

Thinking About Spacetime

DAVID YATES

Centro de Filosofia, Faculdade de Letras, Universidade de Lisboa

Abstract

Several different quantum gravity research programmes suggest, for various reasons, that spacetime is not part of the fundamental ontology of physics. This gives rise to the problem of empirical coherence: if fundamental physical entities do not occupy spacetime or instantiate spatiotemporal properties, how can fundamental theories concerning those entities be justified by observation of spatiotemporally located things like meters, pointers and dials? I frame the problem of empirical coherence in terms of entailment: how could a non-spatiotemporal fundamental theory entail spatiotemporal evidence propositions? Solutions to this puzzle can be classified as realist or antirealist, depending on whether or not they posit a non-fundamental spacetime structure grounded in or caused by the fundamental structure. These approaches place different constraints on our everyday concepts of space and time. Applying lessons from the philosophy of mind, I argue that only realism is both conceptually plausible and suitable for addressing the problem at hand. I suggest a role functionalist version of realism, which is consistent with both grounding and causation, and according to which our everyday concepts reveal something of the true nature of emergent spacetime.

1. Introduction

Several different quantum gravity research programmes suggest that spacetime is not part of the fundamental ontology of physics. The reasons vary from case to case, as does the extent to which they suggest a non-spatiotemporal fundamental ontology.¹ But the central idea, and all I need to set up the arguments of this paper, is that the fundamental entities posited by the theories in question seem to inhabit spaces whose structures are radically different from the familiar spacetime of general relativity (GTR). In loop quantum gravity (LQG), for instance, which is obtained by applying canonical quantization procedures to the Hamiltonian formulation of GTR, it's very hard to identify spatiotemporal properties like locality. In LQG, the fundamental structures are spin networks, which are represented as graph structures in

¹ See Huggett & Wüthrich (2013) for details of these cases, graded according to how radically non-spatiotemporal their fundamental ontologies seem to be.

which the nodes and edges are associated with quantum spin states. These spin states are eigenstates of operators that can be interpreted as area and volume operators, and which have discrete spectra, so the fundamental LQG space is not the smooth differentiable manifold of GTR. Localization in LQG spacetime must be understood in terms of interaction between other quantum fields and the gravitational field as described by LQG, rather than in terms of position in a background spacetime, as in standard quantum field theory.² These properties of LQG—the discrete nature of LQG spacetime and the interactive theory of localization it requires—are certainly radical departures from a continuous background spacetime, but they don't suggest the disappearance of spacetime in LQG. They suggest that spacetime is fundamentally discrete and dynamic.³ However, other aspects of LQG are more problematic.

The nodes of spin networks can be interpreted as quanta of volume and the edges as quanta of area, with adjacency relations definable within a given spin network. However, adjacent nodes in a spin network map onto areas that can be arbitrarily far apart in general relativistic spacetime. Hence, there is no relation corresponding to the classical locality relation in the fundamental LQG structure, and no clear way of identifying the properties and relations of LQG with ordinary spatiotemporal properties and relations. Furthermore, when we factor in that the fundamental space ought to be in a *superposition* of spin network states, rather than an eigenstate thereof, the problem becomes more pressing still, as the terms of the superpositions have different structures, so what is local in one need not be local in another. These problems are jointly known as the problem of *disordered locality*.⁴ As has been noted, disordered locality in LQG seems to preclude straightforward spatiotemporal interpretations of LQG operators,

² Rovelli & Vidotto (2015).

³ See for instance Huggett & Wüthrich (2013), pp. 277-8.

⁴ Following Markopoulou & Smolin (2007).

according to which the nodes and edges in spin networks are quanta of *spatial* volume and area respectively, and adjacency in a spin network is *spatial* contiguity.⁵

Similar issues arise in string theory, in virtue of T-duality. T-dual theories are defined on *compactified* spaces, in which at least one dimension has the topology of a circle, but where the radius assigned by one of the pair of theories is the reciprocal of that assigned by the other.⁶ T-duality is a mapping that preserves all observables, but such that each of the pair of theories attributes a different radius to a compactified dimension of the space—‘target space’—in which the string theoretic dynamics play out. Philosophers of physics typically embrace the *common core* approach, according to which the physical content of T-duals should be seen as that which is common to both theories.⁷ The physical geometry of target space should thus be regarded as indeterminate. Nick Huggett argues on this basis that target space can’t be identified with *phenomenal* space, the space we take ourselves to occupy and about which GTR is our best theory.⁸ GTR spacetime has a *determinate* geometry, and so can’t be identified with target space, because nothing can have both a determinate and an indeterminate geometry. As with LQG, there’s good reason to think that the fundamental ontology of string theory is non-spatiotemporal, in the sense that its dynamics play out in a structure that can’t be identified with GTR spacetime.

1.1. The Problem of Empirical Coherence

The apparent disappearance of spacetime in quantum theories of gravity gives rise to a problem of *empirical coherence*. To a first approximation, the problem can be stated as a question: if the elements of the fundamental ontology don’t occupy spacetime, then how can fundamental

⁵ Huggett & Wüthrich (2013); Wüthrich (2017); Lam & Wüthrich (2018).

⁶ See Huggett & Wüthrich (2013), pp. 280-1 for an informal account.

⁷ See for instance Matsubara & Johansson (2018); De Haro (2018); Le Bihan & Read (2018).

⁸ Huggett (2017).

theories be justified by observation of things that do, such as meters, pointers and dials—the so-called “local beables” of the ontology?⁹ This problem has already been much discussed in the context of wavefunction realism, and it will be useful to frame the present debate in the terms already established there. Briefly, there are good grounds for treating the wavefunction as a physical field, but it can’t be a field on 3D space, because there aren’t enough degrees of freedom there to distinguish between distinct quantum states.¹⁰ Rather, the wavefunction must inhabit configuration space, which has at least $3N$ dimensions, where N is the number of particles in the universe. None of these $3N$ dimensions corresponds to any of the 3 dimensions of space we intuitively take our world to have, so if configuration space is fundamental, how can there be local beables? I shall frame the problem of empirical coherence as follows:¹¹

1. A fundamental theory is empirically coherent only if it entails evidence propositions.
2. Evidence propositions attribute spatiotemporal properties to local beables.
3. The fundamental ontology of physics is non-spatiotemporal.

I assume for the purposes of this paper that (3) is true.¹² Premise (1) is uncontroversial. An evidence proposition as understood here is just a proposition describing a potentially observable state of a physical system, such as ‘the pointer moved to position 5 on the dial’. A theory that failed to entail evidence propositions could not be confirmed by observation. Premise (2) is intuitively plausible: evidence propositions *seem* to be claims involving spatiotemporal configurations of observables. But by (3), physics doesn’t have local beables in its fundamental ontology, so how could it entail evidence propositions? There’s no problem if

⁹ Healey (2002); Maudlin (2007); Huggett & Wüthrich (2013).

¹⁰ See the appendix to Ney (2012) for details.

¹¹ Following Maudlin (2007); see also Ney (2015).

¹² Some solutions to the problem of empirical coherence that purport to show how empirical confirmation might work given non-spatiotemporal fundamental ontologies in my view tacitly presuppose that (3) is false, and so fail to deliver on their promises. I return to this in §2.3.

the fundamental ontology is spatiotemporal, for in that case it contains fundamental local beables, whose dynamics determine the behaviour of the non-fundamental things they compose. By contrast, non-spatiotemporal fundamental theories seem to preclude their own confirmation, and this is the sense in which they are threatened by empirical incoherence.

