Finding Space in a Nonspatial World*

David J. Chalmers

1 The two images

A familiar metaphor from Wilfrid Sellars says that there are two images that we confront and that we have to reconcile when doing philosophy. The manifest image is how things appear to be in everyday perception and thought. The scientific image is how things are according to science.

Numerous debates in philosophy try to reconcile these images. How do you reconcile the manifest picture of consciousness as it presents itself with what science is telling us about the brain and the mind? How do you reconcile the manifest image and the scientific image for free will, for meaning, for value, for color, or for time? Here I will confront that question for space.

How does space appear in the manifest image—the image presented to us relatively pretheoretically in ordinary perception and thought? It appears that space is roughly Euclidean. Perception presents an absolute three-dimensional space evolving in time. Space arguably seems to be fundamental, a kind of fundamental container of all the things presented to us in the concrete world.

Space as revealed in the scientific image, on the other hand, is non-Euclidean. There is a four-dimensional spacetime without absolute dimensions of space or time. On some current theories, spacetime may be non-fundamental.

What is the relationship between space in the manifest image and space in the scientific image? The relevant science here is physics. The idea I'll focus on is that in moving from the manifest

[0]Forthcoming in (B. Le Baptiste, N. Huggett, and C. Wüthrich, eds.) *Philosophy Beyond Spacetime* (Oxford University Press). Sections 1-6 are an edited version of my third Jean Nicod lecture at the Ecole Normale Superieure in Paris on June 25, 2015. (Video at http://savoirs.ens.fr/expose.php?id=2167.) The Nicod lectures will eventually appear as a book from MIT Press. This material was first presented in talks to the ANU Research School of Physics and a Rochester conference on the metaphysics of physics in 2013, and later in talks at Arizona, Brown, Connecticut, Johns Hopkins, Oxford, and Rutgers. Thanks to audiences on all these occasions for their feedback. For comments, thanks also to Baptiste Le Bihan, Trevor Teitel, and David Yates.

image of space to the scientific image, we need to move from spatial primitivism, where we have a primitive conception of space, to spatial functionalism, where space is picked out as whatever plays the space role.

I argued for a functionalist view of space and time in my 2012 book *Constructing the World*. Here I'll motivate and develop spatial functionalism in more depth, using considerations from physics and also analogies with the metaphysics of color and with functionalism in the philosophy of mind. In the final two sections I'll connect these ideas to recent work in the philosophy of physics on spacetime functionalism and on emergent spacetime.

2 Primitivism and functionalism about color

I'll start with an analogy with the case of color. In the case of color, vision appears to present a world of colored objects. Out there in the environment now there are these deeply colored red chairs and there's a blue one and there's a white screen and so on. On the face of things, colors seem to be relatively primitive qualities of objects. The chair's surface is somehow imbued with redness. This chair's surface is somehow imbued with blueness. We seem to have some kind of fairly direct acquaintance with colors on the surface of these objects.

A metaphor that I like in making sense of the manifest image to suppose that in the Garden of Eden everything was exactly as presented in the manifest image. Then Eden is a label for a world where everything is exactly as it seems to be.

How were colors in the Garden of Eden? It's arguable that the phenomenology of color experience presents us with a world of primitive color qualities on the surface of objects. In Eden, colors are just like this. Objects were imbued with primitive colors on their surfaces. We can call these colors Edenic redness, Edenic greenness, and so on. We stood in a direct acquaintance relationship to Edenic colors on the surface of objects.

Of course, when there's Eden there's a fall. Scientists ate from the tree of knowledge and discovered that we dont seem to live in Eden. People examined objects out there in the external world and they dont seem to have these primitive color qualities. All they have is some really complicated physical properties that reflect light in certain ways. Rather than having direct acquaintance with those things, there is a long and complex chain leading from these physical properties to the brain, involving light transmitted to the retina and electrical signals transmitted up the optic nerve.

Many people reacted to this by saying that colors dont really exist. The fact that people have that reaction is some evidence for the role of this primitivist conception of color in our ordinary

conception of the world. Galileo, for example, reacted by saying objects in the world arent really colored at all. Colors exist only in the mind and not out there in reality.

If you take this line, you have to say that apples arent really red, and grass isnt really green. This is a kind of eliminativism about colors. In effect, colors are eliminated from the scientific image. A number of people have reacted this way, but over time it's been the less popular reaction. Theres a more common reaction, which we can call compatibilism. Compatibilism says that even in light of the facts of color science and this fall from Eden, apples are still red, grass is still green. Even though we dont live in Eden, objects still have colors.

This is done in a variety of ways. Sometimes people identify colors with dispositions. Perhaps most common is to identify colors with certain complex physical properties such as surface reflectances and other physical properties such as radiation properties that affect light in appropriate ways. In effect, whats happening is that redness is now identified as that physical quality that plays a certain causal role. In particular, it is that physical quality that is causally responsible for our experiences of redness.

One can think of this metaphysical picture as a sort of imperfect realism. Imperfect realism is a picture that is very common in philosophy. It roughly says there are Xs, they aren't Xs exactly as they seem to be pre-theoretically, but they are still something that deserves well enough to be called Xs. In the case of colors, there are no perfect colors. There are no colors exactly as presented in experience, primitive Edenic qualities on the surface of objects. But there are still imperfect colors, properties that play the color role, for example in bringing about color experience.

One way to see whats going on here is there has been a move from color primitivism, on which colors are regarded as primitive qualities, to color functionalism, on which colors are picked out as what play a certain causal role. In particular, colors are whatever is normally causally responsible for our color experience. As with other familiar varieties of functionalism about which we'll talk more soon, this view doesn't actually require that one be functionalist about the metaphysics of colors. It doesn't require that colors be functional properties. It's consistent with that view, but it's also consistent with the view that colors are physical properties that play that role.

