Spacetime Structuralism or Spacetime Functionalism?

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

DRAFT - COMMENTS WELCOME
PLEASE DO NOT QUOTE OR CITE

10th August 2014

Abstract

I examine some currently popular articulations of ontic structural realism with respect to spacetime *points*, and find them lacking. Instead, I propose functionalism about spacetime itself (with at most derivative consequences for its points). This has much in common with structuralist positions, but has the virtue of usefully applying to emergent spacetimes.

Introduction

Discussions of ontic structural realism (OSR) are all the rage in the philosophy of physics literature. Proponents claim that the conceptual challenges of modern physics necessitate a radical break with the sometimes-scholastic and intuition-driven approaches that dominate contemporary metaphysics. Opponents claim to be befuddled by the sheer number of arguments and positions adopted by the structural realist, and argue that there's no coherent notion of structure to be found.[1] I have some sympa-

[1] James Ladyman and Don Ross's book ([8]) and Kyle Stanford's response to it [17] provide a nice example of the debate.

thy with both groups, especially when it comes to the topic of spacetime. Both general relativity, and, more starkly, theories of quantum gravity involving emergent spacetimes render certain debates obsolete. At the same time it is hard to pick a single, coherent, and truly original position from among the bewildering array of 'spacetime structuralisms' on offer.

This paper advocates spacetime functionalism as a cure for some of the confusion caused by the structuralist hordes. Inasmuch as this simply introduces yet another kind of structuralism (and indeed, functionalism can readily be thought of as a variant of structuralism) into the debate, I apologize. But I'll argue that spacetime functionalism has a great virtue: it is useful. It will not only help us to deal with new and exotic theories involving emergent spacetime(s), but it will also shed light on old debates, such as the debate over the status of the metric field in general relativity. In part, its usefulness stems from the fact that it focusses on spacetime, rather than spacetime points, and thus doesn't prejudge whether spacetime *has* points at all.

The first half of this paper focusses on the most dominant voice in the literature on spacetime structuralism, ontic structural realism about spacetime points. OSR of this kind is often motivated with reference to arguments over substantivalism and relationism, in particular the hole argument. I'll review these motivations along with the hole argument and its solutions, and argue that the positions that result are, where coherent, too weak to be interesting alternatives to more conventional options.

However, our tour of the hole argument will bring into focus another related and important debate over the status of the metric in general relativity: much of the debate between the substantivalist and the relationist concerns whether to classify the metric as matter or spacetime (or, to use a metaphor that I'll later reject as unhelpful, contents or container). I'll argue that the criteria for such a classification must be functional; the metric field in GR (or some entity in a different theory) is spatiotemporal by virtue of the role it plays in a theory, rather than by dint of its intrinsic properties. I'll briefly sketch my preferred way of identifying the spacetime role, and then examine some ways in which the definition can be put to use.

But before beginning our journey through spacetime matters, it will be helpful to say a bit about structuralism and functionalism more generally. I take structuralism to be a loose umbrella term for a family of positions that uphold a central role for *structure*. (Structure itself is of course notoriously ill-defined, but will be applied clearly enough in the approaches under consideration of the body of this paper). Under the structuralist umbrella fall mathematical structuralism, epistemic structural realism and ontic structural realism; we'll shortly see that functionalism also has a structuralist role to play.

Ontic Structural Realism (OSR) is itself a catch all term, and the views it includes are often less clear than one might wish, but I'll take OSR to be a to be a position about *fundamental* physical ontology that gives a prominent role to structures. Ontic structural realists are prone to making a range of pronouncements. Sometimes they claim that structures are 'on a par ontologically' with objects, sometimes that objects lack intrinsic properties. Much discussion focusses around identity conditions - objects are often said to be individuated by their positions in a structure. Needless to say, not every interlocutor finds all of these statements comprehensible, and I'll try to do better when we come to the spacetime case.

The second kind of structuralist position under consideration here will be functionalism (not necessarily about mental states, but functionalism more generally). Inasmuch as functionalism about a property involves identifying that property with a place in a structure, there's clearly some reason to place functionalism under the structuralist umbrella.[2] But does it have anything to do with OSR as described above? OSR is a general thesis about fundamental ontology. Moreover, it concerns objects rather than properties. By contrast, functionalism is usually adopted with respect to merely some subset of higher-level properties. And there's nothing in standard functionalism to say that the realizer of a functional property has to be a structure. Of course, there's nothing to say that the realizer of the property shouldn't be a structure either, but at this stage OSR and functionalism with respect to particular properties seem to be independently adoptable.[3]

[2]Stewart Shapiro puts it thus:

> Functionalism is an in re structuralism of sorts...In present terms the metaphysical functionalist characterizes a structure, and identifies mental states with places in this structure. In other words, a functional state just is a place in a structure. [16, p.106]

[3]Things get more interesting when we consider the prospects for functionalism about *all* properties, or for all the fundamental ones. This kind of functionalism seems to entail a thoroughgoing denial of *quiddities*, and, although I won't have time to discuss it here, we might think that this is exactly the kind of business the ontic structural realist should

But let us turn to the more localized issue of the nature of spacetime. I'll argue that some of the arguments that have been presented in favour of OSR about spacetime points really push us towards a milder (although compatible!) form of structuralism - spacetime functionalism.

