

Physical Relativity from a Functionalist Perspective

Eleanor Knox
King's College London
eleanor.knox@kcl.ac.uk

July 30, 2017

Abstract

This paper looks at the relationship between *spacetime functionalism* and Harvey Brown's dynamical relativity. One popular way of reading and extending Brown's programme in the literature rests on viewing his position as a version of relationism. But a kind of spacetime functionalism extends the project in a different way, by focussing on the account Brown gives of the role of spacetime in relativistic theories. It is then possible to see this as giving a functional account of the concept of spacetime which may be applied to theories that go beyond relativity. This paper explores the way in which both the relationist project and the functionalist project relate to Brown's work, despite being incompatible. Ultimately, these should not be seen as two conflicting readings of Brown, but two different directions in which to take his project.

1 Introduction

Harvey Brown's dynamical relativity has had a considerable impact on the philosophy of spacetime physics since its full articulation in his 2006 book *Physical Relativity: Space-time Structure from a Dynamical Perspective*. It rests on deep and probing questions: *what is it about the detailed dynamics of a theory of matter in spacetime that leads bodies to reflect spacetime structure? Why does the geometry of general relativity's metric field have direct operational significance?* Brown holds that the answers to these questions lie in the details of the dynamics: in special relativity the key to geometry lies in the Lorentz invariance of the laws; in general relativity, the strong equivalence principle is essential to the chrono-geometric significance of the metric.

Many philosophers, myself among them, have found these questions, and Brown's answers, provoking and compelling. But while the literature is clear on the interest of Brown's work, it's less clear on the philosophical upshot: what exactly are the consequences of Brown's view for our understanding of the nature of spacetime itself? Brown himself is more interested in physics than metaphysics, so interlocutors have often been left to themselves to puzzle out how his view might be extended to one that takes an explicit position in traditional debates like that between the substantivalist and the relationist. Many see strong hints in Brown's work that he intends to advocate a new form of relationism. In particular, a focus on the explanatory status of the Minkowski metric makes many of Brown's arguments look like attempts to block an abductive inference to the existence of Minkowski spacetime. For the most part, both arguments against Brown's view, and extensions of the view, take this to be the obvious reading.

But not all aspects of Brown's book suggest relationism. I'll argue in what follows that, while the standard reading makes a great deal of sense when applied to his statements about Minkowski spacetime, it commits Brown to an implausible relationism in the case of general relativity. This leaves us with another option: Brown's dynamical relativity is best read not as a salvo in the substantivalism wars, but rather as an account of what counts as spacetime structure in a given physical theory. Wayne Myrvold argues for just such a reading in his contribution to this volume; according to Myrvold, Brown is best seen as giving an analysis of the concept of spacetime, rather than as defending a view about its reality. This seems right to me. In what follows, I'll articulate a particular version of this kind of conceptual view that I call *spacetime functionalism*. Not everything about spacetime functionalism sits easily with Brown's arguments, but the view owes a great deal to Brown's dynamical relativity, and allows for the application of his ideas to both special and general relativity, and many more theories besides.

As we'll see in what follows, there is some tension between the two readings suggested above. Which reading does Brown intend? The simple answer is: neither. Both relationism position based on dynamical relativity, and a general analysis of the concept of spacetime, are extensions of Brown's work, not interpretations of it. Much of what follows will involve reading more into Brown's work than his tight focus on the details of relativity requires. But it is nonetheless instructive, and a testament to the richness of the work, to explore both extensions and see where they lead us - ultimately to two quite different places.

In section two, I examine the case advancing a relationist thesis based on Brown's work. Not only is this the reading most interlocutors have chosen, but it makes good sense of Brown's interest in explanatory issues.

In section three, I consider the trickier case of general relativity, and argue that relationism based on dynamical relativity here is much less attractive; not

only must relationism about GR be of a different stripe from special relativistic relationism, but it also relies on the idea that a single field cannot share the properties of spacetime and matter, something that is brought into question in the general relativistic context.

Section four will articulate an alternative to this relationism: spacetime functionalism. I'll argue for functionalism about the spacetime *concept*, and for a functional role for spacetime that incorporates key insights from Brown's work. One virtue of this view is that it plausibly applies to both kinds of relativity, as well as to other theories.

Section five will examine the relationship between this version of spacetime functionalism and dynamical relativity. There is no doubt that many of the arguments made by Brown can be read as supporting this particular breed of functionalism, if not functionalism more generally. But at the same time, the position sits uncomfortably with relationist extensions of Brown's work. The view itself is neutral on substantivalism and relationism in general, but it can deliver a substantivalist verdict when applied to a theory with reasonably clear ontological commitments: general relativity is just such a case.

