Eliminating Spacetime

forthcoming in Erkenntnis

Sam Baron*

Abstract

A number of approaches to quantum gravity (QG) seem to imply that spacetime does not exist. Philosophers are quick to point out, however, that the loss of spacetime should not be regarded as total. Rather, we should interpret these approaches as ones that threaten the *fundamentality* but not the *existence* of spacetime. In this paper, I argue for two claims. First, I argue that spacetime realism is not forced by QG; spacetime eliminativism remains an option. Second, I argue that eliminativism provides a useful framework for developing two existing approaches to the metaphysics of QG, involving functionalism and mereology respectively.

1. Introduction

Quantum gravity (QG) is the name of a broad program that seeks to reconcile our two best physical theories at the moment: general relativity (GR) and the standard model of particle physics. It is well known that these theories are in tension with one another. While it is possible to provide a quantum account of the various fields that populate the standard model, it is extraordinarily difficult to quantize the gravitational field in GR. A central goal of quantum gravity is to produce a quantum account of gravitation, one that works for the low energy regime currently described by GR as well as the high energy regime where GR breaks down.

A number of approaches to QG seem to carry a surprising implication: spacetime does not exist fundamentally.[1] This implication has been identified for a range of theories in different ways, including: string theory, loop quantum gravity, causal set theory, and canonical quantum gravity. Some philosophers maintain, however, that the apparent loss of spacetime within physics should not be regarded as total. Spacetime exists, they argue, just as a non-fundamental entity. If that's right, then the central task both for physicists engaged in developing the approaches to QG at issue, and for philosophers interpreting those approaches, is to provide an account of *how* spacetime exists as a non-fundamental or 'emergent' entity. That is, it must be shown "how relativistic spacetimes re-emerge ... from the fundamental structure as postulated by the theory at stake" (Wüthrich 2019, p. 2).

Call the view according to which spacetime exists: spacetime realism. Spacetime realism is to be contrasted with spacetime eliminativism.[2] The spacetime eliminativist

*Dianoia Institute of Philosophy, Australian Catholic University, email: samuel.baron@acu.edu.au

[1]See Huggett and Wüthrich (2013) for an overview.

[2]See Le Bihan (2018b) for more on realism and eliminativism.

denies that spacetime exists: spacetime, on this view, does not re-emerge in a metaphysical sense. Most of the discussion surrounding QG has focused on realism.[3] Explicit arguments against spacetime eliminativism are scarce, but there appear to be three such arguments available. In this paper I show that eliminativism is at least an option by addressing all three arguments. Having shown that spacetime eliminativism is an option, I then argue that it provides a useful framework for developing two existing metaphysical interpretations of QG: a functionalist interpretation, on the one hand, and a mereological interpretation, on the other. In both cases, situating these approaches within an eliminativist framework presents a way to address certain concerns that have been raised for each view.

To be clear, my aim is not to defend eliminativism. The point is just that we cannot rule eliminativism out: it remains an open possibility, and it constitutes a potentially fruitful way of thinking about QG. Whether spacetime eliminativism is the correct interpretation of any approach to QG is a matter I leave for another time.

2. Three Arguments

Must we be spacetime realists? In answering this question, I won't consider any specific approach to QG to see whether spacetime realism or spacetime eliminativism is most appropriate for that approach. Rather, I will focus on general considerations concerning the broad QG program. The idea is to see whether eliminativism is ruled out based on what we know about QG in general. Because the QG program is still very much a work in progress, with new theories likely to be developed in the future and old theories redeveloped, it is useful to consider whether, in a general sense, we should always be thinking in realist terms.

As noted, there are three arguments that seem to offer a mandate in favour of spacetime realism. In this section, I will introduce all three arguments, and provide a response in each case.

2.1. Deducibility

The first argument focuses on the inter-theoretic relationship between any theory within the broad QG program and GR. GR, it is thought, should be recoverable from the underlying physics of QG in a certain sense. As Crowther (2018, p. 75) puts the point:

> QG is understood as any theory that satisfies the set of criteria that are taken to define QG. Currently, there is no well-established, generally agreed-upon set; however, some of the criteria whose inclusion is the least controversial across all the approaches to QG concern the relationships to, and between, current theories. For instance, I take it that the set includes at least the following criteria: that the theory describe the domains where both GR and quantum theory are necessary; that the theory "recover" GR in the regimes where GR is known to be successful; and that the theory take into account quantum theory.

[3] Spacetime realists include Chalmers (forthcoming); Le Bihan (2018b); Lam and Wüthrich (2018); Wüthrich (2017, 2019); Yates (2021).

Recovery requires (at a minimum) derivability: GR is recoverable from QG only if GR can be derived from QG using various bridge laws or approximation procedures.[4]

The requirement for GR to be deducible from a theory of QG can be used to formulate an argument against spacetime eliminativism as follows. First, assume that deduction conducts ontological commitment, in this sense: for any two theories T_1 and T_2, if T_2 is derivable from T_1 and, moreover, we have good reason to believe in the existence of the entities posited by T_1, then we should also believe in the existence of the entities posited by T_2. Next, suppose that, for some approach to QG, we should believe in the existence of the entities posited by that theory. Then it follows that we should also believe in the entities posited by GR because GR can be derived from QG.

Of course, we don't yet know whether we should believe in the particular entities posited by an existing approach to QG, because we don't know which (if any) of the approaches to QG is the correct approach. But in a sense, it doesn't matter. Because, if Crowther is right, we *must* be able to derive GR from an underlying theory of QG, we have no choice but to accept the existence of spacetime, because we know that we will be forced to accept the ontology of some approach to QG at some point.

This argument against spacetime eliminativism can be summarised as follows:

1. For any theory T_1 and any theory T_2, if T_2 is derivable from T_1 then if the entities posited by T_1 exist, so do the entities posited by T_2.
2. GR is derivable from QG.
3. For some theory of QG, the entities posited by that theory exist.
4. Therefore, the entities posited by GR exist.