1 Of Points and Holes: Ontic structural spacetime realism

It would be helpful if I could start this section with a clear articulation of ontic structural spacetime realism (OSSR). Unfortunately, as we'll see, and as Hilary Greaves demonstrates in [5], such a thing is easier said than done. In order to get at OSSR, therefore, it will be helpful to come at things via a consideration of the motivation for the position. Here, at least, there is reasonable agreement among authors: OSSR is primarily motivated by the desire to provide an adequate metaphysics for spacetime in light of Earman and Norton's[4]1987 hole argument [4]. At the end of his canonical book, the ever-influential John Earman expresses his desire for an approach to spacetime that is neither traditional substantivalism, nor traditional relationism, but something else:

> My own tentative conclusion from this unsatisfactory situation is that when the smoke of battle finally clears, what will emerge is a conception of space-time that fits neither traditional relationism nor traditional substantivalism. At present we can see only dimly if at all the outlines the third alternative might take. [3, p.208]

Some claim that OSSR provides the longed-for *tertium quid*. But we'll see that the options on the table for OSSR are not up to the job.

1.1 Peering into the hole

Let's quickly review the hole argument. General Relativity is generally covariant. That is, it is written (of necessity!) in geometrical form in

be in.

[4]Of course, in a sense, the hole argument is really Einstein's; it slowed down the development of General Relativity by two years while Einstein grappled with it! But Earman and Norton coined the term and put the argument in its modern form.

such a way that its equations hold true in any coordinate system whatsoever. But this general covariance has a hidden face, for as well as ensuring invariance under passive coordinate transformations, it also ensures invariance under the class of active diffeomorphisms, where a diffeomorphism is a smooth continuous map from the manifold onto itself. A little more formally: GR has models $\langle M, g, O_i...O_n\rangle$, where M is a four-dimensional smooth manifold, g is the metric field that defines a curved geometry on the manifold, and $O_i...O_n$ are matter fields. It turns out that if $\langle M, g, O_1...O_n\rangle$ is a model of a GR theory, then so is $\langle M, d^*g, d^*O_1...d^*O_n\rangle$, where d^* is induced by some diffeomorphism d. These diffeomorphisms correspond to complete rearrangements of matter and metric fields with respect to the manifold.

But these diffeomorphisms can cause trouble. Suppose we now consider some diffeomorphism that is the identity everywhere except spacetime region R ('the hole'). We now have two solutions to the GR equations that agree everywhere except in the hole. No amount of information about the state of the fields outside the hole will determine the assignment of fields to manifold points within the hole. So it looks as if our theory is in trouble - after all, we can always find a diffeomorphism such that any region we're interested in falls in the hole; the theory fails to yield a unique solution for any region in which we don't already know the distribution of field values!

Earman and Norton present this as a particular problem for the substantivalist, but let us for a moment step back from the question of the reality of spacetime points, and merely consider the solution to the formal problem. When faced with apparent indeterminism or underdetermination, we are nowadays apt to ask whether we can think of the apparent difference in solution as a mathematical artefact.[5] In order to do this, we need to think of the diffeomorphically related solutions as representing one and the same physical situation. To assert that both $\langle M, g, O_1...O_n\rangle$ and $\langle M, d^*g, d^*O_1...d^*O_n\rangle$ represent one and the same physical situation is to assert *Leibniz Equivalence*. Is such an assertion defensible on empirical grounds?

The answer, originally provided by Einstein in 1915, is yes. A diffeomorphism acts on both the metric and matter fields in exactly the same way. Diffeomorphically related solutions agree on what Einstein calls 'point-

[5]For an argument that the apparent difference is not even mathematical, see [19].

coincidences'; the arrangement of metric and matter fields do not change relative to one another. And any observable content of the theory depends only on this relative arrangement. So one straightforward way to escape the hole argument is to assert *Leibniz Equivalence*.