2 Dynamical Relativity as Relationism

In 2006, Harvey Brown and Oliver Pooley published a paper entitled "Minkowski space-time: a glorious non-entity". This paper sets out some of the central arguments that are developed and expanded in Brown's book.¹ The paper, and the book, argue that one can think of special relativity's Minkowski spacetime as a 'codification' of the symmetries of matter fields. All that is required for bodies made of these matter fields to reflect a Minkowskian geometry is that the fundamental equations be Lorentz covariant. This contrasts with geometry-first approaches, which take geometry to be 'prior to' the dynamics in some sense,² and which often come with an implicit or explicit commitment to taking spacetime geometry to be the structure of a substantival spacetime.

Given this description, and the title of Brown and Pooley's paper, one might suppose that dynamical relativity is intended as a form of relationism, albeit one different from traditional forms. Spacetime structure on this account is reduced not to relations between bodies, but to symmetries of the laws themselves.³

¹The themes in this paper, and the book, are also developed elsewhere. See [8, 4, 5, 7].

²Of course, what 'prior to' means here is a matter of some debate. If one means ontologically prior to, one simply assumes a position in the substantivalism/relationism debate without argument. But if one means explanatorily prior to, one enters into debates about the nature of explanation which have no clear conclusion. It is in part because of this difficulty that I prefer in my own work to develop a different strand of Brown's thought.

³Of course, the metaphysical implications of such a view turn on our account of the laws of

The arguments offered in both the book and the paper seem to support this reading. One of Brown's starting points is Einstein's distinction between *principle* and *constructive* theories. Those who find the distinction instructive note that members of the former class of theories are, like thermodynamics, deduced from established empirical principles; they enjoy great empirical support, but their central postulates go unexplained. Members of the latter class explain phenomena in terms of the constituents of matter, and thus offer deep understanding and, perhaps, the hope of metaphysical insight.⁴ Einstein famously noted that his 1905 formulation of special relativity was a principle theory, and therefore, despite enjoying great empirical success, failed to explain *why* relativistic phenomena occur. For this, the thought goes, we need a constructive theory.

Although, prior to Brown's work, this particular issue had been relatively little discussed, there is something like an 'orthodox' view of the geometry of spacetime that suggests a response to the problem above.⁵ It goes along the following lines: Minkowski's 1908 geometrization of special relativity precisely *was* the presentation of a constructive account. It explained phenomena like length contraction by appealing to the geometry of spacetime itself. Balashov and Janssen make this argument explicitly:

length contraction is explained by showing that two observers who are in relative motion to one another and therefore use different sets of space-time axes disagree about which cross-sections of the 'world-tube' of a physical system give the length of the system. [3, p.331]

The idea here is that a phenomenon like length contraction is explained by appealing to the cross-section of a rod whose invariant description is given in a four-dimensional Minkowski geometry. The thought is, in part, that spacetime itself forms part of the explanation, and hence the picture is constructive.⁶ Of course, the fan of spatiotemporal explanation can grant some of

nature. Brown (in conversation, at least) is inclined towards a Humean view, in which case the laws themselves are reducible the behaviour of material bodies. [22] gives more detail as to how one can build dynamical relativity on Humean foundations.

⁴There is room, of course, to question whether this distinction is really coherent, especially when we think of our physics as involving a complex layering of emergent theories with no known fundamental theory. Brown himself now thinks that this distinction has been something of a distraction in the debates that followed. For a critique of the role that the distinction has played in this debate, see [20].

⁵The 'orthodox' view of spacetime geometry is sufficiently orthodox to have been assumed rather than argued for and hence has received relatively little explicit defense outside the literature addressing Brown's position. One exception is Tim Maudlin's book on the philosophy of space and time [18]

⁶For brevity's sake, I gloss over distinctions between those who advocate spatiotemporal constructive explanation here. Balashov and Janssen do not in fact take their argument to commit them to the reality of Minkowski spacetime, whereas Maudlin, for example, is fairly explicit in this commitment.

Brown's agenda: appealing *only* to spacetime geometry to explain phenomena like length contraction gives, at best, a partial constructive explanation and, at worst, no explanation at all. What is lacking is an account of why particular matter systems reflect the geometry. But here the fan of a geometrical constructive theory can appeal to the Lorentz covariance of the fundamental laws; it's just that, contra Brown and Pooley, they see these as explained by the structure of spacetime. Alternatively, (and equivalently), they can appeal to the generally covariant form of those laws in which the Minkowski metric appears explicitly.

By contrast, Brown holds that the best explanation for phenomena like length contraction appeals only to the Lorentz covariance of the laws, which is a brute, unexplained fact. His argument for this rests on scepticism about so-called 'geometrical explanation':

It is doubtful at best whether the geometries of the configuration space in classical mechanics, or the space of equilibrium states in thermodynamics, play the kind of explanatory role that the spacetime interpretation of SR attributes to Minkowski geometry. Why should space-time geometry be any different? It might be thought that space and time are somehow more fundamentally physical than the other space, or more accessible to the senses, or that they combine to form the arena of physical events. In short that they are more real. But is not this reasoning question-begging? [6, p.139]

The merits of Brown's explanatory argument have been much discussed in the literature;⁷ I don't intend to rehash the debate here. Rather, I'd like to note that Brown's argument here may be used to block an inference to the best explanation by the substantialist. The substantialist claims that belief in Minkowski spacetime is warranted by the explanatory utility of Minkowski spacetime; Brown counters that the explanation in question is obscure or illusory. So Brown certainly intends his position to weigh against the substantialist.