The problem with the argument lies with the first premise. To see the issue, let us set QG aside for a moment and consider a related example. The example involves the relationship between GR and a Newtonian theory, specifically Newton-Cartan theory (NC). NC is a reformulation of the core ideas underlying Newtonian mechanics. The basic idea is to take Poisson's field-theoretic formulation of Newton's law of gravitation and show that Newtonian gravitation expressed in this way can be captured in terms of a relationship between geometric curvature and matter distribution. NC is thus a way of reformulating Newton's account of gravitation that renders it analogous to GR. Having generalised Newton's theory in this manner, it is then possible to show that Newton's equations for gravitation constitute a mathematical limit of GR. We can thus derive NC from GR under certain conditions. In particular, for a group of test particles that are moving slowly with respect to the speed of light, and that are situated within a weak, static gravitational field, the GR description reduces to the description provided by the reformulated Newtonian theory.[5]

The ontological commitments of GR and NC appear to be distinct. GR is committed to the existence of a physical structure that can be represented by a differentiable manifold, M, equipped with a metric tensor field, g_{ab} (described by the field equations) which defines the geometry at each point in the manifold. NC, by contrast, describes a physical structure that corresponds to a differentiable manifold, M, equipped with two

[4]See Butterfield and Isham (1999) for discussion.

[5]See Carrol (2019) for details.

orthogonal metrics: τ—a temporal metric, which is used to assign temporal coordinates and distances to vectors on the manifold—and h—a spatial metric, which is used to assign spatial lengths. The manifold structure $\langle M, \tau, h \rangle$ is not isomorphic to the manifold structure of GR, namely $\langle M, g_{ab} \rangle$. The metric of GR is generally covariant, whereas the manifold structure of Newton-Cartan theory obeys the symmetries defined by the Galilea group. *Prima facie* this mathematical difference between the two theories translates into an ontological difference. The physical structure described by GR can be interpreted as a four-dimensional structure in which space and time are woven together. The physical structure described by NC, by contrast, presents a picture in which space and time are distinct, orthogonal dimensions.

Because NC is derivable from GR, we can formulate an analogous argument to the one outlined above in favour of spacetime realism. The argument moves from the assumption that the entities posited by GR exist to the conclusion that the entities posited by NC exist and thus toward what might be called Newtonian realism (by way of analogy with spacetime realism). This argument can be stated as follows:

1. For any theory T_1 and any theory T_2, if T_2 is derivable from T_1 then if the entities posited by T_1 exist, so do the entities posited by T_2.
2. NC is derivable from GR.
3. The entities posited by GR exist.
4. Therefore, the entities posited by NC exist.

This argument seems much less compelling than the argument against spacetime eliminativism, however. Assume, for the sake of argument, that we should accept the existence of the entities posited by GR. Even so, it seems at least permissible to adopt a form of Newtonian eliminativism, whereby the ontology of NC is eliminated in favour of the underlying physical structure described by GR. There is no requirement, that I can see, to believe in the existence of the entities posited by NC. Eliminativism in this case seems to be at least an option. This is not to say that we are forced to deny the existence of the Newtonian structures posited by NC either. It may be possible to make a philosophical case for realism about NC.[6] But the case must be made. We are not compelled to believe in the Newtonian entities posited by NC just because NC is deducible from GR. If that's right, however, then deduction is not in general sufficient for ontological commitment. This undermines the first premise in the argument for Newtonian realism. Since this is also the first premise in the argument involving GR and QG stated above, that argument fails as well.

Now, one might concede that the first premise in both arguments is false and thus that there is no general connection between deduction and ontological commitment, but maintain that in the specific case of GR and QG deduction is sufficient for ontological commitment nonetheless. In order to defend this approach one must also hold that deduction is not sufficient for ontological commitment in the case of GR and NC. This localised approach to the matter thus requires identifying some difference between the GR/QG case and the GR/NC case that might support a differential attitude toward deduction and ontology. Perhaps there is some such difference to be found: it may be

[6] Wallace (2020) seems to hold this view.

that the way deduction works in the case of GR and NC is substantially different to the deduction of GR from QG and perhaps such a difference translates into a difference in ontological commitment. I certainly cannot exclude this possibility. But no such difference has been identified yet, and so we are not in a position to rule eliminativism about spacetime out via deduction alone.

Of course, one might argue that realism about the entities posited by NC is, in fact, forced, and that I am just wrong in my assessment of the Newtonian case. I will provide a reply to this broad line of thought in a moment. Before I do so, however, let us turn to the second argument against spacetime eliminativism, since the relationship between GR and NC is important for that argument as well.

2.2. Empirical Adequacy

Like the first argument against spacetime eliminativism, the second argument rests on a theoretical relationship that is expected to hold between QG and GR. Rather than a relation of deduction, however, the relationship is one of empirical recovery. Empirical recovery involves showing that the confirmed empirical predictions of one theory, T_1, can be reproduced by another theory, T_2. Often, the predictions can only be reproduced by holding T_2 under certain constraints, or by applying it only to certain situations. In the case of GR and QG it is expected that every confirmed empirical prediction of GR will be reproduced by a theory of QG. This generally means that the confirmed empirical predictions of GR ought to be a subset of the empirical predictions of a theory of QG within the low energy regime where GR is successful.

Empirical recovery is related to deduction, but they are not necessarily the same thing. A deductive relationship is likely to ensure some measure of empirical recovery between theories. The deducibility requirement considered above can thus be viewed as one way of meeting a deeper demand for empirical recovery. Empirical recovery does not, however, require deduction. A theory can generally reproduce the empirical predictions of another theory without entailing it. Like deduction, however, one might argue that a relationship of empirical recovery between two theories has ontological implications. Specifically, when the confirmed empirical predictions of one theory are recovered by another theory, then if we accept the existence of whatever is posited by the recovering theory, we are forced to also accept the existence of whatever is posited by the recovered theory. Using this general idea we can formulate an argument against spacetime eliminativism based on the empirical relationship expected to hold between QG and GR:

1. For any theory T_1 and any theory T_2, if every confirmed empirical prediction of T_2 is also an empirical prediction of T_1, then if the entities posited by T_1 exist, so do the entities posited by T_2.