Now enter substantivalism and relationism. Earman and Norton claim that only relationists can help themselves to Leibniz equivalence: substantivalists, who are committed to the existence of spacetime points, must view the diffeomorphically related solutions as representing different things happening at the same spacetime point. Relationists, only committed to the reality of the fields, have no such problem.

But interpreting GR as a relationist theory is not straightforward either. GR, as ordinarily interpreted, is a theory about the dynamical geometry *of spacetime*; a naïve interpretation would include spacetime amongst its objects. Of course, naïve interpretations may be wrong. More pressingly, GR is not formulated like a relationist theory; it is committed to quantification over the points of the manifold. The relationist, therefore, should at the very least be committed to an extremely difficult, perhaps impossible, reformulation.

Thus we arrive at the impasse lamented by Earman. But the modern substantivalist is unlikely to accept this; since 1987, quite a number of 'sophisticated substantivalisms' have flourished, providing new ways for the substantivalist to escape the hole. Of greatest interest to us here are those that endorse Leibniz Equivalence. In particular, consider a response first suggested by Carl Hoefer [6] and subsequently developed by Oliver Pooley [10, 11, 12]. Pooley and Hoefer argue that the hole argument only threatens the substantivalist who endorses *haecceitism* with respect to spacetime points. Haecceitism, following David Lewis, is the idea that possible worlds can differ solely in terms of which objects instantiate which properties. The haecceitist with respect to spacetime points thinks that we can make sense of the idea of two different general relativistic worlds differing solely in terms of which spacetime points possess which particular field values. But anti-haecceitism is independently attractive; Lewis endorses it for any object whatsoever! So anti-haecceitism with respect to spacetime points doesn't entail abandoning our commitment to them. Anti-haecceitist substantivalism provides a way out of the hole for the spacetime realist.

My purpose here is not to rehash the extensive literature on the hole argument; there is quite enough of it to be getting on with! So let us simply note here that some of the most standard responses to the hole argument have the substantivalist and relationist agreeing on Leibniz Equivalence. Indeed, as many have commented, in the context of general relativity, substantivalism and relationism seem to share many features. It's exactly this that leads to the conclusion that the debate is outmoded, and a third way is required. Let us see whether OSSR can provide a modern alternative.

1.2 A third route from the hole?

Proponents of OSSR agree that they should subscribe to *Leibniz Equivalence*. What kind of options are available? I can find only two clear propositions in the literature, and we'll consider each in turn.

One option is to claim that sophisticated substantivalism is itself a form of ontic structural spacetime realism. On this view, endorsed by Steve French and Dean Rickles [13], a move away from haecceitism is very much in the spirit of structural realism; after all, haecceities are archetypal 'object-properties'. The structuralist move here is to reject both primitive object individuation (haecceities) and primitive trans-world identity (haecceitism).[6] So OSSR turns out to be just the view that the substantivalists have been forced to anyway! If this is the case, OSSR adds little new to the debate over general relativity, but ontic structural realism garners considerable support from the hole argument literature.

But a closer look suggests that this kind of position is much too weak. We were promised a radical new metaphysics, and we are given a position that any good Lewisian metaphysician endorses anyway! Moreover, anti-haecceitism is itself entirely compatible with a very conventional view of spacetime points with intrinsic as well as relational properties. If this is OSSR, structural realism is just a new name for an old and popular position.

Let us turn to the second option. What can be added to anti-haecceitism

[6]We should note here that, depending on how much one packs into the notion of a haecceity, commitment to haecceities need not entail commitment to haecceitism. One might, for example, endorse a minimal notion of haecceities in order to resist the identity of indiscernibles and ground primitive numerical distinctness, but deny that this minimal notion can ground claims of determinate trans-world identity of a kind that could ground haecceitism. One might here find a way to drive a wedge between OSSR and sophisticated substantivalism, if we assume that both deny haecceitism but that the latter endorses minimal haecceities. But this is a very thin wedge indeed.

to give OSSR a stronger flavour? One claim often made by ontic structural realists is that objects are individuated solely by their placement in a structure. One part of this claim is that they do not possess primitive, non-qualitative properties that ground their individuation - haecceities. Some form of the principle of the identity of indiscernibles (PII) then plays a role, because objects must be individuated by their qualitative properties. For the structural realist, these qualitative properties should be relations. But any attempt to individuate properties via the relations they stand is met by a class of counter-example. The most familiar of these is Max Black's pair of indiscernible spheres. An example more germane to physics-led OSR is a pair of electrons in a singlet state, which plausibly share all their properties and yet possess opposite spin. A final example, more relevant here and proposed by Chris Wüthrich [20] as a challenge to the spacetime structuralist, is a pair of spacetime points in a spatially symmetric solution of the general relativistic field equations. In each of these cases, we have a pair of putative objects that share all their properties, including the relational ones. In each case there is a concern that a proponent of PII should to be identical two objects that are in fact distinct. In the symmetric spacetime case, Wüthrich suggests that some spacetimes will be reduced to a single point.