1.2. Realist and Antirealist Solutions

As I have set things up, the problem of empirical coherence is one of entailment: how could a non-spatiotemporal fundamental theory entail evidence propositions attributing spatiotemporal properties to local beables? In what follows, I will classify responses to this problem as *realist* or *antirealist*, depending on whether or not they are committed to a non-fundamental spacetime in addition to the ontology of fundamental physics. My usage of ‘realism’ and ‘antirealism’ here departs from traditional usage in the philosophy of science. Some positions classified as antirealist in my usage would be classified as realist according to certain traditional measures, for instance whether or not there are mind independent spatiotemporal truths. My distinction concerns only the existence or otherwise of a non-fundamental spacetime as truthmaker for such propositions, and should be taken to have no further ontological implications. For brevity, I will henceforth use ‘S-concepts’ to refer to spatiotemporal concepts, and ‘S-propositions’ for propositions featuring such concepts.¹³

In order for a non-spatiotemporal fundamental theory to have observational consequences, we will need some way of connecting the language of that theory with the S-concepts that we use to describe possible observations. Clearly, it’s impossible to know what form this connection will take without a final theory. My aim in this paper is to explore the space of possibilities by focusing on the spatiotemporal side of the relation. According to realism, S-propositions

¹³ I shall use underscore to refer to specific concepts.

describe non-fundamental states in the *derivative* ontology of the fundamental theory, with ordinary objects and their spatiotemporal properties *emergent* in relation to the fundamental.¹⁴ Perhaps the fundamental structure causes a non-fundamental spacetime structure to emerge in the appropriate classical limit, or perhaps it synchronically grounds that structure. For realists, the challenge is to give an account of the emergence relation itself, which will then bridge the fundamental and the spatiotemporal, and render (1)-(3) consistent. Antirealism eschews robust emergence in favour of a deflationary approach, according to which S-concepts don't attribute the kind of properties we ordinarily take them to. As we shall see, there are several distinct forms of antirealism, but they all have it in common that they deny premise (2). The challenge for antirealists is to give an account of S-concepts that is both conceptually plausible *and* which offers the promise of a solution to the problem of empirical coherence. My primary aim in this paper is to argue that these aims cannot jointly be met, thereby motivating realism.

2. Spacetime Antirealism

A preliminary defence of spacetime antirealism is to be found in Ney's work on wavefunction realism. Given wavefunction realism, if 3D space is real, then it must somehow be derived from 3N-dimensional configuration space. Ney suggests that we can do without derived space:

I am not arguing that on a straightforward, ontological reading of realist versions of quantum mechanics that there are not rocks....However, I claim that this does not imply that there is anything that is genuinely rock-shaped (if this implies occupying locations in a physical, three-dimensional space) or that moves in precise trajectories through such a space. Since we can allow that the wavefunction simulates the behaviour of

¹⁴ See Ney (2012) for full discussion of this position in the context of the wavefunction realism debate. In my current usage, emergence is consistent with physical realization. If general relativistic spacetime were grounded in a non-spatiotemporal fundamental ontology in a fully transparent way, then it would count as emergent in this usage. This usage of the term 'emergence' is common among philosophers of science, but less so among metaphysicians, who tend to think of the sort of emergence that gives rise to genuinely novel dependent entities as requiring in principle non-deducibility or brute supervenience. See Yates (2016) for more on this issue.

something in a three-dimensional space, we have a way of making sense of our observations in terms of the language of the theory.¹⁵

Ney's suggestion is that we can identify physical objects with states of the wavefunction on configuration space, provided we give up on the idea that they literally have spatiotemporal properties like location, shape and so forth. Thus, Ney is an antirealist in the current sense, seeking to explain empirical confirmation without having to derive a non-fundamental spacetime from fundamental physics. According to Ney, confirmation requires only:

[A] difference that can be grounded in the ontology of the quantum theory between what we refer to as 'the pointer pointing to up' and 'the pointer pointing to down.' But the wave function realist is free to trace the difference in what we refer to as 'the pointer pointing to up' and 'the pointer pointing to down' to differences in the state of the wave function in the two scenarios.¹⁶

On this view, confirmation doesn't require observation of distinct states of local beables, but simply observation of distinct states, which may then be understood as being states of the wavefunction on configuration space. What we ordinarily describe as someone perceiving successive states of a measuring instrument in 3D space is really the wavefunction as a whole evolving from one state to another. Both perceiver and instrument are identified with patterns in the wavefunction, and despite appearances to the contrary, neither is located in 3D space. In order for this kind of confirmation theory to work, we need a theory of S-concepts according to which S-propositions are in principle deducible from the fundamental theory, but are not attributions of spatiotemporal properties to local beables. There are several options available, all of which involve treating S-concepts as *non-transparent*.

¹⁵ Ney (2012), p. 550. Note that in her (2017), Ney defends the possibility of *deriving* ordinary spacetime from configuration space mereologically, so isn't committed to antirealist solutions.

¹⁶ Ney (2015), p. 3210.

2.1. Transparent vs. Non-transparent Concepts

I will say that a concept is *transparent* if deploying that concept involves thinking about its referent *as* the bearer of (at least some of) its essential properties. Transparent concepts, if any such there be, give us a kind of *a priori* access to the natures of their referents, in the sense that thinking about an entity under a transparent concept is thinking about it as it is essentially. Non-transparent concepts pick out their referents in terms of inessential properties.¹⁷ Our everyday concept water is *non-transparent*: to think ordinary water-thoughts is to think about a natural kind in terms of its surface features, such as being colourless, odourless and tasteless. The chemical essence of water is only revealed by scientific investigation, and is not accessible to us *a priori*. Many suppose that in contrast to natural kind concepts, *phenomenal* concepts are transparent. If so, then to think about pain (say) in terms of the way it feels is to represent a specific mental state as it is essentially. The essence of pain is revealed by experiencing what it's like, and this experience enables us to form a phenomenal concept pain, which then refers to the property of being in pain by means of its essential phenomenology.¹⁸

Now S-concepts are intuitively transparent: something of the nature of sphericity is revealed to us, we might suppose, in deploying the concept spherical. Chalmers offers several independent twin-Earth style arguments against the transparency of S-concepts.¹⁹ Chalmers suggests²⁰ that S-concepts might pick out “classical properties in a Newtonian world, relativistic properties in a relativistic world, and string-theoretic properties in a string-theory world.” For Chalmers, S-concepts are functional, and pick out elements of the fundamental

¹⁷ Obviously, on such an account, there can be no transparent concepts of things that don't have essential natures. Goff (2017, pp. 97-100) is particularly clear on transparent concepts.

¹⁸ I need not commit to the transparency of phenomenal concepts for present purposes, as my aim is merely to give the reader a sense of what's involved in affirming or denying conceptual transparency in general. See Lewis (1995) for compelling arguments that materialists ought to reject the transparency of phenomenal concepts.

¹⁹ Chalmers (2019); (2012), pp. 325-36. I lack the space to engage with Chalmers' twin-Earth arguments here, but the position suggested in §3.1 predicts twin-Earth cases for partially transparent S-concepts.

²⁰ Chalmers (2012), pp. 325-6.

Forthcoming in *Philosophy Beyond Spacetime*, eds. Wüthrich, Huggett & Le Bihan. OUP.

ontology as the occupants of contingent roles, and not as the properties they are in themselves. This sort of position is known, variously, as *realizer*, *conceptual*, or *filler* functionalism; I will use the term ‘realizer functionalism’. In the remainder of this section I will consider two versions of realizer functionalism about S-concepts, both of which are attributable to Chalmers: *phenomenal* and *theoretical*, depending on the kind of role specification involved.