2. Every confirmed empirical prediction of GR is also an empirical prediction of QG.

3. For some theory of QG, the entities posited by that theory exist.

4. Therefore, the entities posited by GR exist.

A similar relationship of empirical recovery holds between GR and NC. Like GR, Newton's account of gravitation is immensely successful within a certain regime (as noted: for objects situated in weak, static gravitational fields, that are moving slowly with respect to the speed of light). Indeed, this is one of the central reasons we continue to use Newtonian physics for a range of different real-world applications. Now, GR is capable of reproducing all of the empirical successes of a Newtonian account of gravitation. Within the relevant regime, every confirmed prediction of Newton's account of gravitation is a confirmed prediction of GR as well. Because NC just is a reformulation of Newton's account of gravity, it inherits the same class of confirmed empirical predictions, and thus stands in the same relationship of empirical recovery to GR. We can thus formulate an analogous argument to the one above, except that this argument is designed to rule out eliminativism about the entities posited by NC:

1. For any theory T_1 and any theory T_2, if every confirmed empirical prediction of T_2 is also an empirical prediction of T_1, then if the entities posited by T_1 exist, so do the entities posited by T_2.

2. Every confirmed empirical prediction of NC is also an empirical prediction of GR.

3. The entities posited by GR exist.

4. Therefore, the entities posited by NC exist.

As with the deduction-based argument against Newtonian eliminativism, this recovery-based argument is not compelling. As already discussed, realism does not seem to be required in the case of NC. If realism about the entities posited by NC is not forced here, however, then it follows that empirical recovery is not, in general, sufficient for ontological commitment. Accordingly, the first premise in the empirical recovery argument appears false, and thus the related argument against spacetime eliminativism that uses the same premise fails.

As before, one might respond that spacetime realism is, in fact, required by the relationship of empirical recovery between GR and QG, even though relationships of this kind are not generally sufficient for ontological commitment. In order to hold the line and rule out spacetime eliminativism, some difference must, once again, be identified between the GR and QG case on the one hand, and the GR and NC case on the other. For instance, one might argue that the number of empirical predictions that must be recovered in the case of NC is much smaller than the number to be recovered in the case of GR. Or, in the same vein, one might argue that the types of empirical predictions to be recovered in the case of NC are different in kind to those that must be recovered for GR. At present, however, we lack an account of how many or what types of confirmed empirical predictions must be recovered for realism. In the absence of a convincing story along these lines, spacetime eliminativism remains open.

2.3. Perfect and Imperfect Realism

In a moment, I will consider the third and final argument against spacetime eliminativism. My response to that argument does not rely on an analogy between QG/GR, and GR/NC. Before proceeding, then, it is worth briefly considering a response to the proposed parallel between the two cases. I have claimed that we are not forced to adopt a realist attitude toward the entities posited by NC; eliminativism is at least open in this case. As noted, however, one might simply disagree with this assessment: realism is required in the case of NC after all.

At first glance, this kind of response might not seem all that plausible. NC has, after all, been fully superseded by GR. Precisely what we learn in the shift from Newtonian mechanics to relativistic physics one might argue is that we have no need for the broadly Newtonian structures posited by NC. Indeed, one might go a bit further and argue that realism is not even an option in the case of NC, let alone a requirement. We should, rather, be realists about GR, and eliminativists about NC despite the inter-theoretic relationships that obtain between the two theories.

This assessment of the situation, however, is too quick. To see why, it is important to draw a distinction between two kinds of realist attitude that one might adopt toward a given theory: perfect realism and imperfect realism (cf. Chalmers (forthcoming)). According to perfect realism, the entities posited by a theory T exist exactly as that theory describes them. Perfect realism requires there to be an exact isomorphism between the structure of a given theory and the structure of the world. Imperfect realism, by contrast, eschews perfect matching between a theory and the world, and thus requires at best a partial isomorphism between theory and world. According to imperfect realism, we can adopt a realist attitude toward a theory, and thus admit that the entities named by the theory exist, while denying that those entities are exactly as the theory describes them. The imperfect realist thus allows that what exists can be approximated by a theory, and that this approximation is sufficiently close to warrant the claim that what the theory posits exists.

In the case of GR and NC, perfect Newtonian realism does seem to be false. The entity posited by GR is not perfectly described by NC. A physical entity corresponding to the metric structure of GR is, at best, approximately described by NC, and only under very specific circumstances. If one adopts a form of perfect realism about NC anyway, then it seems that one is forced to 'double' one's ontology in an unattractive manner, by accepting the existence of the ontology of GR plus some new entity that is perfectly described by the Newtonian theory. Not only would positing such a structure bring no extra benefits from the perspective of physics, positing the structure would incur a significant theoretical burden. One would need to explain how the extra Newtonian entity and the relativistic structure of GR relate to one another.

Imperfect Newtonian realism, by contrast, does not face the same difficulties. The imperfect Newtonian realist does not 'double' the ontology. The imperfect Newtonian realist maintains that there exists some physical structure—whatever structure is needed for GR—and this physical structure is approximately described by the Newtonian theory. The approximation is sufficiently close to warrant a realist attitude toward the Newtonian theory. In this way, the entities posited by NC exist well-enough for realist purposes.

The central reason, then, why we should not be too quick to adopt an eliminativist attitude toward NC is that imperfect realism about NC might be true. We thus cannot dismiss the idea that realism about NC might be required on the grounds that perfect realism is false or otherwise implausible. The strongest form of the objection outlined at the beginning of this section, then, holds that imperfect realism is forced in the case of NC. If imperfect realism about NC is required, then the entities posited by NC exist well-enough and the parallel between the QG/GR case and the GR/NC case fails to undermine the first premise in each of the two arguments against spacetime eliminativism considered above. If, by contrast, NC eliminativism remains open even against the imperfect realist option, then the parallel between the QG/GR case and the GR/NC case continues to be probative, in so far as it still shows that spacetime eliminativism cannot be ruled out by deduction or empirical recovery alone.

The question before us, then, is whether imperfect realism about NC really is required. Now, as noted, I am willing to concede that imperfect realism about NC is open. But I maintain that we simply don't know enough about how imperfect realism works to warrant the claim that we *must* be imperfect realists about NC. Imperfect realism relies crucially on the notion of approximation, and closeness. The imperfect realist maintains that some x described by a theory T exists well-enough so long as there exists some y that is approximately described by T and the approximation is sufficiently close.