Does Brown's anti-substantialism mean that he is proposing a new form of relationism? Brown rarely uses this language, but his collaborator, Oliver Pooley, certainly thinks of the view this way; in [22] he lists dynamical relativity as a form of relationism. And, at least in special relativity, the shape of the view is easy to see: in Brown's relationism, spacetime structure is reducible to facts about the symmetries of the dynamical laws.

⁷See, for example, [3][13][12][11] or [22].

3 General Relativity as a Problem Case

In special relativity, then, Brown's arguments have some force against the substantialist, and can be extended to a kind of relationism, albeit one very different from the kind of direct reduction of spacetime to the relations of material bodies that Mach envisaged. But what of general relativity? Here the reading is much harder to maintain. What can we make, in the GR context, of the idea that the geometrical structure of spacetime is reducible to the symmetries of the matter fields? Not only do matter fields in general relativity not generally *have* global symmetries,⁸ but the geometrical structure of general relativity doesn't supervene on *any* aspect of the non-gravitational fields; the metric field has its own degrees of freedom.

It goes without saying that Brown, who has a deep understanding of the intricacies of general relativity, knows this perfectly well. Brown's interpretation of general relativity diverges from his view on special relativity; he attaches great significance to features of the former that are not apparent in the latter. In particular, he points out that the availability of a geodesic theorem for GR means that GR offers an explanation of inertial behaviour that need not simply assume a connection between spacetime structure and the movement of bodies:

Inertia, in GR, is just as much a consequence of the field equations as gravitational waves. For the first time since Aristotle introduced the fundamental distinction between natural and forced motions, inertial motion is part of the dynamics. It is no longer a miracle. [6, p. 163]

Since the publication of Brown's book, further aspects of the geodesic theorem have come to light. Proving the geodesic theorem requires substantive assumptions about the matter fields, and similar theorems seems possible for other theories.⁹ But Brown's interest in the geodesic theorem is part and parcel of a view that pays great attention to those aspects of general relativity that connect the metric to the behaviour of non-gravitational bodies. The *strong equivalence principle* is, by Brown's lights, absolutely key to the interpretation of general relativity.

A more far-reaching claim is the *strong equivalence principle* (SEP), which will be defined here as follows. There exist in the neigh-

⁸Of course, depending on what one thinks about the issue of 'substantive general covariance', one can always think of the diffeomorphism group as being the symmetry group of GR's matter fields, but there's no question of the *metrical* structure of GR being inferred from diffeomorphism invariance.

⁹See [17] for a discussion of the assumptions behind the geodesic theorem and [25] for a similar theorem in a Newtonian context.

bourhood of each event preferred coordinates, each called locally inertial at that event. For each fundamental non-gravitational interaction, to the extent that tidal gravitational effects can be ignored, the laws governing the interaction find their simplest form in these coordinates. [6, p.169]

Brown holds that the “chronogeometric significance of the $g_{\mu\nu}$ field is not an intrinsic feature of gravitational dynamics, but earns its spurs by way of the strong equivalence principle”[6, p.151].¹⁰ The central point here is that it is the fact that $g_{\mu\nu}$ defines inertial structure that means that it represents geometry. This fact is not determined by the mathematical form of $g_{\mu\nu}$ but rather by its relationship to the rest of the dynamics. In particular, the minimal and universal coupling of the matter fields to the metric field means that, in those coordinates (the Fermi Normal Coordinates) in which the metric diagonalises and the connection coefficients vanish, the laws take a simple and universal form, at least to the extent that curvature terms do not appear explicitly in the laws themselves.¹¹ We shall see in the next section that there is a way of reading this that suggests a kind of functionalism: the metric field is spacetime because of what it does (via the strong equivalence principle) and not by way of what it is. But how might we interpret this kind of claim in a way that is consistent with the relationist themes in the previous section?

There is, as far as I am aware, just one way to defend the idea that general relativity is a relationist theory, and that is to insist that the metric field is a matter field *and not spacetime*. Carlo Rovelli famously takes this view, illustrating it with a memorable thought about gravitational waves:

A strong burst of gravitational waves could come from the sky and knock down the rock of Gibraltar, precisely as a strong burst of electromagnetic radiation could. Why is the [second] matter and the [first] space? Why should we regard the second burst as ontologically different from the first? Clearly the distinction can now be seen as ill-founded. [24, p.193]

Brown quotes this passage approvingly, and Rovelli’s relationism appears in some ways to be close to his own position; Brown, for example, warns against “conflating $g_{\mu\nu}$ with spacetime itself”(p.159). But Rovellian relationism differs significantly from the kind of dynamical relationism discussed above in the case of special relativity. In the special relativistic case, the geometry of spacetime reduced to symmetries of the dynamics. In this case, the geometry

¹⁰In recent work with Dennis Lehmkuhl and James Read, [23] Brown endorses this message, but notes that there are considerable subtleties here once one acknowledges that, even in the presence of minimal coupling, dynamical equations in GR do not always take their special relativistic form in freely falling coordinates.