2.2. Phenomenal Realizer Functionalism

Most would accept that if we are in the Matrix, then almost all of our beliefs are false—it’s not true that I’m sitting here typing this sentence if I’m attached to a computer that’s feeding me all the sensory inputs that cause me to have experiences as of sitting here typing it. Spacetime realizer functionalism, for Chalmers, forms part of a challenge to this orthodoxy. In *phenomenal* spacetime realizer functionalism, sphericity is the concept of whatever normally causes sphere-ish experiences. If phenomenal realizer functionalism is true, then if I am in the Matrix and have an experience as of a spherical ball, I can truly say that I see something spherical, provided whatever part of the Matrix program it is that normally causes that kind of phenomenology is the cause of my experience. The bearers of spatial properties are genuine physical objects, albeit objects whose fundamental proper parts are bits of code in the Matrix mainframe. The things we think of as fundamental—electrons, photons, and quarks, for instance—are themselves composed of bits, and compose into further digital objects. Some of these latter have the complex property that normally causes sphere-ish experiences in Matrix dwellers, and are therefore spherical objects.

An immediate worry is that things that aren’t spatiotemporally located can’t cause our experiences, rendering phenomenal realizer functionalism of no use in addressing the problem of empirical coherence. I won’t address this issue here, and assume that on a broad enough

notion of causality, it makes sense to suppose that our experiences as of things located in spacetime can be caused by things that are not themselves spatiotemporally located. Can S-concepts be functionalized in the way this version of phenomenal realizer functionalism requires? For my part, I don't think they can. Readers familiar with the history of the philosophy of mind will be aware of the reasons for the decline of logical behaviourism. According to that theory, mental states have *a priori* dispositional analyses—the belief that it's raining is the disposition to take your umbrella if you go outside, to say “yes” when asked if it's raining, and so on. The problem with such analyses stems from the holism of the mental: even if you believe it's raining, you're only disposed to take your umbrella with you if you go outside *if you also want to stay dry*. But wanting to stay dry, *qua* mental state, will then be analysed in terms of the disposition to take your umbrella...if you also think it's raining. Which dispositions you have depends on the totality of your mental states, so no individual mental state suffices *a priori* for any specific behavioural disposition. A similar problem arises on the simple version of phenomenal realizer functionalism sketched above.

What phenomenology is associated with a given S-concept? What does being in contact look like, or in relative motion? Focus on the latter. The phenomenology associated with being in motion relative to a perceiver depends on many things, including: (i) what shape the object is, (ii) where the light sources are located in relation to object and perceiver, (iii) the direction of motion relative to the perceiver, (iv) the current distance between perceiver and object, and (v) the spatial relations between the object and other things in the perceiver's environment. But (i)-(v) all involve further S-concepts. Another example: what do two cylindrical rods look like, placed a metre apart endwise? It depends where the perceiver is located in relation to the rods, the relative size and motion of the rods and perceiver, and so forth. In short, the way things that fall under a given S-concept appear to a given perceiver depends on the other S-concepts that

object and perceiver fall under. A given S-concept C can't refer to whatever normally causes C-ish phenomenology, because there's no way to specify C-ish phenomenology independently of the other S-concepts that the perceiver and the object perceived fall under. Just as behavioural dispositions can't be paired up 1-1 with mental concepts, so the phenomenology associated with falling under S-concepts is associated not with individual concepts, but with complex conjunctions thereof.

This leaves room for an holistic version of phenomenal realizer functionalism, according to which S-concepts are jointly defined in terms of the phenomenology associated with falling under many S-concepts at once. On this theory, we refer to complex fundamental states by means of complex S-propositions, with the level of complexity dependent on how much we have to build into the proposition to describe a situation with a unique spacetime phenomenology for a perceiver embedded within it in a particular way. Chalmers offers remarks that are consistent with holistic phenomenal realizer functionalism:

[S]patial concepts pick out that manifold of properties that serves as the normal causal basis of a corresponding manifold of properties in our spatial experience. Temporal concepts pick out that manifold of properties that serves as the normal causal basis of a corresponding manifold of properties in our temporal experience.²¹

Gather the manifold of properties (intrinsic and relational) that jointly cause complex spacetime phenomenology together into a set M. Schematically, to think of something as spherical is then to think of it as having a certain property $P \in M$ such that P coinstantiated with the properties in $X \subseteq M$ normally cause sphere_x-ish phenomenology; and so on for the other subsets of M. This is quite some level of conceptual sophistication to attribute to the folk. Deploying a concept like sphericity involves thinking of an object in terms of how it would appear if it fell under that concept together with certain other concepts of the same kind—enough to correspond to a

²¹ Chalmers (2012), p. 335.

determinate phenomenology and secure unique reference to the relevant cause property. This isn't a knock-down objection, but it does make holistic phenomenal functionalism look somewhat implausible. These are *our* S-concepts, and while on realizer functionalism we don't have *a priori* access to the natures of their referents, we should surely have (at least some) *a priori* access to the manner in which they refer. So, when you think of something as spherical, are you thinking about it in terms of the various ways it would look if you and its bearer also fell under various combinations of further S-concepts? That doesn't seem to me to be what I am doing. Phenomenal realizer functionalism in any form also ties cognitive contents to specific sensory modalities, and so entails (for instance) that blind or partially sighted people can't think the same spatiotemporal thoughts as normally sighted people, because their concepts involve different phenomenological reference-fixers.²² A more plausible theory, to my mind, is that when we attribute sphericity to something, we are thinking about what the results would be if we measured the distance from its outer boundary to its centre—but not, of course, about the phenomenology associated with those measurement outcomes.

A further problem for phenomenal functionalism, considered as a solution to the problem of empirical coherence, is that it requires bridge principles connecting the non-spatiotemporal fundamental theory with the relevant phenomenology, but it's very difficult to see how such principles could be supplied without already having solved the problem at hand. In order to avoid a possible misunderstanding, first consider a *prima facie* similar problem that arises in spatiotemporal theories, which of course don't suffer from the problem of empirical coherence. In cases where we are able to deduce evidence propositions from a theory, we still need ways of determining, on the basis of our conscious experiences, whether those propositions are true.

²² In terms of 2D semantics, the spatiotemporal thoughts of blind folk at our world will have the same secondary intensions as, but different primary intensions from, the thoughts of sighted folk. I assume here that for cognitive purposes, the primary intensions are what matter.

This conceptual competence need not take the form of phenomenal analyses of the relevant concepts, but grant for the sake of argument that it does. The connection between theory and phenomenology can then be established in two steps: (i) deduction of evidence propositions from the theory, and (ii) functional analysis of the concepts features in those propositions in phenomenal terms. The problem for *non*-spatiotemporal theories, however, runs deeper.

If phenomenal analysis is to render S-propositions deducible from a non-spatiotemporal theory, we need a way of correlating arbitrary states of that theory's ontology with the phenomenology that those states typically bring about. It's no use analysing sphericity as the concept of whatever normally causes sphere-ish experiences if we don't know which state of the fundamental ontology that is. The problem now is that this correlation must be established *prior* to the derivation of any S-propositions from the theory: our apparent inability to deduce such propositions in quantum theories of gravity is the problem that phenomenal functionalism is supposed to solve. However, everything we think we know about perception relies on our being agents with sensory receptors located in a spatiotemporal world. Our perceptual systems, we think, work by means of sensory receptors that convert physical inputs—for instance photons or air pressure waves—into electrical signals that our brains use to construct representations of their causes. Following Chalmers,²³ suppose that a neural correlate of consciousness is “[a] neural representational system N such that representation of a content in N directly correlates with representation of that content in consciousness.” Suppose now that we have made sufficient progress in neuroscience to identify the neural correlates of arbitrary phenomenal states. To predict how a *non*-spatiotemporal fundamental state will appear to creatures like us,

²³ Chalmers (2000), p. 20. Direct correlation is understood in terms of minimal sufficiency: representation of the given content in N is sufficient for its representation in consciousness, and no proper part of N is sufficient. I lack the space to go into detail, and appeal to this notion here only for illustrative purposes.

we would *already* need a way of deducing S-propositions from the fundamental theory, because the neural correlates of phenomenal states are spatiotemporal.