What should the structural realist say in response? Interestingly, a group of structural realists have made a virtue of necessity, and grounded their version of OSR in a modified version of the PII. Following Saunders [15] call two objects *weakly discernible* just in case they share all their relational and intrinsic properties but bear some symmetric but irreflexive relation to one another. According to Saunders, this is enough to ground their distinctness. And indeed, from a perspective that grants priority to relations, as OSR does, this seems to be the case. There are two of Max Black's spheres rather than one precisely because a certain kind of relation exists.[7] So the way in which one version of OSSR makes a virtue of necessity is that it takes a putative counterexample to the PII and points out that it can be understood to obey a version of the PII just in case we adopt an approach on which relations are more fundamental than objects.

How can this idea be elevated into a full-blown account of Ontic Struc-

[7]Wüthrich [20, p.1048] finds this response to be circular because it presupposes individuation of the objects in question, but misses the structural realist's point, which is just that if we take the properties of relations, including the property of irreflexivity, to be fundamental, then we can ground numerical distinctness in these properties.

tural Spacetime Realism? The best worked out such account, indeed, in my view the most precise articulation of OSSR, is given by Fred Muller [9]. Muller takes OSSR to be committed to four principles or features, which I present here in simplified form, and with some notation modified:

0. *The Structural Characterisation of GR*: GR is characterised by its models $\mathcal{M} = \langle M, g, O_i ... O_n \rangle$.[8]

1. *Principle of Automorphic Properties and Relations*: For a given model $\mathcal{M}$, all genuine properties and relations of the spacetime points of $\mathcal{M}$ must be invariant under automorphisms of $\mathcal{M}$. Call the set of all such properties $\mathfrak{F}$ and all such binary relations $\mathfrak{R}$.

2. *PII**: $(AutAbsInd(p,q) \wedge AutRelInd(p,q)) \rightarrow p = q$, where:

   - $AutAbsInd(p,q)$ iff $\forall F \in \mathfrak{F}(Fp \leftrightarrow Fq)$.[9]
   - $AutRelInd(p,q)$ iff $\forall R \in \mathfrak{R}(\forall r(Rpr \leftrightarrow Rqr) \wedge (\forall s Rsp \leftrightarrow Rsq)$.

3. *Structuralist Representation Thesis*: If some model $\mathcal{M}$ represents the universe, then every model isomorphic to $\mathcal{M}$ also represents the universe.

My issue with the position that results from adopting all four of Muller's principles is not that it is wrong. On the contrary, once we're in the game of analysing the individuation criteria for spacetime points, then this seems to me to be an attractive way to go. But my first real issue is that the spacetime structuralist should not be in this game in the first place. More on this later. The second issue, not unrelated to the first, is that this kind of spacetime structuralism fails to provide either a radical new metaphysics, or indeed, an innovative third way out of the hole argument. In fact its motivations have little to do with GR or the hole argument at all. Let us examine Muller's principles to see how this complaint holds up.

Principle 0 amounts to the identification of the GR (or, perhaps more weakly), its characterisation, by its class of models. I take it this is a common view. Principle 3 amounts to the assertion of *Leibniz Equivalence*, albeit couched in terms that carefully avoid talk of possible worlds. Both

[8] Muller calls the models structures, and includes topological and other geometrical features in the construct, but nothing I say here hangs on this.

[9] The properties in question can be relational properties with only one free variable, so this is intended to capture what is sometimes called the 'weak' PII.

of these are principles to which the sophisticated substantivalist and relationist might happily subscribe.

Principle 1 has more content. The set of automorphisms of $\mathcal{M}$ is the set of maps from the manifold of M to itself that preserve all of the structures of $\mathcal{M}$, that is, the set of maps that preserves the arrangement of *both metric and matter fields*, reflecting the symmetries of the particular model at hand. It is very natural for the spacetime structuralist to restrict permissible spacetime properties and relations to those that are invariant under these transformations; it is exactly this criteria that rules out some additional individuation of spacetime points over and above that provided by the structure of GR itself. It is this rule, for example, that precludes the addition of an individuating property "distance from the center of the universe" in a model that is spatially infinite and homogeneous. It also, needless to say, rules out the introduction of haecceities, distinct and unique intrinsic properties of spacetime points. But this principle does not appear to be motivated by the hole argument, nor by any feature particular to GR. The group of automorphisms reflect the particular symmetries of a given model, not the symmetries of the theory as a whole. (If we choose to think of the symmetry group of GR as a theory is the diffeomorphism group,[10] then the automorphism group of some model is a tiny subgroup of this.) Rejecting properties that might individuate points beyond what is empirically possible has a long and illustrious history; one way of thinking about what is wrong with postulating absolute space in Newtonian theories is that one introduces a property (well-represented by a timelike vector field) that is not invariant under the automorphisms of the theory. So if Principle 1 is independent of both GR's specific formulation and the hole argument, and could be adopted by those of a traditional metaphysical bent, it can hardly provide the essential ingredient in a radical new metaphysics.