¹¹See [23] for detailed discussion of minimal coupling and its consequences for higher order equations.

of spacetime is given by a single matter field, which represents the geometry only because of its role in the dynamics. Although the two positions share a commitment to relationism and an interest in the importance of dynamics to geometry, there is no straightforward sense in which the position in special relativity simply carries over to the general relativistic case.¹²

If this were the only issue, Rovellian relationism might nonetheless be defensible; it's hardly surprising that the move to GR should bring with it a change in one's view of spacetime metaphysics. But the proposed relationism for GR is much less attractive than Brown's dynamical relationism in the case of special relativity. In the earlier case, one could read Brown as arguing for eliminativism with respect to the Minkowski metric. In the GR case, spacetime geometry is now held to be instantiated by a matter field. But this only has force against the substantivalist if one also assumes that something with the characteristics of a matter field cannot *also* be spacetime. One finds no argument in Brown's work, or in the work of others, for such a position. Indeed, Brown's emphasis on the chronogeometric significance of the metric means that his account explains in detail exactly how the metric comes to represent spatiotemporal quantities.

Given all of this, it may be worth seeing Brown's endorsement of Rovelli in a different light: when Brown emphasises the matter-like properties of the metric field, he is interested in highlighting the sense in which this field, in the absence of its relation to other fields, does not possess any intrinsically spatiotemporal characteristics. This sets the scene for an account which explains just why, given the dynamics of general relativity, this field is in fact spacetime. And *this* argument points not to the reduction of spacetime to a matter field and hence to a form of relationism, but to an explanation of just why and how it is that general relativity is a theory of spacetime. The next two sections will explore these ideas in more detail.

4 Spacetime Functionalism

This section will describe an alternative way to think about the interpretation of spacetime theories: spacetime functionalism. I'll advocate a version of the position that incorporates several of Brown's insights, although I'll leave it to the next section to fully flesh out the relation between this position and Brown's own.

Spacetime functionalism, as the name suggests, describes any position that takes spacetime, spacetime regions, or spatiotemporal relations to be picked out by their functional role. Spacetime functionalisms of various kinds might

¹²Oliver Pooley also makes this point when discussing the merits of Brown's relationism in GR [22, p.578].

have some things in common: for example, they promise to help to make sense of theories of quantum gravity in which spacetime is emergent; it is relatively straightforward to make sense of the idea that functionalized objects or properties may be non-fundamental. They also share a certain ability to cut across old substantialist/relationist distinctions. But beyond this, positions might¹³ vary widely. There is little more to be said until one has picked out the kind of functionalism one has in mind; one must say exactly what is being functionalized, and what the functional role is. One may then approach the task of working out what plays the functional role; in the kind of functionalism I have in mind, the results of this process will then further depend on the theory under consideration.

I'll argue here for a functionalism about the spacetime *concept* that ties spacetime to a certain role within a broader theory. But one might instead, for example, hold that spatiotemporal regions are functionally picked out by the causal role that they play. Alyssa Ney [21] and David Albert [1][2] defend a position like this in the context of space in quantum mechanics; they think that space can be recovered from wavefunction realism (in 3N-dimensional configuration space) via a functional characterization of spatial regions. David Chalmers [9] advocates a kind of functionalism about spatiotemporal relations aimed at recovering the content of our experiences.

These strains of functionalism differ markedly from the kind of spacetime functionalism I advocate here. For one thing, I'm interested in recovering theoretical spacetime structure, rather than phenomenological space (although one would expect there to be a relation between the two). My functionalism instead involves the claim that *spacetime* is a *functional concept*. A structure that plays the spacetime role in a theory just *is* spacetime; once one has analysed the role and understood what fills it, there are no further questions to be asked about the 'real' spacetime structure. Spacetime is spacetime by virtue of what it does, not what it is. In his entry in this volume, Wayne Myrvold notes

...that a full-blown defense of [Brown's] view requires considerations of a sort that are very much out of fashion in the contemporary philosophical landscape. These involve considerations of the way that our concepts - in this case, spacetime concepts - gain purchase on the physical world.[19]

My view involves just this kind of unfashionable conceptual consideration.

This kind of concept functionalism is defended in broad strokes by the kind of consideration that Myrvold appeals to: *what else could spacetime be?* If our

¹³I say 'might' because the literature on spacetime functionalism is in its infancy; the phrase appears very little outside of my work. Nevertheless, there's an obvious reading that takes it to denote a position wider than my own.

conceptual grasp on spacetime is not exhausted by the role it plays in our theory, what might the extra ingredient be? As Myrvold notes, those who favour a stronger flavour of metaphysics may insist that a functional role fails to capture the *metaphysical* role played by spacetime: that of a container or stage. But I agree with Myrvold that any non-metaphorical content of this claim disappears when we strip away the functional role of spacetime. Intuitions about spacetime's container nature either boil down to thoughts about the kind of functional role that spacetime plays, or to obviously false ideas about preferred reference frames and inapt comparisons to familiar spatial containers.