Phenomenal realizer functionalism does have the resources—setting aside the question of its conceptual plausibility—to explain how S-concepts could refer to the elements of a non-spatiotemporal fundamental ontology. However, if we have no way of finding out *which* fundamental properties they refer to, we are no closer to explaining how S-propositions could be deducible from the fundamental theory. Phenomenal realizer functionalists need to explain how, in principle, we could identify which fundamental properties bring about the various types of spacetime phenomenology that define our S-concepts, but without relying on spatiotemporal theories of perception. In response to this, it might be suggested that holistic phenomenal realizer functionalism could achieve this correlation by mapping the structure of our spacetime phenomenology directly onto the structure of the fundamental theory. In effect, this would treat S-concepts as defined by means of a phenomenological *theory*, and explain the deducibility of S-propositions from the fundamental theory in terms of inter-theoretic reduction. I shall now discuss what I take to be a far stronger version of theoretical realizer functionalism, based not on phenomenology but on GTR, and argue that it fails as a solution to the problem of empirical coherence for reasons that would also apply to the phenomenal theory just sketched.

2.3. Theoretical Realizer Functionalism

Chalmers hints at a theoretical version of realizer functionalism about S-concepts, but doesn't really flesh it out.²⁴ In this section I suggest a way to do so. Einstein's field equation, in its standard formulation, is a set of tensor equations that relate the matter distribution to the local

²⁴ See for instance Chalmers (2012), pp. 332-4, pp. 418-9. A version of theoretical spacetime functionalism is also developed in Lam and Wüthrich (2018); I return to their views presently.

spacetime curvature. The *metric tensor* describes all the geometric properties of spacetime, and can be used to define properties such as volume, distance and intrinsic curvature. The field equation of GTR relates the metric to the distribution of mass-energy by means of the *stress-energy tensor*, and expresses a mutual lawlike dependence between the matter distribution throughout spacetime and the metric structure of spacetime itself. Whether there is a causal relationship between the matter distribution and the metric is controversial, but if there is, it's not that the matter distribution *brings it about* that spacetime has a certain metric structure, because the specification of the matter distribution in the stress-energy tensor *presupposes* a metric. A more natural interpretation is that the field equation tells us that:

[G]iven both a certain intrinsic geometry for spacetime and a specification of the distribution of mass-energy throughout this spacetime, the joint description is the description of a general-relativistically possible world only if the two descriptions jointly obey the field equation".²⁵

Eleanor Knox appeals to GTR to defend functionalism about the concept spacetime.²⁶ On her view, spacetime is the concept of "whatever occupies the role of defining a structure of inertial frames". Inertial frames are those with respect to which bodies not acted on by a force travel at uniform velocity, and within which the laws of physics take a uniquely simple form. Knox argues that the way the metric field couples to matter fields guarantees that the local symmetries of the metric are the same as the local symmetries of the dynamical laws. This in turn explains why the metric governs the behaviour of bodies that obey those laws. Given that the metric field occupies the spacetime role, Knox's functionalism entails that it should be identified with spacetime. Now if the fundamental ontology is non-spatiotemporal, then the metric field isn't fundamental. Knox's brand of realizer functionalism therefore identifies spacetime with a non-fundamental, derived structure, and so doesn't provide an antirealist solution to the problem of

²⁵ Sklar (1976), p. 75

²⁶ Knox (2017).

empirical coherence. To do that, we need to embrace theoretical realizer functionalism about the S-concepts that appear in our evidence propositions, and not just the concept spacetime.

According to the Ramsey-Lewis semantics, theoretical terms are defined by their places in a theoretical structure.²⁷ We can write a theory T schematically as follows: $T[\langle t_1, \dots, t_n \rangle]$, where the ordered n-tuple $\langle t_1, \dots, t_n \rangle$ denotes the theoretical (or t-) terms to be defined by means of T. The remainder of T's vocabulary (the o-language) consists of terms whose meanings are given independently of T. We get the Ramsey-sentence of T by replacing $\langle t_1, \dots, t_n \rangle$ with bound variables: $\exists! \langle x_1, \dots, x_n \rangle T[\langle x_1, \dots, x_n \rangle]$. The meaning postulates of T are the Carnap-sentence: $\exists! \langle x_1, \dots, x_n \rangle T[\langle x_1, \dots, x_n \rangle] \rightarrow T[\langle t_1, \dots, t_n \rangle]$, and a sentence saying that $\langle t_1, \dots, t_n \rangle$ are empty if T is not uniquely realized. Lewis appeals to these sentences to offer an explicit definition of the t-terms: $\langle t_1, \dots, t_n \rangle =_{df} \iota \langle x_1, \dots, x_n \rangle T[\langle x_1, \dots, x_n \rangle]$. Less formally, $\langle t_1, \dots, t_n \rangle$ refer to the unique n-tuple satisfying T, or (less formally still) to the n-tuple that occupies the T-roles. S-concepts appear as theoretical terms in GTR, so why not Ramsify them relative to that theory and say that they refer to the fundamental occupants of the GTR roles?

This proposal doesn't entail that S-concepts are *entirely* theoretical, defined in ways that are beyond the ken of anyone who doesn't know GTR. S-concepts already appear in our folk conception of the world, and in some respects at least, GTR can be seen as an extension of that conception. Concepts like duration, distance, contiguity, shape, area, and volume all figure in our everyday conception of space and time, by way of the range of folk physical platitudes we are inclined to accept—for instance, the transitivity of 'later than', or the idea that only contiguous objects can interact. If there's enough in common between our folk spacetime theory and GTR, then the S-concepts in question can plausibly be seen as extended by GTR,

²⁷ See Lewis (1972) for details.

not re-defined wholesale. The issue of how wrong a theory can be before reference failure occurs is of course controversial, but it's at least plausible that our folk conception of spacetime gets *enough* right for GTR to count as a conservative extension thereof.²⁸ On the present account, then, when we deploy ordinary S-concepts to note the position of the meter on the dial, we refer to fundamental non-spatiotemporal properties, by means of the theoretical roles the corresponding scientific concepts assign to something via GTR.

Theoretical realizer functionalism explains the in principle deducibility of S-propositions from a non-spatiotemporal fundamental theory, on the assumption that states expressible in the language of that theory occupy GTR roles. We don't need psychophysical bridge principles in this case—an S-concept such as locality refers to the fundamental occupant of the locality role as defined in GTR, so deducibility is a matter of mapping the fundamental role occupants onto the GTR roles.²⁹ The problem for theoretical realizer functionalism, as a solution to the problem of empirical coherence, is that it ceases to work just when (and if so because) we have good reason to regard the fundamental ontology as non-spatiotemporal in the first place. If the terms of GTR are to refer to non-spatiotemporal properties in a quantum gravity ontology, then we will need to recover enough of the nomic structure of GTR from the fundamental theory that GTR role specifications pick out fundamental role occupants. But the more closely the structure of GTR can be recovered from a quantum theory of gravity, the less reason we will have for regarding that theory as non-spatiotemporal. The same issues arises, *mutatis mutandis*, for the phenomenal version of theoretical realizer functionalism mooted at the end of §2.2. An example from the philosophy of mind will clarify this problem.