Let us turn to Principle 2. This introduces a new criterion of individuation, and hence a new category of object, which Muller calls *relationals*, and deems 'a neglected metaphysical category' [9, p.8]. Relationals are those objects which are absolutely indiscernible, and hence would be deemed identical on the traditional PII, but which this version of OSR deems to

[10]More precisely, since the group of diffeomorphisms itself depends on the manifold: if we choose to think of the symmetry group of GR on a manifold as the diffeomorphism group for that manifold.

be nonetheless individuated by the (irreflexive) relations in which they stands. Max Black's spheres, electrons in the singlet state, and spacetime points in GR models with global symmetries are relationals. Max Black's sphere's are individuated by the irreflexive "two miles from" relation in which they stand. The electrons are individuated by the relation they hold as a result of having total spin 0.[11] This idea, I think, has some claim to novelty. If one is in the business of worrying about how objects are individuated, then there is much to be said for introducing a new option. And, in my view, individuating objects via fundamental features of the relations they enter in to seems a great deal less mysterious than postulating haecceities, however minimal, if only because the properties of relations are rather more familiar than haecceities. So, again, my issue is not that PII* is wrong. The trouble is that it's neither particularly radical, nor motivated by the need to find an alternative to relationism and substantivalism in the light of the hole argument or particular features of GR.

To see that the introduction of PII* has little to do with GR, note that it by no means follows from Muller's account that all spacetime points are relationals. Indeed, in many models of GR, all spacetime points will be absolutely indiscernible. It is only in special, globally symmetric solutions that one finds points that require PII* for their individuation, and such globally symmetric solutions are no more common in GR than in other spacetime theories; indeed, if one thought that such an idea was meaningful, one should think that such models will be much *more* common in theories where spacetime has generic symmetries. It is helpful here to follow Oliver Pooley [11] in pointing out that the situation is very different in quantum mechanics, where symmetries of the theory itself, rather than of a given solution, demand that fermions exist in anti-symmetric states, and hence that certain quantum particles are relationals. It seems quite plausible that quantum mechanics forces on us new ideas about particles, and that some of these ideas are helpfully cashed out in terms of the structural realist's individuation criterion. In the case of spacetime points, it is rather than they were rather pared down objects all along, and PII* simply allows the believer in spacetime points a new option for individuation.

As mentioned, there is another worry at play here. The ontic structural realist claims to endorse a less object-centric ontology than tradi-

[11]The individuating relations for spacetime points in GR are a little more complicated; see [9] for details.

tional views. And yet, consideration of OSSR dragged us ever deeper into discussions of spacetime points![12] One might think that the spacetime-point individuation game is not the arena for a new approach to spacetime well-suited to modern theories, and this worry might seem even more acute once we consider the fact that some contemporary theories of quantum gravity are not formulated on a manifold at all.

1.3 The status of the metric

At this point it is quite clear that the ontic structural realist's promissory note has not been cashed. Rather than new metaphysics, we are given small variations on existing substantivalist positions. But although the above options are the only ones clearly developed in the literature, much more has been said about OSSR and what it hopes to achieve. Let's reexamine the motivations for the position to see if anything can be gleaned.

OSSR seeks a third way between relationism and substantivalism. In light of the hole argument, we have considerable agreement between the two positions; in particular, both agree that we should be committed to Leibniz Equivalence. What, then, is the substance of the debate? In very large part, it seems to boil down to a debate over the status of g, the metric field. The metric field of general relativity is an interesting beast. In many ways, it plays much the same role as, say, the Euclidean spatial metric in Newtonian theories, or the Minkowski metric of special relativity. It defines distances and angles, and gives geodesic structure, and hence the paths of force-free bodies. But it has a new feature: it is dynamical. It enters into the fundamental equations of the theory, curving in response to the presence of matter as dictated by Einstein's field equations. It also has independent degrees of freedom - vacuum solutions involve gravitational waves that carry energy of their own of a certain kind. As Carlo Rovelli [14] famously pointed out, we can smash the rock of gibraltar just as well with gravitational waves as we can with electro-magnetic ones. So the metric field seems to have some properties traditionally associated with spacetime, and some properties traditionally associated with matter fields. The substantivalist insists that it represents spacetime itself, the relationist insists that it should be counted among the matter fields.