There is another brand of opponent to be defended against: those who think that identifying spacetime structure is a simple matter of looking for objects with a certain mathematical structure. Here, arguments given in Brown's book help us to see off the challenge.¹⁴ As the passage quoted on p.4 of this paper attests, Brown [6, ch. 8] describes a number of examples of physically useful geometries that are not spatiotemporal, among them velocity spaces, projective Hilbert spaces and the phase space of thermodynamics. Geometry is a flexible tool; with so many non-spatiotemporal uses of geometry there is no question that mere geometrical form could be a sufficient condition for representing spacetime. But nor can we plausibly give any necessary geometrical conditions for spacetime; the geometries of Newtonian and relativistic theories are diverse enough to leave little common ground for such a criterion. And insisting that, for example, spacetime can *only* be represented by a rank 2 metric tensor not only rules out the spacetimes of past theories (*spacetime* in Newtonian theories is not represented by a single tensor), but unhelpfully restricts future theorizing as well.

If there are no plausible non-functional ingredients in the spacetime concept, then we should be spacetime functionalists. Might we simply stop there? There's a long and illustrious history in philosophy of committing to functionalism without articulating the functional role under consideration. But in this case, much of the usefulness of spacetime functionalism comes in its application to individual theories; the idea here is, in part, that spacetime functionalism will allow us to pick out spacetime structure in theories where there is some debate or ambiguity. In order to do this, we need to be explicit about the functional role. I propose that the spacetime role is played by whatever defines a structure of local inertial frames. I'll say a bit more about this, and about the definition of an inertial frame shortly, but, first, let's look at the considerations in play when thinking about the role that spacetime plays.

Much of Brown's work is directly relevant to the question of defining a role for spacetime structure. In particular, two key themes emerge from his book that can serve as desiderata when seeking a concise way of expressing the spacetime role. First, Brown is concerned with the operational significance

¹⁴Brown does not use these examples for the same purpose I do here; instead he's concerned with the explanatory status of spacetime geometry.

of the spacetime metric; the spacetime role had better ensure that the behaviour of rods, clocks, light rays and test particles appropriately (if not exactly) reflects the metric structure. Second, and relatedly, he notes that ensuring such operational significance is a matter of dynamics.

Brown's own account is aimed at establishing the 'chronogeometric significance' of the metric: he is silent on the subject of whether any object with chronogeometric significance must be spatiotemporal. But we may take his account and apply to the concept of spacetime itself: in order for any structure or object to count as spacetime, it had better be related in the right kind of way to the behaviour of physical systems (and particularly to those physical systems that measure spatiotemporal variables); securing this relation will involve postulating an appropriate connection to the dynamics. What kind of connection to the dynamics? At the very least, the spacetime role had better be such that the symmetries of the dynamical laws match the symmetries of the structure that we take to be spacetime. This idea falls out of Brown and Pooley's thoughts about the importance of Lorentz covariance in ensuring appropriate behaviour of clocks and rods, but it of course it also has an older provenance. The idea that spacetime symmetries and dynamical symmetries must be appropriately coordinated is sometimes called *the Earman p prescription* because it is particularly clearly expressed in John Earman's 1989 book [10, p.46]. Given its methodological importance in spacetime foundations, whatever role we assign to spacetime ought to imply the Earman Prescription.

Might we simply conjoin the considerations above in order to come up with a suitable functional role? Perhaps; I have no doubt that there is more than one way to pick out the role played by spacetime in existing theories. Good formulations might agree on our standard spacetime theories, but need not agree in their attributions of spacetime structure to all theories; in my view disagreements over the interpretation of spacetime theories should precisely turn on disagreements over the functional role. But a cobbled together conjunctive definition seems undesirable. We would ideally like a formulation that entails appropriate phenomenological behaviour without requiring us to model the behaviour of complex systems. In addition, application of the Earman prescription is, in fact, not possible unless we first understand which variables are spatiotemporal: there are some symmetries of the dynamics, such as the gauge symmetry of the vector potential in electromagnetism, that we take to be unrelated to spacetime structure and not subject to the Earman prescription. We sometimes say that the prescription applies only to external symmetries, such as the symmetries of relativistic dynamics under Lorentz transformations, and not to internal symmetries like the aforementioned electromagnetic gauge symmetry. Arguably, identifying the external symmetries of the dynamics requires a prior handle on spacetime structure.