²⁸ I return to this point in §3.1.

²⁹ Note that this will require inter-theoretic bridge principles in order to connect GTR roles with the nomic roles available in the fundamental theory. My point is that no problematic *psychophysical* bridge principles are needed in order to explain how S-propositions are deducible from the fundamental theory.

Suppose we Ramsify folk psychology, so that (for instance) pain refers to the brain state that's caused by tissue damage, and causes both the various forms of physical and verbal pain behaviour, and other psychological states such as anxiety. Similarly, anxiety will refer to the brain state that occupies the anxiety role, part of which involves being caused by the brain state that occupies the pain role. In short, the entire nexus of folk psychological state terms refers holistically to a network of brain states that stand in the causal relations specified by the theory, both in relation to each other and to physical and verbal behaviour. In Lewisian philosophy of mind, given that pain = the occupant of causal role R, if the occupant of the pain role is brain state B, then pain is contingently *identical* to brain state B, because B contingently occupies pain's defining role.³⁰ That is to *reduce* pain to a state that is contingently mental, not to show how pain could refer to a *non-mental* state. Now suppose that no brain state comes close to occupying the folk psychological pain role, or the roles of beliefs and desires. It follows that folk psychological terms *won't* pick out brain states as the occupants of folk psychological roles, because those roles aren't occupied at all, and Ramsified folk psychology is a set of false existential claims.³¹ In that case, the brain may well be fundamentally non-mental, but for that reason, we won't be able to refer to brain states using Ramsified folk psychological terms.

Theoretical realizer functionalism faces a dilemma. If Ramsified GTR terms succeed in referring to the fundamental properties of a theory like LQG or string theory, then each GTR role is (approximately) occupied by such a property. But in that case, the fundamental properties in question can be *identified* with GTR spatiotemporal properties, in which case the relevant fundamental ontology can't be non-spatiotemporal. Conversely, if the ontology is non-spatiotemporal, then Ramsified GTR terms won't pick out its elements. It's instructive to

³⁰ Lewis (1994).

³¹ Those familiar with the literature on eliminative materialism will recognise the inference from the falsity of folk psychology to reference-failure via the Ramsey-Lewis semantics; see for instance Churchland (1981).

consider at this point a proposal due to Lam and Wüthrich on how a version of theoretical realizer functionalism helps us to recover spacetime in LQG.³² They suggest that *weave states* of LQG spin networks do have properties that approximately occupy the roles of locality, area and volume in GTR. Ramsified GTR terms can thus pick out the properties of such states as the (approximate) occupants of GTR roles, with many different weave states capable of occupying the same roles. However, Lam and Wüthrich are clear that this class of models is unified in having at most only slightly disordered locality. In effect, the disordered locality problem that motivates the claim that LQG is non-spatiotemporal must already be solved in order for Ramsified GTR terms to pick out fundamental LQG properties. What Lam and Wüthrich have shown, then, is that there is a class of models of LQG in which a theoretical realizer functionalist strategy could be made to work. But for precisely that reason, I submit, there's no obvious reason to regard the models in question as non-spatiotemporal.³³

What if we abandon the idea that S-concepts are semantically functional and try to treat them as directly referring natural kind terms? Suppose for illustration that some version of the causal-historical theory of reference is true.³⁴ On that theory, we first baptise a natural kind such as water by ostension, and thereby fix the reference of water for all subsequent users of the concept, by means of causal relations to the initial baptism. If S-concepts refer directly, then we no longer need to recover the nomic structure of GTR from the fundamental theory; deducibility of S-propositions will be secured when we can identify the fundamental kinds to which S-concepts directly refer. But how are we supposed to do *that*, if the fundamental ontology is non-spatiotemporal? Conversely, if concepts like contiguity, being 1m apart, sphericity, and so forth, refer to kinds expressible in the language of a fundamental theory of

³² Lam & Wüthrich (2018).

³³ I suspect that a similar dilemma arises for the claim that Ramsified GTR terms could refer to string theoretic properties in a space of indeterminate geometry, but I omit discussion for brevity.

³⁴ Kripke (1972).

quantum gravity, then why should we regard that theory as non-spatiotemporal? The disordered locality problem in LQG, for instance, seems to preclude identifying locality properties of spin networks with spatiotemporal locality, so if that problem is serious, we should expect there *not* to be fundamental LQG kinds corresponding to GTR spatiotemporal concepts. As before, it may be plausible that there are models of the relevant theories in which we can pick out fundamental kinds by means of ostension to spatiotemporal states of affairs, but again, it won't be clear why we should think of those models as non-spatiotemporal.

I conclude that antirealist approaches to the problem of empirical coherence are untenable. Phenomenal realizer functionalism is conceptually implausible, and whereas theoretical realizer functionalism has greater plausibility, as a solution to the problem at hand it depends on a tacit rejection of the claim that the fundamental ontology is non-spatiotemporal.³⁵

3. Spacetime Realism

Despite my misgivings about antirealism as a solution to the problem of empirical coherence, I think it's plausible that S-concepts are semantically functional and defined in terms of GTR. This idea is consistent with realism, because Ramsifying a theory T doesn't settle the issue of what kind of properties its t-terms refer to.³⁶ In particular, Ramsifying T doesn't commit us to treating its t-terms as referring to the fundamental *occupants* of roles specified by T. As Lewis himself notes,³⁷ we can also treat them as referring to a domain of second-order functionally individuated properties that are *realized* by the fundamental role occupants. On this theory, t-

³⁵ This tacit move could be made explicit, and theoretical realizer functionalism recast as a position according to which, despite appearances to the contrary, the fundamental ontology is spatiotemporal after all. The resulting position—suitably renamed—would then *dissolve* the problem of empirical coherence by rejecting (3) of §1.1.

³⁶ Whittle (2008) argues at length that the Ramsey-Lewis method, while useful for giving functional analyses of concepts as part of a broader functionalist theory of properties, does not in itself constitute such a theory.

³⁷ Lewis (1994), pp. 307-8. Realizer functionalism is a matter of theory choice, for Lewis, based on the uncommon assumption that second-order functional properties are causally excluded by their realizers. That assumption, however, can be challenged in various ways; see Yates (2012) for full discussion.

terms are still semantically functional, because they pick out their referents by means of their places in a structure, but they will also now be *transparent*, at least to the extent that they get the structure right. In short, the claim that the t-terms of a theory T are semantically functional leaves open whether their referents are role-occupants or role-properties,³⁸ and in the latter case, on the assumption that T is true, those concepts will be transparent, since they will pick out their referents by means of their defining roles. In the remainder of this section, I will first outline a version of role spacetime functionalism that's consistent with the partial transparency of S-concepts, and then consider whether a role functional spacetime is the sort of structure that could be synchronically grounded in the non-spatiotemporal.

3.1. Role Spacetime Functionalism

There are many different versions of role spacetime functionalism, depending on which theory supplies the role specifications. For instance, we might say that spherical refers not to the unique property that gives rise to sphere-ish phenomenology, but to the second-order property of having *some* property that occupies that role. On this view, which we might call simple phenomenal role spacetime functionalism, spatiotemporal properties themselves are defined in terms of giving rise to spacetime phenomenology, in contrast to the corresponding realizer functionalist theory, in which the phenomenology acts merely as a reference fixer to an underlying mind-independent reality.³⁹ An holistic version of phenomenal role functionalism about spacetime could also be given, in which spatiotemporal properties are defined in terms of the phenomenology they give rise to in various combinations with each other.