[12] I am grateful to David Wallace for pointing out this irony in conversation a few years ago.

Looking at this debate, we might well find ourselves frustrated. The substantivalist and relationist agree that the metric field exists. They agree on its properties and its role in the dynamics. They simply disagree on its classification. So perhaps here is our tertium quid. OSSR is distinguished by the rejection of the traditional categories of matter and spacetime. This seems to be what Mauro Dorato suggests when he comments:

> Having established that the interpretative task raised by the substantivalism/relationism debate essentially involves the metric field, the undeniable ambiguity between matter and space that such a field embodies suggests dropping both substantivalism insofar as it requires empty spacetime and relationism insofar as it requires an antirealist stance about spacetime and embracing a new position as an often invoked tertium quid to the debate, which I will refer to as structural spacetime realism. [2, p.1611-12]

I am sympathetic. But the rejection of traditional categories does not itself distinguish a positive philosophical position; OSSR must find more to say for itself. Furthermore, before rejecting traditional categories, we must take care that we have not thrown out useful theoretical tools. It is clear, in the light of GR, that the neat division entailed by the metaphor of contents (matter) moving in the passive container of spacetime is no longer tenable; spacetime has taken on some of the properties of matter. But does this render the concept of spacetime unable to do useful work in current and future theories? Is the very notion of spacetime otiose?

I think not. Whether or not one finds the container/contents metaphor contentful or illuminating, the notion of spacetime does more work than merely endowing us with a familiar metaphysical picture. To identify a piece of structure as spacetime structure is to tell us something useful about the application and formulation of a theory; if we know that a geometrical field is spatiotemporal, we know that it relates in a particular way to the motions and arrangements of bodies, and to the behaviour of regular periodic processes, and other physical processes capable of underpinning clocks. We also know that its structure will help us to write simple equations in particular coordinate systems, something we must do if we are ever to apply the theory in practice. Identifying spacetime structure is essential to the practice of physics.

At the same time, quantum gravity has thrown up ever greater puzzles in identifying spacetime structure. Both String Theory and Loop Quantum Gravity seem to agree that the spacetime of our macroscopic theorizing is non-fundamental. This means it must emerge in some limit from the fundamental structure, which in turn means that extracting predictions from our theory will depend on recognizing emergent spacetime structure when we see it. Moreover, even at more fundamental levels, spacetime structure, where it exists, resists easy identification. The kind of duality proposed by the AdS/CFT conjecture, and echoed even in non-string theoretic contexts, suggests the existence of equivalent descriptions of reality that appear to posit spacetimes of different dimensionality. An understanding of what spacetime is and what it does will be important in the analysis of these problems and may therefore play a role in the development of a truly fundamental theory.

My proposal is, therefore, to abandon the containment metaphor, but get on with the important business of giving an analysis of spacetime. That analysis will be functional. And in this spacetime functionalism we will find at least some of the features that John Earman, and the ontic structural spacetime realist, desire.

2 Spacetime Functionalism

I propose that we identify spacetime not by what it *is*, but by what it *does*. After all, as the very possibility of the relationist position suggests, we don't have any direct access to the intrinsic properties of substantival spacetime. Spacetime is an unobservable entity. But, as David Wallace convincingly argues in his [18] even if it weren't, there are compelling reasons to think of a variety of observable structures as defined by their functional roles. Spacetime is an entity that plays a very important role in our theories. Let us then, identify *being a spacetime* with filling this role. The proposal is clearly for a kind of functionalism.

2.1 Articulating spacetime functionalism

What kinds of considerations go into determining the spacetime role? We are here firmly in the grip of what has been called 'scientific functionalism'; I don't know what a folk-theoretic platitude about spacetime might look like, and I wouldn't trust it if I did. The spacetime role is one that is defined by theoretical practice. But here we have a rich tradition. Spacetime governs the behaviour of the special entities that we deem to be rigid bodies and regular processes usable as clocks. It also, importantly, defines a standard of natural, unaccelerated motion, and thus allows acceleration to be defined relative to this standard. And it tells us which coordinate systems (the inertial coordinate systems) will be the simplest in which to do physics, and how to transform between these coordinate systems.