So we should look for some simpler way of expressing the functional role that nonetheless ensures an appropriate link to the dynamics, including the

external dynamical symmetries. Here, Brown's work once again rides to the rescue. Recall Brown's formulation of the strong equivalence principle:

There exist in the neighbourhood of each event preferred coordinates, each called locally inertial at that event. For each fundamental non-gravitational interaction, to the extent that tidal gravitational effects can be ignored, the laws governing the interaction find their simplest form in these coordinates. [6, p.169]

According to Brown, the fact that the metric couples to the matter fields in such a way as to make this principle true is the key to its chronogeometric significance. But in a functionalist setting, we can take this further, and say that the metric *is* spacetime precisely because it plays this role in our theory. By defining a structure of local inertial frames in the way described by the strong equivalence principle, the metric succeeds in filling the desiderata set for spacetime in the discussion above: the local coupling ensures that the local symmetries of the dynamics coincide with the local symmetries of the metric, and hence ensure that the metric governs the behaviour of rods and clocks which obey those dynamical laws. Given that both the Minkowski metric and the affine structure of Newtonian theories serve to define a structure of inertial frames, the definition seems to have wide-ranging applicability. Moreover, in relativistic theories, inertial structure fixes projective and conformal structure, and hence metrical structure¹⁵ so the definition does justice to the full geometrical significance of the theory.

Let us say then, that spacetime is whatever serves to define a structure of inertial frames, where inertial frames are those in whose coordinates the laws governing interactions¹⁶ take a simple form (that is universal insofar as curvature may be ignored), and with respect to which free bodies move with constant velocity.

There is, of course, much packed in here to the notion of 'a simple form'. As previously mentioned, and discussed in [23], simplest form here need not mean special relativistic form. Some higher order equations may involve curvature terms even in their simplest form, which is why we can only claim that the form of such equations is universal to the extent to which curvature terms may be ignored. But this does not mean that the notion of a 'simplest form' is contentless: first order equations will take their special relativistic form in in-

¹⁵Projective and conformal structure in fact only fix the metric up to a global scale factor. But there is no need for this kind of conceptual functional definition to pick out every feature of the metric field. In addition, the empirical significance of a global scale factor in relativistic theories is far from clear.

¹⁶In the case of general relativity, one might specify that this applies to the non-gravitational interactions, but there's a sense in which this isn't necessary, if we take interactions to be associated with the fundamental forces. Once we acknowledge the spatiotemporal status of the metric field, gravity no longer classifies as an interaction in the relevant sense.

ertial frames, and higher order equations involving curvature terms will take a poincaré invariant form in these frames.

This gives us a ‘thick’ definition of an inertial frame that builds in a connection to the dynamics, and to the dynamical symmetries, as well as to a notion of natural motion that has historically been central to ideas about space and time. It thus captures the concept of spacetime used in contemporary theories well. Concepts are slippery things, however, and prone to changing over time. It’s been suggested to me, both by Wayne Myrvold and by Harvey Brown that I should appeal to a thinner definition and only talk about inertial motion rather than inertial frames, or perhaps only about natural motion. Such a move would allow more things to come under the umbrella of my functionalist definition; Myrvold, in particular, is concerned about the status of Aristotelian theories. According to my analysis, inasmuch as Aristotelian theories lack inertial structure, they lack spacetime structure.

But there are costs as well as benefits to a thinner definition. Much of my work on functionalist themes involves putting this criterion to use to solve apparent underdetermination problems; I’ve argued elsewhere [14][15][16] that, although both Newtonian gravitation and teleparallel gravity are usually formulated in such a way that their explicit geometry is flat, both really involve a curved spacetime. These arguments depend on the idea that the explicit flat connection in these theories fails to define full inertial structure. But only a sufficiently rich notion of inertial structure can play this role. If spacetime structure only requires defining natural motions, then by redefining which class of bodies we take to be force free, we can easily generate alternative spacetime structures. But describing force laws that couple to the new geometry in an appropriate way is a much trickier business; the thick notion of inertial structure mentioned here is capable of blocking conventionalist moves.

5 Dynamical Relativity as Spacetime Functionalism

Much more needs to be said in defense of the position above, but my main purpose here is to understand the connection between these ideas and Brown’s own. As a matter of intellectual history, the kind of spacetime functionalism I advocate owes a deep debt to Brown’s work. It arose directly from my attempt to square the cogency and interest of the kinds of questions that Brown asks with the conviction that general relativity was a substantivalist theory in a fairly straightforward way. And the debt, of course, goes deeper than just reading Brown’s work: not only did Harvey teach me most of what I know about the foundations of spacetime, but it was Harvey who prompted me to think about teleparallel gravity, the theory that first made me think that identifying spacetime structure might involve more than looking at the explicit geometrical form of a theory. Given the importance of the inspiration, one might won-

der whether my version of spacetime functionalism just *is* Brown's dynamical relativity read in a certain way.

It is clear, of course, that Brown does not think of his theory as giving a breed of functionalism. The word 'function' appears in his book primarily in its mathematical sense, and explicit discussion of the nature of the concept of spacetime is not Brown's style. Moreover, some things that could go under the title of spacetime functionalism are anathema to Brown; he's unlikely to engage in debates about the metaphysical status of spatiotemporal relations in the context of wavefunction realism anytime soon. Furthermore, the scope of Brown's project is different from that of the functionalist. Brown is concerned to interpret relativistic theories; the functionalist is concerned to provide and analysis of spacetime that might apply much more broadly across both pre and post-relativistic theories.