³⁸ It should be noted that these options are not mutually exclusive: on a powers theory, for instance, semantically functional concepts such as charge pick out fundamental properties as the occupants of the causal roles that individuate them. For more on the relational individuation of fundamental properties, see Yates (2018).

³⁹ Schwitzgebel (2019) argues for a Kantian phenomenological conception of Matrix spacetime based on simple phenomenal role functionalism about spacetime.

Functionalist theories are often divided into analytical (or common-sense) and scientific, depending on whether or not the theory that supplies the role specifications is *a priori*. This distinction is orthogonal to that between realizer and role functionalism. In Lewis' realizer functionalism, it's supposed to be *a priori* that pain refers to the unique occupant of a certain causal role because folk psychology is a *common-sense* theory of mind.⁴⁰ The theory in question, for Lewis, supplies reference-fixing descriptions that pick out brain states, but we can easily define a second-order version of Lewisian functionalism. According to this theory, the mental concept pain refers *a priori* to the property of being in some state that occupies the folk pain role. It follows that our ordinary mental concepts are fully transparent to their referents, because when we deploy those concepts, we thereby think of mental properties in terms of their essential natures, without remainder. Now consider *psychofunctionalism*,⁴¹ according to which mental properties are defined by their places in an *a posteriori* scientific psychological structure. In psychofunctionalism, it is often said that the role of folk concepts is not to capture the true natures of mental properties, but to fix the reference of mental concepts, which refer to properties whose true natures are given by deep scientific psychology. But now suppose that scientific psychology is a (relatively) *conservative extension* of folk psychology—suppose, that is, that the folk know something about the essential natures of the mental properties we attribute. On this view, not all of folk psychology will be abandoned as cognitive science progresses; some folk claims will be preserved (in some form) as components of mature scientific psychology. It follows that folk psychological concepts are *partially* transparent: they pick out mental states by specifying parts of their individuating roles. I shall now suggest that the same could be true of S-concepts.

⁴⁰ The role specifications are *a priori* in the limited sense that no scientific investigation is required in order to know them; possession of folk concepts is sufficient. Nothing turns on epistemic priority for my purposes.

⁴¹ See Block (1978).

Forthcoming in *Philosophy Beyond Spacetime*, eds. Wüthrich, Huggett & Le Bihan. OUP.

In the Minkowski spacetime of STR, properties like shape, size and duration vary depending on the relative motion of the observer and the objects or events in question; and in the pseudo-Riemannian spacetime of GTR, the intrinsic curvature of spacetime depends on the matter distribution. However, such radical departures from common-sense don't rule out that a theory is conservative, in the sense outlined above. As is familiar, special relativistic effects such as Lorentz contraction and time dilation are only significant between inertial frames at very high relative velocities. For observers that are not too far from being in the same inertial frame, shapes, sizes, and durations are (approximately) the same, and the same events are (approximately) simultaneous. For observers at relative velocities far from the speed of light, our common-sense conception of space and time, as reflected in our pre-theoretic S-conceptual structure, is approximately *true*. GTR is more revolutionary than STR, as it treats spacetime as a dynamic structure rather than a background container for the dynamics. However, the point still stands: locally and for weak gravitational fields, GTR spacetime approximates Minkowski spacetime, which in turn departs from common-sense only at high relative velocities.

Relativity doesn't so much refute common-sense as reveal its *limits*. Because common-sense posits like absolute spatial and temporal intervals, and flat Euclidean geometry, are recoverable in the appropriate limits of STR and GTR, it's plausible that at least some of our common-sense theory of space and time is preserved therein. Now suppose that ordinary S-concepts refer to role functional properties, individuated by their places in a structure described by GTR.⁴² Partial transparency follows. Our ordinary ways of thinking about space and time are a mixture of semi-formal and causal principles, for instance: the transitivity of 'longer than' and 'later than'; the idea of an absolute present moment; the claim that A and B are equidistant

⁴² The Minkowski metric is an approximate solution to the field equations of GTR in the limit of zero gravitational field, i.e. empty space, so STR is typically seen as reducible to GTR. For that reason I refer only to GTR as the theory that specifies the defining roles for a role-functionalist conception of spacetime. See Lam & Wüthrich (2018) for a functionalist analysis of (*inter alia*) spatiotemporal locality based on its role in GTR.

from C iff it takes the same time to travel between A and C as between B and C at constant velocity; spatiotemporal contiguity as a necessary condition on interaction; the principle that causes must precede their effects; and so on. If some of this structure is preserved in GTR, then our ordinary ways of thinking about spacetime represent their referents as relata in (a proper substructure of) the general relativistic structure that individuates them.⁴³

As long as the folk aren't completely wrong about spacetime, there's room for a functionalist theory that combines the partial transparency of S-concepts with the idea that spatiotemporal properties are scientific kinds. This entails the possibility of twin-Earth cases involving S-concepts, for we can always imagine worlds with properties that instantiate the folk structure, but at which the scientific extension is different, or even absent. Imagine a twin-Earth, for instance, at which spacetime is a classical background container for the dynamics with zero curvature everywhere. It's intuitively plausible that twin-Earthlings would have the same everyday S-conceptual framework as Earthlings, and also that the twin-Earth concept T-separation, for instance, refers to a different property than our separation. Our concept picks out a property of a dynamic structure whose geometry depends on the matter distribution, whereas theirs does not. However, given the possibility of a version of role functionalism in which at least some of the folk roles are incorporated into scientific theory, this doesn't suffice to show that our S-concepts are *fully* non-transparent. The fact that we can construct twin cases for a family of concepts suggests that there's more to their referents than meets the eye, but it doesn't follow that what meets the eye is merely superficial.⁴⁴

⁴³ Note that this does not amount to what Chalmers (2012, p. 325) calls spatiotemporal *primitivism*, the claim that we have direct unmediated grasp of spatiotemporal properties. Our grasp of such properties, on the position outlined here, is mediated by our grasp of the structure that individuates them.

⁴⁴ I don't claim here that role functionalism can handle the twin-cases offered in Chalmers (2019). As a matter of fact I think it can, but I will defend that claim in future work.

3.2. Grounded or Caused?

In this section I will first consider the conditions under which it is possible for spacetime to be grounded in a non-spatiotemporal physics, then briefly consider the implications of causal spacetime emergence. Role functionalist theories need a theory of realization, because they are theories intended to describe a domain of dependent properties such as mental or other special science properties. In traditional *flat* theories of realization, a functional property F is defined as the property of having some property that occupies role R (usually causal), and a property P realizes F by filling R . This theory presupposes that P and F are instantiated in the same object, and for that reason won't help us with the problem of empirical coherence. The properties whose realization we want to make sense of are spatiotemporal properties of local beables, but the fundamental objects in the putative realization base are, by hypothesis, not bearers of spatiotemporal properties. To account for the realization of spatiotemporal properties, we need a *dimensioned* theory of realization:

A property-instance $F_\phi(x)$ is realized by properties and relations $P_1, \dots, P_n, R_1, \dots, R_m$ iff (i) x or its proper parts possess P_1, \dots, P_n and R_1, \dots, R_m in some combination; and (ii) x meets the specification ϕ definitive of F_ϕ in virtue of (i), but not *vice-versa*.⁴⁵