Interestingly, it turns out that all of the above roles will be filled if we manage to fill the last one; a structure will play the spacetime role in our theories just in case it describes the structure of the inertial frames, and the coordinate systems associated with these. In the interests of brevity, and not getting bogged down in a technical discussion, I won't argue for that claim here.[13] Let's leave the details of how we might abbreviate the functional rile for the moment, and simply note that any abbreviation must ensure the features mentioned in the paragraph above. And any object that plays these kinds of role must relate in a very special way to the rest of physics. In particular, the inertial frames are those frames in which all of the laws of physics take a uniform and particularly simple form. So whatever serves to define these frames has to interact with every part of the rest of the dynamics in a universal and uniform way, and must also bear a strong relation to the symmetry group of the dynamics. Any old geometrical object will not do; in order for a geometrical object to play the spacetime role, it must appear in the dynamics in a very particular way.

Filling the spacetime role thus defined is not a trivial matter. But is it enough? Spacetime functionalism involves a deliberately light metaphysical touch which will doubtless leave many arch-substantivalists feeling unsatisfied. But spacetime functionalism turns out to be the only real option for identifying spacetime structure in new theories. Let us consider the alternatives.

First, we might think that any functionalist analysis will fail to mention the importance of geometry for spacetime. Spacetime, after all, is a geometrical entity! Can't we identify spacetime structure in a theory just by finding the right kind of geometrical object? The metric field in general relativity, for example, turns out to be just the right kind of thing to de-

[13] For those who are interested, it's possible to read Harvey Brown's 2006 book [1] as giving a long argument for my claim.

scribe a variably curved 3+1 dimensional space; isn't it just obvious that it represents spacetime? But on further refection it is still more obvious that geometrical considerations alone don't pick out spacetime structure. The metric field is a rank-2 metric tensor field of Lorentzian signature. But this feature is neither necessary nor sufficient to represent spacetime. In Newtonian theories, spacetime structure is represented not by a single metric field, but by 2 metrics and a covariant derivative operator; nothing like the metric field exists. And, as Harvey Brown has pointed out, there are all kinds of uses of geometrical fields (including rank-2 tensor fields) that have nothing whatsoever to do with spacetime (fibre bundles, velocity spaces and phase spaces spring to mind; for more examples, see [1, pp.154-160]). Geometrical form is not a good way to identify spacetime structure, except insofar as geometrical objects prove to be suitable fillers of the spacetime role.

What other considerations might we bring to bear? One might want to argue that we have a handle on a notion of spacetime that is at the very least prior to the kind of functional role described above. Spacetime is, after all, the arena in which events play out, the container for matter, and the background for the rest of our physics. Surely we can imagine a spacetime that fails to relate in quite the right way to matter, or a theory in which some imposter serves to determine inertial structure while another entity really serves as the container?

These moves are, I suppose, the analog of qualia-based arguments against functionalism with respect to mental states. Just as the believer in Zombies, or the proposer of homunculus-head thinks that states could fill the functional role of mental states, but nonetheless be missing something, the proponent of the container thinks that spacetime functionalism fails to capture the essential nature of spacetime. But here the spacetime philosopher is in a far worse boat than the philosopher of mind. There is no argument that we have privileged access to the non-functional features of spacetime, whatever they might be. Where the fan of qualia has introspection, the fan of the container has only metaphor.

3 Conclusions: Applying spacetime functionalism

How does the position above impact on the debate between the substantivalist and the relationist? As it turns out, in the case of general relativity, it rules rather conclusively in the substantivalist's favour. For the metric field of general relativity plays the subtle spacetime role to perfection; it turns out to couple to matter (via the 'minimal coupling prescription')[14] in just such a way as to ensure that objects built of matter fields conform to the geometry of the metric field. So on the functionalist account, the metric field is spacetime, whatever we might say about the matter-like properties of gravitational waves. At the same time, on any realist approach to GR that considers it as a fundamental theory, the metric field is very much part of the ontology. Substantivalism is vindicated.

This means that, in the context of GR, spacetime functionalism is not quite the tertium quid that proponents of OSSR desire. But it does do justice to at least some structuralist instincts. For one thing, as I pointed out in the introduction, functionalism is itself a structuralism of sorts. For another, it allows us to cast off what James Ladyman and Don Ross [8] call 'the containment metaphor'. The substantivalism that emerges from the hole argument on the functionalist account is leaner and meaner than traditional substantivalism, insofar as traditional substantivalism has been tangled up with the containment metaphor, so John Earman's desire is satisfied. Most importantly, functionalism is a flexible tool, and thus promises to be at least compatible with, if not demanding of, the radical new metaphysics that the structural realist believes us to require.