This last part is particularly important. While imperfect realism allows for a certain degree of approximation, even imperfect realism must have its limits. When the difference between what exists and the description provided by a given theory is too great, realism no-longer seems appropriate. The difficulty is that we currently lack a general account of how close is close enough for realist purposes when it comes to imperfect realism. What would such an account look like? Well, imperfect realism relies on a notion of similarity. Because similarity comes in degrees and can be measured in different respects, an account of how close is close enough for realism requires two things. First, it requires the specification of a threshold: a degree of similarity such that only similarity to that degree or greater implies realism. Second, some account must be given of the dimensions of similarity that matter when it comes to imperfect realism including, presumably, an account of any weights applied to those dimensions.

Without an account along these lines—and to my knowledge one is yet to be provided—it is difficult indeed to make the case that imperfect realism about NC is required. Moreover, providing an account of imperfect realism along the right lines seems to be a difficult task. Any such account will need to walk the line between too much similarity and not enough similarity. For if one specifies the requirements for closeness too tightly then imperfect realism will be too hard to establish; if one specifies the requirements for closeness too loosely, imperfect realism will be too easy. As matters stand, then, eliminativism about NC cannot be ruled out. If that's right, however, then the GR/NC case continues to serve as a counterexample to the first premise in each of the two arguments against spacetime eliminativism considered above.

Indeed, what the discussion of imperfect realism shows us is that there is a rather direct response to those arguments available. First, note that QG puts perfect realism about GR under threat as well. Indeed, it is generally expected that GR will be at best an approximate description of whatever structure is ultimately posited by QG. In so far as any inter-theoretic relationship between QG and GR can force a realist attitude toward the ontology of GR, it can at best force a commitment to imperfect realism. As noted, however, there is no agreement on how close is close enough for imperfect realism, or on what the relevant degrees of similarity might be that matter for realism. So we are not in a position to know whether either deduction or empirical recovery force a realist attitude of this kind in the case of QG and GR. It really depends on exactly how imperfect realism works, and on the account one gives of closeness.

The analogy between the QG/GR case and the GR/NC case is not strictly required for this direct response to work. The direct response relies only on our lack of understanding concerning imperfect realism. This, one might worry, makes the GR/NC case into an idle wheel in much of the discussion thus far. While the availability of the direct response does tend to diminish the importance of the GR/NC case, that case still has a role to play in the current dialectic. What the GR/NC case reveals is just how limited our understanding of imperfect realism really is. We don't know enough about imperfect realism to be able to rule out eliminativism via inter-theoretic relations between theories in even fairly well-understood cases, where we have a deep understanding of the two theories at issue. This tends to cast doubt on our prospects for being able to rule out eliminativism in much harder cases, such as the relationship between QG and GR.

The GR/NC case can also be used to fend off a further objection. I have argued that we don't know enough about imperfect realism to be able to rule out spacetime eliminativism. One might worry, however, that the problem doesn't have much to do with imperfect realism. The problem, or rather the appearance of a problem, is simply due to the the fact that I have abstracted away from any particular theory of QG. From such an abstract standpoint it is indeed difficult to tell whether imperfect realism is required, but this is unsurprising. It is only by looking at a specific theory of QG that a determination about spacetime realism can be made.

The GR/NC case allows me to sidestep this objection. In the GR/NC case we are dealing with particular theories, and theories that we understand quite well, and still it is hard to make the case that imperfect realism is forced. This suggests that the problem lies with our understanding of imperfect realism, rather than our understanding of particular theories. Filling out the discussion with a specific approach to QG won't clearly make a difference. Besides, if it is the case that whether imperfect realism is true can only be determined by first identifying a specific theory of QG and focusing on that, then it is difficult to see how we could be in a position to rule out eliminativism anyway. It really depends on what the correct theory of QG turns out to be, and how similar its ontology ultimately is to spacetime.

One final point before I move on. Once we have shifted the discussion to focus on imperfect realism about spacetime, the distinction between realism and eliminativism seems much less significant. The distinction is, at best, a matter of degree. Realism is true if what exists is sufficiently close to spacetime; eliminativism is true if what exists is not sufficiently close. We may find that there are two quite similar approaches to QG that fall on different sides of the realist/eliminativist divide, depending again on what exactly it takes for imperfect realism to be true. This, in itself, makes it difficult to see how eliminativism could be ruled out or even why we would want to. It also makes eliminativism seems much less objectionable. In considering the distinction between realism and eliminativism we are not necessarily considering two radically divergent pictures of reality. We may, in the end, be considering two quite similar ontologies.

2.4. Empirical Coherence

I come now to the third and final argument against spacetime eliminativism. Whereas the first and second arguments focus on inter-theoretic relationships between GR and QG, the third argument focuses directly on QG, and on a specific constraint that has been proposed for the entire QG program. The constraint is one of empirical coherence. Following Barrett (1999), a theory is empirically incoherent when the *truth* of that theory undermines any prospect for empirically justifying it.

According to Huggett and Wüthrich (2013), approaches to QG in which spacetime is absent at the fundamental level face a potential threat of empirical incoherence. On one way of interpreting this threat, it is possible to formulate an argument against spacetime eliminativism (whether this is how Huggett and Wüthrich intend for the threat of empirical incoherence to be formulated is an issue I return to below).

Here's the idea. In order to empirically confirm a theory, one must gather observations that provide evidence for the theory in question. Suppose, however, that observation is linked to spatiotemporal location, in this sense: it is possible to observe an entity only if that entity is located in a region of spacetime. Now, suppose that, according to some theory T, spacetime does not exist at all. Then it follows that nothing is located in a region of spacetime according to that theory because there are no such regions. From this it follows that there is nothing to be observed. In this situation, T is empirically incoherent: if the theory is true, then it cannot be empirically justified. Conversely, a theory is empirically coherent only if, according to that theory, spacetime exists.

Based on this line of reasoning, a third argument against spacetime eliminativism can be stated as follows:

1. A theory T is empirically coherent only if, according to T, spacetime exists.

2. A theory of QG is viable only if it is empirically coherent.

3. So, a theory of QG is viable only if, according to that theory, spacetime exists.

Assuming that theories of QG are viable and thus empirically coherent, it follows that spacetime must exist according to those theories. Notice that this third argument sidesteps any need to provide a full account of imperfect realism. No matter how imperfect realism is ultimately spelled out, realism about spacetime must be true if any theory of QG is to be viable. This third argument does not therefore suffer from the issues identified with imperfect realism in the previous section.