Now suppose local beables are composed of elements of the fundamental ontology—field quanta, modes of vibration of cosmic strings, or whatever they might turn out to be—that are not located in ordinary spacetime. The fundamental particulars don't have ordinary spatial or temporal properties and relations, but perhaps local beables can have their spatiotemporal properties in virtue of the *non*-spatiotemporal properties and relations of their proper parts. By way of illustration, consider again the Matrix. Suppose that what Matrix dwellers perceive as

⁴⁵ Gillett (2003). The 'not *vice-versa*' clause secures the asymmetry of realization. I have adapted Gillett's position so that F is no longer understood as causally individuated, and realization need not consist in bestowing powers. Realization of causally individuated properties by causal power bestowal is a special case of the definition above. For role functional spatiotemporal properties as understood here, ϕ will be given by GTR, and may well specify broadly causal roles, but I find it implausible that realizing GTR spacetime could consist in causal power bestowal.

physical objects are composed of states of computer programs—the digital objects that combine to cause Matrix dwellers’ experiences as of seeing an apple are parts of a *virtual* apple, whose fundamental nature is computational.⁴⁶ On a role functionalist conception of spacetime, there’s no obvious reason why Matrix virtual objects can’t have computational properties and relations in virtue of which they satisfy the specifications that define spatiotemporal properties. If we can flesh out the details of the GTR roles in a way that explains how they could be jointly occupied by computational properties of digital objects, then GTR will be literally true in the Matrix, and not merely simulated, provided the roles in question are occupied there. There are of course many complications—not least the fact that the GTR roles will include the way in which the mass-energy distribution curves spacetime, so realizing GTR spacetime in a Matrix environment also requires that mass-energy is digitally realizable. Global role functionalism about fundamental physical properties, as we arguably find in a pure powers theory, would seem to have this consequence, but I cannot develop this proposal here for reasons of space.⁴⁷

The position just sketched suggests that local beables need not instantiate the same locality properties as their fundamental proper parts, and that’s just what is needed in order to make sense of the emergence of spacetime. Locality in GTR spacetime, conceived as a role functional property, is instantiated by non-fundamental local beables, with the fundamental proper parts of such things having different locality properties in the space that they inhabit.⁴⁸ Suppose global role functionalism about physical properties, and consider a Matrix world that is a qualitative physical duplicate of actuality. Within this Matrix’s spacetime, the virtual Earth really is 149.6 million km from the virtual Sun, but they are composed of digital objects that

⁴⁶ Chalmers (2005); see Chalmers (2017) for more on the ontology of virtual reality.

⁴⁷ See Yates (2018).

⁴⁸ Le Bihan (2018) suggests that locality is relative to a structure, and distinguishes *fundamental* from *emergent* locality on that basis. The idea that a complex object could be located relative to GTR spacetime, while its proper parts are not, is a consequence of allowing composition and realization to connect distinct locality structures.

are stored within the same supercomputer and (let's say) 1nm apart relative to our spacetime. If we take seriously the claim that the spatiotemporal properties of local beables are realized by the properties and relations of their fundamental proper parts, this kind of consequence should not be any more surprising than the idea that the mental is grounded in the non-mental.

The problem with this proposal is that classical mereology is, in Paul's terminology, *spatiotemporalist*:⁴⁹ we build complex local beables from more fundamental local beables, with the spatiotemporal properties of the whole dependent on the spatiotemporal properties and relations of the parts. In the Matrix example given above, it's hard to see how Matrix objects could be composed of computer program states unless those states themselves instantiate *spatiotemporal* properties and relations, in virtue of which they compose into complex objects. If we want build things that are in spacetime from things that are not, we need an alternative mereology. Paul herself defends a non-spatiotemporalist version of mereological bundle theory, in which local beables are primitively restricted fusions of properties:

Objects may have their locations in virtue of being fused with whatever location properties and relations there are that define the actual space of the world, and many objects will have a physical structure in virtue of having location properties and relations as parts of their fusions, or in virtue of being part of a larger fusion which has location properties and relations as parts. The character of the space might not be what we take the character of ordinary spacetime to be, but the structure of the space is generated by fusing qualitative properties with relevant properties and relations that define the space as determined by modern physics.⁵⁰

Paul intends mereological bundle theory to be maximally flexible, so that objects, events, facts, states of affairs, etc., can be generated by bundling together fundamental qualitative and location properties, without the requirement that the location properties are spatiotemporal.⁵¹ It's not entirely clear to me how to flesh out the details, but the general idea is that by

⁴⁹ Paul (2012).

⁵⁰ Op. Cit. p. 242.

⁵¹ Op. Cit. pp. 244-5.

abandoning spatiotemporalist mereology, we make room for an ontology in which fundamental things that are not spatiotemporally located can be proper parts of non-fundamental things that are. If we embrace a version of Paul's mereology with fundamental particulars as well as properties, we can maintain role functionalism about spatiotemporal properties without giving up on the dimensioned account of realization sketched above. The spatiotemporal properties of local beables will then be properties they have in virtue of the non-spatiotemporal properties and relations of their proper parts. A version of this claim may also be defensible on Paul's own view, although of course in that case the proper parts *just are* properties and relations.⁵²

What if GTR spacetime is not grounded in a non-spatiotemporal fundamental structure at all, but emerges from such a structure *causally*? Suppose for the sake of argument that LQG is the correct approach to spacetime. If GTR spacetime emerges causally from a spin foam, then there's no need to explain how the latter, complete with its disordered localities, could realize the former. Rather, the problem is to explain why the relevant non-localities disappear at low energies, and that's a problem for physicists, who need to supply a *causal* bridge principle in place of the metaphysical principle provided by realization.⁵³ As we saw in §2.2, if a LQG spin foam evolves into a low energy state in which the disordered localities disappear, then its states may be suitable to occupy the roles of GTR spatiotemporal properties. In this condition, I suggested, what we have is a spatiotemporal model of a theory many of whose models will be non-spatiotemporal due to disordered locality. Role functionalism about spacetime is also consistent with this kind of possibility. If LQG states can occupy the GTR roles, then they can realize a role functionalist spacetime structure defined by those roles.

⁵² See Le Bihan (2018) for further discussion of the application of Paul's mereology to spacetime emergence.

⁵³ Wüthrich (2017); Lam & Wüthrich (2018).

Now if it turns out that the functional roles associated with S-concepts are indeed occupied by properties specifiable in the language of the fundamental theory, then we may well decide that a realizer functionalist approach is best. The issue there will depend, as it does (*mutatis mutandis*) in the philosophy of mind, on whether we think our concepts give us at least some epistemic access to a higher-order, multiply realizable spacetime structure, or simply enable us to refer to whatever occupies the relevant functional roles. However, if our goal is to develop a metaphysical framework that doesn't presuppose any particular theoretical developments in physics, then role functionalism is clearly preferable. It's not so clear, for example, that a string-theoretic target space of indeterminate geometry could evolve causally into GTR spacetime, in the way we might imagine the non-localities of LQG disappearing at low energies. Role functionalism, if combined with suitable theories of realization and composition, leaves open that the fundamental ontology is synchronically non-spatiotemporal, and that gives us ample reason to work out the metaphysical details.