Consider, for example, what physicists call theories of emergent spacetime.[15] Even in the absence of a theory of quantum gravity, we have good reason to think that GR will be an effective, not fundamental theory. Quite general considerations lead to the thought that fundamental structure must, at the very least, be discrete. A famous and intriguing paper

[14] The minimal coupling prescription, roughly speaking, tells us how to rewrite the equations for matter fields in curved spacetime, and thus how the non-gravitational interactions interact with the curved geometry of GR. It's sometime's summarised as the "comma goes to semi-colon" rule, which replaces derivatives with covariant derivatives in the relevant equations, but in fact it's rather more subtle than this. See [1, p.?] for more details.

[15] Philosophers may want to quibble over the term emergent.

by Ted Jacobsen [7] claims to derive the field equations of GR from thermodynamic considerations, and thus takes them to be equations of state, rather than fundamental laws. Functionalism, it seems, can cope with this. There is nothing in the functional definition of spacetime to suggest that the realizer of the spacetime role must be fundamental; if an effective field fills the role in the macroscopic realm, it is no less a spacetime for doing so.

Equally, there is nothing in the spacetime role to suggest that the role can't be multiply instantiated. One kind of thing might realize the role with respect to one level of description while another kind of thing realizes the role with respect to another. If the AdS/CFT conjecture is correct, it may be that certain physical situations can be described using spacetimes of differing dimensionality. Traditional metaphysics would put these in competition with one another, but functionalism may have the ability to reconcile them.

Granted, there will be much difficult metaphysical work to be done in understanding the new physics, if and when it is established. It may be that some kind of interesting, substantive, OSSR can be articulated once we move away from classical theories. Nothing I have said here entails that the OSSR is false; I only claim that its proponents are incorrect in thinking that it fills the role of Earman's third alternative. For the time being, the structural realist's desire to accommodate a radical new metaphysics seems perhaps best served by putting older, less radical ideas about functionalism to new uses.

References

[1] Harvey R. Brown. *Physical Relativity: Space-time structure from a dynamical perspective*. OUP, 2005.

[2] Mauro Dorato. Substantivalism, relationism, and structural spacetime realism. *Foundations of physics*, 30(10):1605–1628, 2000.

[3] J. Earman. *World enough and space-time*. MIT Press Cambridge, Mass, 1989.

[4] J. Earman and J. Norton. What Price Spacetime Substantivalism? The

Hole Story. *The British Journal for the Philosophy of Science*, 38(4):515–525, 1987.

[5] Hilary Greaves. In search of (spacetime) structuralism. *Philosophical Perspectives*, 25(1):189–204, 2011.

[6] Carl Hoefer. The metaphysics of space-time substantivalism. *The Journal of Philosophy*, 93(1):5–27, 1996.

[7] T. Jacobson. Thermodynamics of spacetime: the Einstein equation of state. *Physical Review Letters*, 75(7):1260–1263, 1995.

[8] J. Ladyman and D. Ross. Every Thing Must Go: Metaphysics Naturalised, 2007.

[9] Fred A Muller. How to defeat wüthrichs abysmal embarrassment argument against space-time structuralism. *Philosophy of Science*, 78(5):1046–1057, 2011.

[10] Oliver Pooley. *The Reality of Spacetime*. Unpublished DPhil thesis, University of Oxford, 2002.

[11] Oliver Pooley. Points, particles and structural realism. In French Rickles and Saatsi, editors, *The Structural Foundations of Quantum Gravity*. OUP, 2007.

[12] Oliver Pooley. Substantivalist and relationalist approaches to spacetime. In Robert Batterman, editor, *The Oxford handbook of Philosophy of Physics*. OUP, forthcoming.

[13] Dean Rickles and Steven French. Quantum gravity meets structuralism: Interweaving relations in the foundations of physics. In *The structural foundations of quantum gravity*, pages 1–39. Oxford University Press, USA, 2006.

[14] Carlo Rovelli. Half way through the woods. *The cosmos of science*, pages 180–223, 1997.

[15] S. Saunders. Physics and Leibnizs Principles. *Symmetries in Physics: Philosophical Reflections*, pages 289–308, 2003.

[16] Stewart Shapiro. *Philosophy of mathematics: Structure and ontology*. Oxford University Press, 1997.

[17] K Stanford, Paul Humphreys, Katherine Hawley, James Ladyman, and Don Ross. Protecting rainforest realism. *Metascience*, 19:161–185, 2007.

[18] David Wallace. *The emergent multiverse: Quantum theory according to the Everett interpretation*. Oxford University Press, 2012.

[19] James Owen Weatherall. Regarding the hole argument. Forthcoming.

[20] Christian Wüthrich. Challenging the spacetime structuralist. *Philosophy of Science*, 76(5):1039–1051, 2009.