The argument fails for a different reason. The particular picture of observability needed to make the argument work is implausible. On this picture, observation can only occur in the presence of spacetime. But then this would have the bizarre consequence that *any theory* that does not imply the existence of spacetime is empirically incoherent, not just a theory of QG. Given that spacetime appears so late in the history of science, it would follow that every theory before roughly the 20th century was empirically incoherent, since those theories did not carry such an implication. But that's implausible.

It is also doubtful that this is the picture of empirical coherence that Huggett and Wüthrich endorse.[7] For them, observation is a matter of observing local beables. The notion of a local beable, however, is not tied to spacetime. A local beable, very roughly is just an entity located in space that can be observed at a particular time (cf. Bell (1987)). Huggett and Wüthrich argue that for a theory to be empirically coherent, it must be able to support observations of local beables in this sense. This requirement ties observation to the existence of entities located in space and time, but not to location in spacetime. Their concern, then, is that theories of QG that don't employ spacetime at the fundamental level also seem to lack any fundamental spatial or temporal structure. Such theories therefore don't seem to have the capacity to underwrite observation, but not, in the end, because spacetime goes missing. Rather, it is because space and time are placed under threat.

[7] This was noted by a referee.

The true threat of empirical coherence thus relies on a conception of observability that is linked to space and time, not spacetime. Once this is realised, it becomes clear that there is no straightforward case to be made against spacetime eliminativism based on the threat in question. Even if it is correct that the approaches to QG that Huggett and Wüthrich focus on lack space and time at the fundamental level, it does not follow that a metaphysical commitment to emergent spacetime is required to secure the empirical coherence of those approaches. In order to secure empirical coherence one needs only an emergent structure of locations that possess spatial and temporal properties.

Now, it could be argued that, within QG at least, the existence of some physical structure in which spatial and temporal locations may be defined just is sufficient for the existence of spacetime.[8] The idea being that any emergent physical structure that can provide spatial and temporal locations for beables (and thus support empirical coherence) must also have the same observational consequences of spacetime. That's because any such emergent structure would be operating at the level of description corresponding to GR, and so should be empirically indistinguishable from spacetime itself.

This line of thought, however, brings us back to the argument from empirical adequacy and thus to the issues identified above for imperfect realism. To defend the view that the existence of something empirically equivalent to spacetime implies realism, one must provide an account of how close is close enough for realist purposes when it comes to spacetime. That, I have noted, is something we don't yet have, and so it remains open that even if there exists a fundamental structure with emergent spatial and temporal properties, spacetime does not exist. We should be very careful about moving between realism regarding space and time and realism regarding spacetime for this reason.

3. An Eliminativist Framework

Thus far I have argued that spacetime eliminativism remains open as a way of interpreting QG. This result is of limited significance, however, if eliminativism is false or uninteresting. As discussed in §1, my goal is not to argue that eliminativism is true. I will, however, argue that the view serves as a potentially useful framework for developing two existing approaches to the metaphysics of QG. The view is, if nothing else, worth considering.

The first approach to the metaphysics of QG I will consider appeals to mereology. The second approach appeals to functionalism. Both approaches are typically stated against a realist backdrop. I will show that reformulating these accounts using spacetime eliminativism enables one to avoid concerns that have been raised for each view.[9]

3.1. Mereology

I will start with the mereological approach. Roughly speaking, according to this approach, QG describes a fundamental ontology of non-spatiotemporal parts which, in various combinations, build up the rest of reality. A view along these lines is attractive because of the role that mereology plays in the rest of science. Mereological relations are often found 'connecting the levels', as it were (such as the macro and micro levels), and so it is natural to suppose that they connect the more fundamental entities described by QG

[8] For instance, Le Bihan and Linnemann (2019) take the presence of space and time to justify the existence of spacetime in a 'minimal' sense.

[9] These views are not necessarily in competition, and may be held together.

to other, less fundamental entities. For this reason, the mereological approach is rather elegant. It proposes to analyse all transitions from the more fundamental to the less fundamental in the same broad metaphysical terms, and thus paints a unified picture of the structure of physical reality.[10]

One way to develop a mereological view along these lines involves adopting spacetime realism. On this view, spacetime exists and it is composed of a range of non-spatiotemporal parts that are described by a theory of QG. We can thereby do justice to the way in which the fundamental ontology is not spatiotemporal, while nonetheless making room for the existence of spacetime by linking it to some more fundamental structure.

Baron (forthcoming) raises an objection against a mereological picture of QG along these lines. The problem, roughly speaking, is that there are various plausible constraints on the mereology of physical entities that seem to undermine the idea that spacetime might be composed of non-spatiotemporal parts. In particular, Baron focuses on the following four mereological principles:

> **Inheritance of Location:** If x is part of y, then y is located wherever x is located.
>
> **H5:** For any x and any y, x is a part of y iff x's location is a subregion of y's location.
>
> **Smaller Than:** Proper parts are smaller than the wholes they compose, where for any x and y, x is smaller than y iff there is a region r at which x is exactly located that is a proper sub-region of the region r∗ at which y is exactly located.
>
> **Compositionality of Extension:** For any spatiotemporal object y and for any $x_1 ... x_n$, if $x_1 ... x_n$ compose y, then the spatiotemporal extent of y is a function of the spatiotemporal extensions of $x_1 ... x_n$ and the spatiotemporal relations between $x_1 ... x_n$.

The inheritance of location is drawn from work by Sider (2007), who maintains that it is one of the core aspects of parthood that differentiates it from identity. H5, by contrast, is one of the principles of mereological harmony identified by Saucedo (2011), which conceptually tie mereology and location together. The smaller-than principle is similar to a principle discussed by Donnelly (2011), who takes it to be a core feature of the mereology of physical objects. The compositionality of extension, is so far as I can tell Baron's own, but it seems plausible enough.

Baron argues that all four principles are false if spacetime has non-spatiotemporal parts. This is easy enough to see in a rough sense. Each principle transmits location either upward from parts to wholes, or downward from wholes to parts. Suppose, then, that spacetime exists. If spacetime exists, then so do spacetime regions. Each spacetime region, on the mereological picture described above, will be composed of non-spatiotemporal parts, because all of spacetime is. Each spatiotemporal region, however, occupies a particular location within the spacetemporal manifold.

