalism debate (see e.g. [33]) have been greatly exaggerated.

References

[1] Pablo Acuña. Minkowski spacetime and Lorentz invariance: The cart and the horse or two sides of a single coin? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 55:1–12, August 2016.

[2] David John Baker. Knox's inertial spacetime functionalism (and a better alternative). *Synthese*, March 2020.

[3] Gordon Belot. Rehabilitating relationalism. *International Studies in the Philosophy of Science*, 13(1):35–52, March 1999.

[4] Richard N. Boyd. Realism, Underdetermination, and a Causal Theory of Evidence. *Noûs*, 7(1):1–12, 1973.

[5] Harvey Brown. *Physical Relativity*. Oxford University Press, Oxford, 2005.

[6] Harvey Brown and Oliver Pooley. Minkowski space-time: A glorious non-entity. *Philosophy and Foundations of Physics*, 1:67–89, 2006.

[7] Harvey R. Brown and Oliver Pooley. The Origins of the Spacetime Metric: Bell's Lorentzian Pedagogy and its Significance in General Relativity. In Craig Callender and Nick Huggett, editors, *Physics Meets Philosophy at the Plank Scale*, pages 256–72. Cambridge University Press, 1999.

[8] Tim Button. *The Limits of Realism*. Oxford University Press, Oxford, 2013.

[9] Tim Button and Sean Walsh. *Philosophy and Model Theory*. Oxford University Press, Oxford, 2018.

[10] Shamik Dasgupta. The Bare Necessities. *Philosophical Perspectives*, 25(1):115–160, December 2011.

[11] Shamik Dasgupta. Substantivalism vs Relationalism About Space in Classical Physics: Substantivalism vs Relationalism About Space. *Philosophy Compass*, 10(9):601–624, September 2015.

[12] Shamik Dasgupta. Constitutive Explanation: Constitutive Explanation. *Philosophical Issues*, 27(1):74–97, October 2017.

[13] John Earman. *World Enough and Spacetime*. MIT Press, Cambridge, MA, 1989.

[14] Samuel C Fletcher. Approximate local poincaré spacetime symmetry in general relativity. In *Thinking about space and time*, pages 247–267. Springer, 2020.

[15] Michael Friedman. *The Foundations of Spacetime Theories*. Princeton University Press, Princeton, NJ, 1983.

[16] Nick Huggett. The Regularity Account of Relational Spacetime. *Mind*, 115(457):41–73, 2006.

[17] Christopher Isham. *Modern Differential Geometry for Physicists*. World Scientific Publishing Company, London, 1999.

[18] Michel Janssen. Drawing the line between kinematics and dynamics in special relativity. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 40(1):26–52, January 2009.

[19] Harold Kincaid. Supervenience and Explanation. *Synthese*, 77(2):251–281, 1988.

[20] Dennis Lehmkuhl. The Metaphysics of Super-Substantivalism. *Noûs*, 52(1):24–46, 2018.

[21] Dennis Lehmkuhl. The equivalence principle (s). In *The Routledge companion to philosophy of physics*, pages 125–144. Routledge, 2021.

[22] David Lewis. Putnam's paradox. *Australasian Journal of Philosophy*, 62(3):221–236, September 1984.

[23] David Lewis. *Philosophical Papers, Volume 2*. Oxford Univeristy Press, New York, 1986.

[24] Tim Maudlin. *The Metaphysics within Physics*. Oxford University Press, Oxford, 2007.

[25] M. H. A. Newman. Mr. Russell's "Causal Theory of Perception". *Mind*, 37(146):137–148, 1928.

[26] Jill North. *A New Approach to the Relational–Substantival Debate*, volume 1. Oxford University Press, September 2018.

[27] Oliver Pooley. Substantivalist and Relationalist Approaches to Spacetime. In *The Oxford Handbook of Philosophy of Physics, Robert Batterman (Ed)*. Oxford Univeristy Press, Oxford, 2013.

[28] Hilary Putnam. Realism and Reason. *Proceedings and Addresses of the American Philosophical Association*, 50(6):483, August 1977.

[29] Willard van Orman Quine. *Word and Object*. MIT Press, Cambridge, MA, 1960.

[30] James Read. Geometrical Constructivism and Modal Relationalism: Further Aspects of the Dynamical/Geometrical Debate. *International Studies in the Philosophy of Science*, 33(1):23–41, January 2020.

[31] James Read, Harvey R Brown, and Dennis Lehmkuhl. Two miracles of general relativity. *Studies in history and philosophy of science Part B: Studies in history and philosophy of modern physics*, 64:14–25, 2018.

[32] Gideon Rosen. Metaphysical Dependence: Grounding and Reduction. In *Modality: Metaphysics, Logic and Epistemology, Bob Hale and Aviv Hoffman (Eds.)*, pages 109–136. Oxford University Press, Oxford, 2010.

[33] Robert Rynasiewicz and Journal of Philosophy Inc. Absolute Versus Relational Space-Time: An Outmoded Debate?:. *Journal of Philosophy*, 93(6):279–306, 1996.

[34] James Owen Weatherall. Conservation, inertia, and spacetime geometry. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 67:144–159, August 2019.

[35] James Owen Weatherall. Two dogmas of dynamicism. *Synthese*, 199(2):253–275, 2021.