But, despite these explicit differences, it should be clear from the previous section that the kind of spacetime functionalism proposed rests crucially on arguments made by Brown; the idea that *all there is* to being a spacetime is playing a role in defining inertial structure is right there in quotes like the following:

...it is of course non-trivial that inertia can be given a geometrical description, and this is connected with the fact that the behaviour of force-free bodies is universal: it does not depend on their constitution... It is more natural in SR to consider the 4-connection as a codification of certain key aspects of the behaviour of particles and fields. [6, p.142]

or

It is because of minimal coupling and local Lorentz covariance that rods and clocks, built out the matter fields which display that symmetry, behave as if they were reading aspects of the metric field and in doing so confer on this field a geometric meaning.[6, p.176]

Like the spacetime functionalist, Brown sees his position as being helpful for the emergence of spacetime in theories of quantum gravity:

... the notion that the $g_{\mu\nu}$ field is the fabric of space-time, rather than a field in space-time, may be misleading. It may serve to hinder recognition of the possibility that Einstein gravity is an emergent phenomenon...[7, p.9]

But despite having much in common with the spacetime functionalism described above, the first part of this quote reminds us of the difference between

the views. For the spacetime functionalist, the $g_{\mu\nu}$ field *is* spacetime. Given that, from a realist perspective, and considering GR as a fundamental theory,¹⁷ $g_{\mu\nu}$ is undoubtedly a real field, and neither reducible nor eliminable, the view is committed to the idea that spacetime is part of GR's ontology.

By my lights this makes the spacetime functionalist a substantivalist with respect to general relativity. This is because I take any theory in which spacetime is part of the fundamental ontology to be a substantivalist one. This is, of course, 'substantivalism-lite'; it is substantivalism stripped of the containment metaphor. But the containment metaphor's content was always illusory; simple realism about spacetime provides the only defensible substantivalism.

Myrvold, I think, disagrees; he claims that the conceptual view (of which spacetime functionalism is a variant) is neither relationist nor substantivalist. This seems to me right in the case of special relativity; the functionalist thinks that spacetime is a 'thing' just in case a 'thing' fills the spacetime role. In special relativity, there is an independent debate to be had about whether the Minkowski metric is a 'thing', or is reducible to something else. Brown-style relationism is a live option for the functionalist here; it may be that the structure that plays the spacetime role is nomological.¹⁸

But again, in the case of GR, I fail to see how belief in a real field that is spacetime is not belief in a real spacetime. And here, there is a tension between those parts of Brown's work that appear to endorse Rovellian relationism and the functionalism proposed. For Brown, $g_{\mu\nu}$ spatiotemporal issue is not the main point; as noted earlier, Brown is silent on the issues of whether a field with chronogeometric significance is therefore spacetime. Moreover, even if Brown's position falls short of Rovellian relationism, Brown is certainly anti-substantivalist, and the spacetime functionalism in discussion here implies substantivalism in the case of general relativity.

It is possible here that we could interpret Brown as having a thicker conception of substantivalism than the functionalist; in the quote above, he insists that $g_{\mu\nu}$ is not *the fabric of* spacetime. Myrvold also denies that the term 'substantivalism' should be applied to any realist conception of spacetime. It may be that Brown assumes that any belief in the reality of space brings with it the baggage of substantivalist container metaphors. If this were the case, his anti-substantivalism, if not his GR relationism, might be very much in line with functionalist concerns.

But it is hard to see this disagreement as entirely semantic. As Myrvold

¹⁷Acknowledging that GR is almost certainly an emergent theory muddies the water here. I am inclined to think it makes sense to be an emergentist and a substantivalist about spacetime in a certain sense, because I think emergent entities should be admitted to our ontology. Others will disagree.

¹⁸Myrvold also seems to think that what I am calling 'Brown-style relationism' here isn't really relationism, because if we're not Humean about laws, then symmetries of the laws aren't reducible to facts about actual matter.

also notes, a reading of Brown that emphasises the analysis of the spacetime concept is in tension with Brown's interest in questions about spatiotemporal explanation. If spacetime just *is* inertial structure, then the fact that spacetime reflects relevant symmetries of the dynamics is an analytic truth. It makes no sense to ask whether one explains the other. The spacetime functionalist here should agree with Pablo Acuña [20] that Lorentz covariance and Minkowski structure are two sides of the same coin, and the analytic connection between them does not allow for sensible arguments about the arrow of explanation.

In short, a functionalist reading of Brown requires us to ignore parts of the project, namely exactly the features described in section two that have traditionally supported a relationist reading of his work. At the same time, the idea that spacetime is conceptually intertwined with inertial structure, an idea at the heart of the functionalism described above, is inherited very much from Brown's work.