[10] The mereological approach is developed by Le Bihan (2018a,b) in the context of QG. See also Ney (2020) and Paul (2012).

Each region is itself spatiotemporally located. If each region is spatiotemporally located and none of its parts are, then it follows that location cannot be transmitted from part to whole or from whole to part. That is sufficient to undermine each of the four 'transmission' principles stated above.[11]

Notice, however, that Baron's argument against this broad mereological picture only really works if there are wholes that have a spacetime location and those wholes have parts that lack a spacetime location or vice versa. In short, what is needed to falsify each of the four principles above is a difference in location between parts and wholes. If parts and wholes are located in the same way, even if that way is not a spatiotemporal way, then all four principles can be sustained. Location is free to transmit from parts to wholes and back again. Baron's argument can be avoided entirely, then, if we simply give up on the existence of spacetime, and thus on the notion that anything is spatiotemporally located. Once this is done, we are free to accept that the more fundamental entities described by a theory of QG can compose less fundamental entities. This can be done without sacrificing even a single mereological principle involving location.

Of course, for this to work, we need a viable notion of location in QG such that parts and wholes can both be located in the same way. It is possible, then, that even within the eliminativist framework, a version of Baron's worry can be reformulated. For as we saw when discussing the threat of empirical incoherence, there seems to be some pressure to admit the existence of locations with spatial and temporal properties at the non-fundamental level. As Huggett and Wüthrich (2013) argue, however, space and time may go missing at the fundamental level in QG. But a transition from parts that are not spatially or temporally located to wholes that are will undermine the four mereological principles that Baron identifies once again. Accordingly, the mereological, spacetime eliminativist picture likely requires the presence of locations with spatial and temporal properties at both fundamental and non-fundamental levels.

This need not sink the eliminativist proposal, however. For while it is plausible that spacetime is missing at the fundamental level in QG, it is ultimately less clear that space and time are lost. Indeed, as Le Bihan and Linnemann (2019) argue, a division between space and time may persist for a range of approaches to QG. The core of their argument focuses on Lorentz symmetry. The Lorentz group features an in-built asymmetry which can be interpreted as a split between two orthogonal vectors, that appear to be broadly spatial and temporal in nature. The presence of Lorentz symmetries is thus often associated with a division between space and time. According to Le Bihan and Linnemann, many approaches to QG currently available still obey the Lorentz symmetries, and so implement a division between space and time at the most basic level.

I admit, however, that the existence of fundamental spatial and temporal properties remains controversial in QG, and note this as a limitation of the mereological, spacetime eliminativist framework. I also recognise that adopting spacetime eliminativism is not the only way to address Baron's worry. One might, for instance, develop a mereology

[11] If a spacetime region R has parts $p_1 ... p_n$ that are not spatiotemporally located, then **Inheritance of Location** is false: the p_n are all parts of R and yet none are located where R is. **H5** is also false: the p_n occupy no spatiotemporal regions and so *a fortiori* occupy no sub-region of R. **Smaller than** fails for much the same reason: some of the p_n will be proper parts of R. However, none of the p_n are smaller than R because none of them occupy any sub-region of R. Finally, the spatiotemporal extent of R is not a function of the spatiotemporal extent of the p_n because they are not spatiotemporally extended, and so **Compositionality of Extension** is false.

that does not constrain parthood with principles that enforce the transmission of location between parts and wholes. Still, it remains an advantage of the eliminativist approach that—in principle at least—it presents a path toward adopting a mereology that includes a broad range of mereological principles.

3.2. Functionalism

This brings me to the second broad approach to QG: spacetime functionalism. The basic idea behind spacetime functionalism is that there are certain functional roles associated with spacetime, and these roles are filled by the entities described by a theory of QG. Spacetime functionalism has been proposed by a number of philosophers, and there are various forms of the view available. My focus, however, is on the specific functionalist approach advocated by Lam and Wüthrich (2018). Lam and Wüthrich's version of spacetime functionalism is modelled closely on functionalism in the philosophy of mind. This version of functionalism, they maintain, generally involves the following two stages:

> (FR-1) The higher-level properties or entities, which are the target of the reduction, are 'functionalized', that is, they are given a functional definition in terms of their causal or functional role.
>
> (FR-2) An explanation is provided of how the lower-level properties or entities can fill this functional role. (Lam and Wüthrich, 2018, p. 43)

Applied to the case of spacetime, (FR-1) involves providing a functional specification for spacetime, or for spatiotemporal properties such that if something stands in the relations or possesses the properties, then that thing functionally realises spacetime. Having provided a functional specification for spacetime, the next step is to show that, for a given approach to QG, there is something in the ontology of the theory that stands in the relevant relations or possesses the relevant properties. In this way, spacetime may be functionally realised by a physical structure described by a theory of QG, in much the same manner that a mental state, like pain, might be realised by the physical structure described by some neurophysiological theory.

According to Yates (2021), spacetime functionalism faces a dilemma. In what follows, I will offer a reconstruction of Yates' argument.[12] The reconstructed argument differs from Yates' version in some respects, but the basic idea is the same. First, some set-up is in order. To begin with, it is important to draw an ontological distinction between *fundamental* and *non-fundamental* entities. A non-fundamental entity is an entity that depends on something for its existence. A fundamental entity, by contrast, is an entity that does not depend on anything for its existence.

The notion of fundamentality I employ is *theory-relative*. Thus, an entity x is fundamental in my sense when it is fundamental according to a theory T. Which is to say that by the lights of T, there is nothing that x depends upon for its existence. Similarly, x is non-fundamental according to a theory T when, according to T, there is something that x depends on for its existence.

Next, suppose that for some theory of QG, no fundamental entities in the ontology of that theory have spatiotemporal properties. Rather, what is fundamental according to

[12] I am grateful to a referee for pressing me to clarify this argument.

that theory is some non-spatiotemporal entity or entities. Finally, assume that there is something in the ontology of that theory that plays the spacetime role.