4. Conclusion

Working out how a non-spatiotemporal fundamental physical theory raises the probability of propositions featuring spatiotemporal concepts is a problem for physicists, but metaphysicians can fruitfully think about the conditions under which this is possible. Antirealism, according to which spatiotemporal concepts pick out elements of the fundamental ontology, is either untenable on conceptual grounds, or unsuitable for addressing the problem at hand. The task then is to work out a version of realism—the theory that spatiotemporal concepts refer to derivative, non-fundamental properties and relations—that is both conceptually plausible and consistent with the broadest possible range of relations between the fundamental ontology and grounded spacetime. GTR-based role spacetime functionalism allows spatiotemporal concepts to be scientific without denying their intuitive transparency, and allows GTR spacetime to be

Forthcoming in *Philosophy Beyond Spacetime*, eds. Wüthrich, Huggett & Le Bihan. OUP.

synchronically grounded in a non-spatiotemporal fundamental structure. This is clearly not a fully worked out theory, and many questions remain, for instance: (i) What exactly are the roles assigned to spatiotemporal properties in GTR? (ii) Are those roles causal, nomic, structural...? (iii) What theory of realization is most appropriate for second-order properties defined by the roles in question? (iv) What metaphysics are required to allow for the composition of local beables by non-spatiotemporal proper parts? I leave these questions for future work.⁵⁴

⁵⁴ Versions of this paper were presented at a seminar at the University of Geneva, at a symposium on spacetime emergence (organized by Baptiste Le Bihan) at ECAP9 in Munich, at a conference on metaphysics at the Federal University of Rio de Janeiro (organized by Guido Imaguire), and at a seminar at the University of Lisbon. Particular thanks to: Baptiste Le Bihan, David Chalmers, Karen Crowther, Natalia Deng, Nick Huggett, Bryan Pickel, Benjamin Schnieder, Chris Wüthrich, Elia Zardini, and two anonymous referees. Funded by the Fundação para a Ciência e a Tecnologia, award nos. IF/01736/2014 and PTDC/FER-HFC/30665/2017.

REFERENCES

- Block, N. (1978). 'Troubles with Functionalism', *Minnesota Studies in the Philosophy of Science* 9, pp. 261-325.
- Chalmers, D. (2000). 'What is a Neural Correlate of Consciousness?', in T. Metzinger, (ed.), *Neural Correlates of Consciousness*, pp. 17-39. MIT Press.
- Chalmers, D. (2005). 'The Matrix as Metaphysics', in C. Grau (ed.), *Philosophers Explore the Matrix*, Oxford University Press, pp. 132–76.
- Chalmers, D. (2012). *Constructing the World*. Oxford University Press.
- Chalmers, D. (2017). 'The Virtual and the Real', *Disputatio* 46, pp. 309-352.
- Chalmers, D. (2019). 'Three Puzzles About Spatial Experience', in A. Pautz and D. Stoljar (eds.) *Blockheads!*, pp. 109-38. Oxford University Press.
- Churchland, P. M. (1981). 'Eliminative Materialism and the Propositional Attitudes', *Journal of Philosophy* 78, pp. 67-90.
- De Haro, S. (2018). 'The Heuristic Function of Duality', *Synthese*.
<https://doi.org/10.1007/s11229-018-1708-9>.
- Gillett, C. (2003). 'The Metaphysics of Realization, Multiple Realizability and the Special Sciences', *Journal of Philosophy* 100, pp. 591-603.
- Goff, P. (2017). *Consciousness and Fundamental Reality*. Oxford University Press.
- Healey, R. (2002). 'Can Physics Coherently Deny the Reality of Time?', in C. Callender (ed.), *Time, Reality and Experience*, pp. 293–316. Cambridge University Press.
- Huggett, N. & Wüthrich, C. (2013). 'Emergent Spacetime and Empirical (In)coherence', *Studies in History and Philosophy of Modern Physics* 44, pp. 276-285.
- Huggett, N. (2017). 'Target Space \neq Space', *Stud. Hist. Phil. Mod. Phys.* 59, pp. 81-88.

Forthcoming in *Philosophy Beyond Spacetime*, eds. Wüthrich, Huggett & Le Bihan. OUP.

Knox, E. (2017). 'Physical Relativity from a Functionalist Perspective', in *Studies in History and Philosophy of Modern Physics*. SI: *Physical Relativity*, 10 years on. <https://doi.org/10.1016/j.shpsb.2017.09.008>.

Kripke, S. (1972). *Naming and Necessity*. Harvard University Press.

Lam, V. & Wüthrich, C. (2018). 'Spacetime is as Spacetime Does'. *Studies in History and Philosophy of Modern Physics* 64, pp. 39-51.

Le Bihan, B. (2018). 'Priority Monism Beyond Spacetime', *Metaphysica* 19, pp. 95-111.

Le Bihan, B. & Read, J. (2018). 'Duality and Ontology'. *Philosophy Compass*. 13:e12555. <https://doi.org/10.1111/phc3.12555>.

Lewis, D. (1972). 'Psychophysical and Theoretical Identifications', *Australasian Journal of Philosophy* 50, pp. 249-258.

Lewis, D. (1994). 'Reduction of Mind', first published in S. Guttenplan (ed.) *A Companion to the Philosophy of Mind*, Oxford: Blackwell (1994); repr. in D. Lewis, *Papers in Metaphysics and Epistemology*, Cambridge University Press (1999), pp. 291-324.

Lewis, D. (1995). 'Should a Materialist Believe in Qualia?', *Australasian Journal of Philosophy* 73, pp. 140-44.

Markopoulou, F. & Smolin, L. (2007). 'Disordered Locality in Loop Quantum Gravity States', *Classical and Quantum Gravity* 24, pp. 3813-3824.

Matsubara, K. & Johansson, L. (2018). 'Spacetime in String Theory: A Conceptual Clarification', *Journal for General Philosophy of Science* 49, pp. 333-353.

Maudlin, T. (2007). 'Completeness, Supervenience, and Ontology', *Journal of Physics A: Mathematical and Theoretical* 40, pp. 3151-3171.

Ney, A. (2012). 'The Structure of our Ordinary Three Dimensions in a Quantum Universe', *Noûs* 46, pp. 525-60.

Forthcoming in *Philosophy Beyond Spacetime*, eds. Wüthrich, Huggett & Le Bihan. OUP.

Ney, A. (2015). 'Fundamental Physical Ontologies and the Constraint of Empirical Coherence', *Synthese* 192, pp. 3105-3124.

Ney, A. (2017). 'Finding the World in the Wavefunction: Some Strategies for Solving the Macro-object Problem', *Synthese*. <https://doi.org/10.1007/s11229-017-1349-4>.

Paul, L. A. (2012). 'Building the World from its Fundamental Constituents', *Philosophical Studies* 158, pp. 221-256.

Rovelli, C. and Vidotto, F. (2015). *Covariant Loop Quantum Gravity*. Cambridge University Press.

Schwitzgebel, E. (2019). 'Kant Meets Cyberpunk', *Disputatio*. <https://doi.org/10.2478/disp-2019-0006>.

Sklar, L. (1976). *Space, Time, and Spacetime*. University of California Press.

Whittle, A. (2008). 'A Functionalist Theory of Properties', *Philosophy and Phenomenological Research* 77, pp. 59-82.

Wüthrich, C. (2017). 'Raiders of the Lost Spacetime', in Lehmkuhl, D., Schiemann, G. & Scholz, E. (eds.) *Towards a Theory of Spacetime Theories*, pp. 297-335. Springer.

Yates, D. (2012). 'Functionalism and the Metaphysics of Causal Exclusion', *Philosophers' Imprint* 12(13), pp. 1-25.

Yates, S. (2016). 'Demystifying Emergence', *Ergo* 3, pp. 809-841.

Yates, D. (2018). 'Inverse Functionalism and the Individuation of Powers', *Synthese* 195, pp. 4525-4550.