6 Conclusions

How can we square these two extensions of Brown's work? I am not sure that we need try to render them compatible. The considerations advanced here, alongside those raised Myrvold and Acuna, reflect an increasing awareness that there are two strands present in Brown's work. On the one hand, Brown advances arguments for an interesting and novel form of relationism in special relativity. On the other, he develops a subtle picture of the spacetime concept, one in which claims about spacetime are tied intimately to claims about inertial structure. It is increasingly clear that a great deal can be gained by disentangling these two very different contributions to the foundations of spacetime theories. The second strand lays the groundwork for an analysis of the concept of spacetime itself, and it is on this groundwork that a kind of spacetime functionalism may be built.

7 Acknowledgements

Needless to say, this paper owes its greatest debt for its ideas to Harvey Brown, and I'm also grateful for comments on an earlier draft. Discussions with Oliver Pooley over many years have also had a considerable impact on the views presented here. I'm also grateful to Wayne Myrvold for helpful correspondence, and for sharing an early draft of his contribution to this volume, and to James Read and Tushar Menon for extensive comments on an earlier draft.

References

- [1] David Albert. Wave function realism. In Alyssa Ney and David Albert, editors, *The Wave Function: Essays in the Metaphysics of Quantum Mechanics*. New York: Oxford University Press, 2013.
- [2] David Albert. *After Physics*. Cambridge: Harvard University Press, 2015.
- [3] Yuri Balashov and Michel Janssen. Presentism and relativity. *British Journal for the Philosophy of Science*, 54:327–346, 2003.
- [4] Harvey Brown. Michelson, fitzgerald and lorentz: the origins of special relativity revisited. *Bulletin Delelott Louisiana Societe des Sciences et des Lettres Delelott Lodz*, LIII:23–35, 2003.
- [5] Harvey Brown. Einstein’s misgivings about his 1905 formulation of special relativity. *European Journal of Physics*, 26:S85–S90, 2005.
- [6] Harvey R. Brown. *Physical Relativity: Space-time structure from a dynamical perspective*. OUP, 2005.
- [7] Harvey R Brown. The behaviour of rods and clocks in general relativity, and the meaning of the metric field. *arXiv preprint arXiv:0911.4440*, 2009.
- [8] Harvey R. Brown and Oliver Pooley. The origin of the spacetime metric: Bell’s ‘lorentzian pedagogy’ and its significance in general relativity. In Craig Callendar and Nick Huggett, editors, *Physics meets Philosophy at the Planck Scale*. CUP, 2001.
- [9] David Chalmers. *Constructing the World*. Oxford University Press, 2012.
- [10] J. Earman. *World enough and space-time*. MIT Press Cambridge, Mass, 1989.
- [11] Laura Feline. Scientific explanation between principle and constructive theories. *Philosophy of Science*, 78:989–1000, 2011.
- [12] Matthias Frisch. Principle or constructive relativity? *Studies in History and Philosophy of Modern Physics*, 42:176–183, 2011.
- [13] Michel Janssen. Drawing the line between kinematics and dynamics in special relativity. *Studies in History and Philosophy of Modern Physics*, 40:26–42, 2009.
- [14] Eleanor Knox. Newton-cartan theory and teleparallel gravity: the force of a formulation. *Studies in History and Philosophy of Modern Physics*, 42:264–275, 2011.
- [15] Eleanor Knox. Effective spacetime geometry. *Studies in History and Philosophy of Modern Physics*, 44:346–356, 2013.

- [16] Eleanor Knox. Newtonian spacetime structure in light of the equivalence principle. *British Journal for the Philosophy of Science*, 65:863–880, 2014.
- [17] David B Malament. A remark about the “geodesic principle” in general relativity. In *Analysis and Interpretation in the Exact Sciences*, pages 245–252. Springer, 2012.
- [18] Tim Maudlin. *Philosophy of Physics: Space and Time*. Princeton University Press, 2012.
- [19] Wayne Myrvold. How could relativity be anything other than physical? *Studies in History and Philosophy of Modern Physics*, 2017.
- [20] Pablo Acuña. Minkowski spacetime and lorentz invariance: The cart and the horse, or two sides of the same coin? *Studies in History and Philosophy of Modern Physics*, 55:1–12, 2016.
- [21] Alyssa Ney. Finding the world in the wave function: Some strategies for solving the macro-object problem. *Forthcoming in Synthese*, 2017.
- [22] Oliver Pooley. Substantivalist and relationist approaches to spacetime. In Robert Batterman, editor, *The Oxford Handbook of Philosophy of Physics*, pages 522–586. Oxford University Press, 2013.
- [23] James Read, Harvey Brown, and Dennis Lehmkuhl. Two miracles of general relativity. *Unpublished*, 2017.
- [24] Carlo Rovelli. Halfway through the woods: contemporary research on space and time. In John Earman and John Norton, editors, *The Cosmos of Science*, pages 180–223. University of Pittsburgh Press, 1997.
- [25] James Owen Weatherall. On the status of the geodesic principle in newtonian and relativistic physics. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 42(4):276–281, 2011.