Here, then, is the dilemma: either what plays the spacetime role in the theory of QG at issue is fundamental or it is non-fundamental. If what plays the spacetime role is *fundamental* then it follows that the theory of QG bears an ontological commitment to spacetime as a fundamental entity. This is true, at least, for the two main forms of functionalism: realiser functionalism and role functionalism. According to *realiser* spacetime functionalism, spacetime is identical to whatever plays the spacetime role. Thus if that entity is fundamental, then so is spacetime. According to *role* spacetime functionalism, by contrast, spacetime is to be identified with a property possessed by a physical structure, namely that of being in a state that plays the spacetime role. If some fundamental entity is in such a state—and it should be if it plays the spacetime role—then that entity has spatiotemporal properties understood in the relevant sense (as higher-order functional properties).

Since we are assuming that the theory of QG at issue does not include any fundamental entities with spatiotemporal properties, we have a problem: no fundamental entity can play the spacetime role. So some non-fundamental entity must be doing this work. This brings us to the second horn of the dilemma. If what plays the spacetime role is some non-fundamental entity or entities, then a further question arises. Namely, what is the relationship between the fundamental non-spatiotemporal entities and the non-fundamental spatiotemporal entities? To be sure, the relationship is one of dependence. But the question is *how* do the spatiotemporal entities at issue depend, for their existence, on more fundamental, non-spatiotemporal entities? Call this: the dependence question.

In light of this question, there is a stronger and a weaker way to put the second horn of the dilemma. The stronger way is to maintain that answering the dependence question is what spacetime functionalism was originally introduced to do. So if some other answer to that question is needed, then functionalism is otiose. This way of putting the point is too strong, however. For even if a further answer to the dependence question is needed, functionalism is still useful. What the functionalist has shown is that there is something in QG that plays the spacetime role, even if the connection between that entity and the more fundamental entities within that theory's ontology must be further explained.

The weaker way of putting the second horn of the dilemma is just to point out that answering the dependence question is difficult. For it is unclear how non-spatiotemporal entities might give rise to spatiotemporal ones. To see *why* this question is difficult, consider the philosophy of mind case. The ontology of a neurophysiological theory contains individual neurons which are fundamental, and non-fundamental neural states which are large assemblages of neurons configured in a certain way, and that depend on neurons for their existence. Now, functionalism about the mind will identify something in the ontology of the neurophysiological theory with mental states. Clearly, we don't want to identify individual neurons with mental states. So we should assume that it is non-fundamental neural states that realise mental states.

In the mental state case we also have to explain the relationship between non-fundamental neural states and individual neurons. But doing so appears straightforward: we can simply appeal to mereology and say that individual neurons compose the neural states that realise mental states. A similar move is available for other cases of functional reduction in science, such as the functional reduction of fluids to groups of molecules. Here too, we can say that fluids are composed of molecules in certain configurations.

The trouble in the case of spacetime functionalism and QG is that the mereological option is not clearly available. This is essentially because of the problem discussed in §3.1. It is not clear that our standard mereological tools will work to connect spatiotemporal entities with non-spatiotemporal ones. It is also unclear how else we might answer the dependence question in this case.

Of course, that the dependence question is *difficult* does not mean it is *intractable*. Still, it would be nice to avoid it. If spacetime eliminativism can be combined with a functionalist approach to QG, then the dependence question can be avoided. For, clearly, if spacetime does not exist in any sense, then it can hardly be objected that we lack an account of how it exists.

At first glance, however, it is difficult to see how one might combine a functionalist approach to QG with a form of spacetime eliminativism. The trick to combining these positions is to note that functional realisation can come in degrees. Often a functional specification of some P will involve the identification of a range of different properties and relations which, together, fill out the functional role. This is true in the case of pain, and equally true in the case of spacetime, where the functional specification itself is likely to have a number of moving parts. A form of spacetime eliminativism can be upheld then by, first, providing a functional specification for spacetime and, second, maintaining that the full functional specification is not, in fact, satisfied by any entity in QG. Rather, the functional specification at issue is satisfied in an attenuated sense only: a number of the properties or relations associated with the spacetime role are possessed by something described by a theory of QG, but not enough to establish realism.

Now, it pays to be careful here. Recall the distinction between perfect and imperfect realism introduced above. The perfect realist maintains that for some P named by a theory T, P exists exactly as that theory describes it. The imperfect realist, by contrast, maintains that for some P named by a theory T, P does not exist exactly as that theory describes it, but it exists well-enough to say that there are Ps. Perfect and imperfect realism apply in the case of functionalism as well. A perfect realist who endorses a functional account of P, will maintain that P exists only if the full functional role for P is satisfied. An imperfect realist, by contrast, will allow for a certain degree of approximation with respect to the satisfaction of a functional role for P. So long as enough of the functional specification is realised, we can say that P's are functionally realised well-enough, and thus that they exist.

By saying that the functional role for spacetime is partially filled, I am not simply committing to imperfect realism. The idea, rather, is that while some part of the functional specification of spacetime is realised, not enough is satisfied to force a commitment to even imperfect realism. Such a view gains whatever benefits are associated with partially realising the functional role for spacetime, without the need to adopt a realist attitude toward spacetime.

As was the case with imperfect realism above, the difference between an imperfect realist approach to spacetime functionalism and an eliminativist approach is a matter of degree. As noted, this makes the distinction between spacetime eliminativism and imperfect realism less impressive than it might otherwise be, and this is no less true in the functionalist case. In the present context, however, the distinction is substantive enough to avoid having to answer the dependence question. And so falling on the eliminativist side of the distinction carries with it a distinct advantage.

What exactly are the benefits associated with partially realising the functional role for spacetime? Well, the chief reason Lam and Wüthrich recommend a functionalist approach to QG in the first place is to provide a solution to the problem of empirical incoherence. This is just the problem, discussed in §2.4, of explaining how a theory of QG can be empirically coherent given that space, time and spacetime do not seem to exist at the fundamental level. As already argued, however, the existence of spacetime is not necessary for solving this problem. So long as a theory admits of a well-defined notion of location—one that is linked to observability—the relevant notion of location need not be spatiotemporal in nature. Translated into a functionalist key, what this means is that the full functional specification for spacetime need not be satisifed in order for the problem of empirical incoherence to be met. All one needs to show is that there exists some physical structure that is capable of realising those functions of spacetime that relate to location.[13]

We can put the point in terms of a localising role. One of the functions of spacetime, we may suppose, is to provide locations for entities. A functional specification for spacetime is thus likely to include this aspect, but is likely to include a good deal more in order to fully flesh out what it is to be spacetime. From the perspective of solving the problem of empirical incoherence, however, all we need to do is show how the localising role of spacetime can be played by something within the ontology of a theory of QG. Showing that every aspect of the full functional specification for spacetime is filled, or even most aspects are filled, is surplus to requirements. To solve the problem of empirical incoherence, we don't need the functional specification for spacetime to be satisfied up to the point where imperfect realism is then true.

Note that the view I am proposing is compatible with much of what Lam and Wüthrich say. When setting out their functionalist approach, they note that "spacetime need not be fully recovered in some strong ontological sense ... in order to provide the grounds for empirical evidence ... but only certain functionally relevant features" (Lam and Wüthrich, 2018, p. 40) which seems broadly in line with the view I am proposing. They also focus on the localising role of spacetime in much the way I have suggested when they demonstrate their functional analysis. Indeed, they suggest that a non-spatiotemporal notion of location may be present at the level of GR already, noting that "localization ... may primarily be understood as relational and dynamical rather than fundamentally spatiotemporal." (Lam and Wüthrich, 2018, p. 47).[14] That being said, in more recent work Lam and Wüthrich (2020) appear to defend spacetime realism, maintaining that spacetime functionalism understood in a realist sense provides a way to connect QG to "a straightforwardly realist understanding of GR". As they recognise, this places them in danger of falling on Yates' dilemma, a problem that functionalism, if interpreted in an eliminativist spirit, can avoid.

4. Conclusion

I have not argued that spacetime eliminativism is true. Not for QG in general, nor for any particular approach to QG. In fact, I suspect that establishing eliminativism is

[13]This may require showing how spatial and temporal properties can be realised. Functionalist spacetime eliminativism may therefore be limited in so far as it requires the existence of *fundamental* spatial and temporal properties (on pain of facing a new dependence question about how spatial and temporal entities arise from non-spatial, non-temporal ones).

[14]This was pointed out by a referee.

about as difficult as establishing imperfect realism. In order to know whether a realist or eliminativist attitude is warranted, we must first limn the boundaries of realism. A better understanding of imperfect realism should thus be high on the agenda for everyone. Still, as I have argued here, spacetime eliminativism is at least a possible way of interpreting QG, and doing so yields certain benefits when it comes to two specific metaphysical approaches that one might take. We should thus take seriously the idea that spacetime may not survive the unification of GR and the standard model, not fundamentally, and perhaps not at all.

Acknowledgements: I am grateful to Baptiste Le Bihan, Kristie Miller and Jonathan Tallant for useful discussion of eliminativism. Research on this paper was funded by two grants from the Australian Research Council, a Discovery Early Career Researcher Award (DE180100414) and a Discovery Project (DP180100105).

References

Baron, Sam (forthcoming), "Parts of spacetime." *American Philosophical Quarterly*.

Barrett, Jeffrey A. (1999), *The Quantum Mechanics of Minds and Worlds*. Oxford University Press, New York.

Bell, John S. (1987), *Speakable and Unspeakable*. Cambridge University Press, Cambridge.

Butterfield, Jeremy and Chris Isham (1999), "On the emergence of time in quantum gravity." In *The Arguments of Time*, 111–168, Oxford University Press, Oxford.

Carrol, Sean (2019), *Spacetime and Geometry: An Introduction to General Relativity*. Cambridge University Press.

Chalmers, David J. (forthcoming), "Finding space in a nonspatial world." In *Philosophy Beyond Spacetime* (Baptiste Le Bihan and Nick Huggett, eds.), Oxford University Press.

Crowther, Karen (2018), "Inter-theory relations in quantum gravity: correspondence, reduction, and emergence." *Studies in History and Philosophy of Modern Physics*, 63, 74–85.

Donnelly, M. (2011), "Using mereological principles to support metaphysics." *The Philosophical Quarterly*, 61, 225–246.

Huggett, Nick and Christian Wüthrich (2013), "Emergent spacetime and empirical (in)coherence." *Studies in History and Philosophy of Modern Physics*, 44, 276–285.

Lam, Vincent and Christian Wüthrich (2018), "Spacetime is as spacetime does." *Studies in History and Philosophy of Modern Physics*, 64, 39–51.

Lam, Vincent and Christian Wüthrich (2020), "Spacetime functionalism from a realist perspective." *Synthese*, https://doi.org/10.1007/s11229-020-02642-y.

Le Bihan, Baptiste (2018a), "Priority monism beyond spacetime." *Metaphysica*, 19, 95–111.

Le Bihan, Baptiste (2018b), "Space emergence in contemporary physics: Why we do not need fundamentality, layers of reality and emergence." *Disputatio*, 10, 71–95.

Le Bihan, Baptiste and Niels Linnemann (2019), "Have we lost spacetime on the way? narrowing the gap between general relativity and quantum gravity." *Studies in the History and Philosophy of Modern Physics*, 65, 112–121.

Ney, Alyssa (2020), "Finding the world in the wave function: some strategies for solving the macro-object problem." *Synthese*, 197, 4227–4249.

Paul, L. A. (2012), "Building the world from its fundamental constituents." *Philosophical Studies*, 158, 221–256.

Saucedo, Raul (2011), "Parthood and location." In *Oxford Studies in Metaphysics Vol. 6* (Karen Bennett and Dean Zimmerman, eds.), 223–284, Oxford University Press.

Sider, Theodore (2007), "Parthood." *Philosophical Review*, 116, 51–91.

Wallace, David (2020), "Fundamental and emergent geometry in newtonian physics." *British Journal for the Philosophy of Science*, 71, 1–32.

Wüthrich, Christian (2017), "Raiders of the lost spacetime." In *Towards a Theory of Spacetime Theories* (D. Lehmkuhl, G. Schiemann, and E. Scholz, eds.), 297–335, Birkhäuser, Basal.

Wüthrich, Christian (2019), "The emergence of space and time." In *Routledge Handbook of Emergence* (Sophie Gibb, Robin Finlay Hendry, and Tom Lancaster, eds.), Routledge.

Yates, David (2021), "Thinking about spacetime." In *Philosophy Beyond Spacetime* (Christian Wüthrich, Baptiste Le Bihan, and Nick Huggett, eds.), Oxford University Press, Oxford.