This view tends to suggest functionalism about color concepts. Roughly we're conceiving of colors as whatever causes color experience. In principle, one could run this without functionalism about color concepts by just appealing to a metasemantic causal theory of reference where we pick out whatever plays this role without our conceiving of colors that way. But I think reflection on how we go about picking out the colors in virtue of their serving as the basis of certain color appearances suggests that the strongest story involves functionalism about the relevant concepts.

When you've have these Edenic colors and non-Edenic colors around, it's arguable there are also two kinds of concepts in the vicinity. There are at least two kinds of concept that somebody could have and arguably two kinds of concepts that we possess and that regulate our thought and talk about color. There's an Edenic concept, which we can call capital-R Red, for which primitivism and eliminativism are correct. We've got a concept of Edenic redness. It's a primitive concept, but nothing in the world turns out to satisfy it. We also have a non-Edenic concept, small-r red, for functionalism and realism are correct. For small-r red, we just pick our colors as whatever causes the relevant experiences. It's a functional, non-primitive conception, but this is one for which realism is correct, because there are things that really play that role. Partly because of this, the functionalist concept is most useful in making sense of the world.

3 Primitivism and functionalism about space

The move from primitivism to structuralism or functionalism is a common picture throughout various areas of science when we try to reconcile a categorical manifest image with a more structural scientific image.

In my view, the same story applies to space. How was space in Eden? How is it presented to us pretheoretically? The kind of space that we have an intuitive grasp on in our experience seems to have a certain categorical qualitative nature that serves as the container of everything we're presented with. In Eden, there were these perfect spatial properties—Euclidean distances, squares, and so on– all embedded in a primitive space and time. Then we ate from the tree of knowledge again. This time the tree of knowledge came in a number of stages.

The first really important step in the tree of knowledge for space was taken by Newton. Newtonian physics did much less to undermine the Edenic picture of space than corresponding steps around the same time did in the case of color. Newton suggests a fundamental Euclidean space, which at least seems consistent with the Edenic model of space.

More radical developments came later. Around the start of the last century, Einstein suggested a non-Euclidean space in which spatial properties are relative. Nothing is absolutely square. Things are just square relative to a reference frame. That's already a spatial fall from Eden. Using the intuitive concept of Edenic space that one gets directly from one's experience, I think it's arguable you can't really make sense of Edenic space being something that holds relatively rather than absolutely. In this move to relativize space, there's a fall from Eden going on.

It's arguable that the fall gets worse when it comes to quantum mechanics. The interpretation

of quantum mechanics is controversial and contested, but it's pretty widely believed that what's fundamental in quantum mechanics are not things like the positions of particles, but wave functions. Wave functions are generally held to inhabit a high dimensional configuration space. This is not a three or four-dimensional space but roughly a 3n-dimensional space where n is the number of particles. That space that corresponds to the space of possible configurations of those particles, with every configuration getting an amplitude. Instead of having fundamental 3-space, you have a fundamental 3n-space or infinite-dimensional space if the world is big enough. It starts to look like three-dimensional space isn't fundamental but arises derivatively from configuration space. At least so it goes on many popular interpretations of quantum mechanics.

Things get even worse in string theory and other theories of quantum gravity, such as loop quantum gravity and causal set theory. These theories try to integrate quantum mechanics and relativity. Of course that is much more speculative physics now. These theories all postulate various models on which space isn't fundamental at all. It's now common to suggest that space and possibly time are emergent in that they pop out at a derivative level of the theory. Nathan Seiberg (2007) put it this way:

> "Spacetime is an emergent concept. The fundamental formulation of the theory will not have spacetime and it will emergence as an approximate classical concept which is valid only macroscopically" (Seiberg 2005)

Getting even more speculative, there's even digital physics, involving what people sometimes call the It from Bit hypothesis. That slogan is sometimes used for various different views. I'm concerned with views on which it turns out that underlying the apparent structure of physics including spatial structure is some computational level, perhaps something like a cellular automaton governed by an algorithm in a more fundamental space, which isn't our familiar 3-space. That the spatial structure of physics emerges from that. Stephen Wolfram, Edward Fredkin and others have famously speculated along these lines.

Of course quantum gravity theories and especially digital physics are admittedly speculative physics and the interpretation of quantum mechanics is controversial. But the kinds of interpretations of quantum mechanics that go in this direction are at least reasonably popular, and of course relativity is solid.

One reaction to all this is to say that space is an illusion if things are as presented that way in physics. We have the experience that things are located in 3-space when in fact, they aren't. That's spatial eliminativism, analogous to color eliminativism. It seems to require that all spatial

experience is illusory and that all ordinary spatial claims are false. One can find a version of this claim in David Albert's 1996 paper on elementary quantum mechanics, though these days he accepts something closer to compatibilism (see Albert 1996 and 2015). I take it that is a much more popular reaction in this case, as it is in the case of color. The compatibilist says that even if physics is this way, there is still space. It's just not exactly as we took it to be. It's going to be relative rather than absolute, and it's going to be a derivative rather than fundamental property. As Seiberg puts it, space-time is an emergent concept. This is basically capturing the idea that space is non-fundamental and derivative.

The question that arises, if space is not fundamental, how do we pick out the complex properties in this world that are spatial properties? I think the answer is very much analogous to the answer in the case of color. We pick out spatial properties as those properties that play the right role in the structure of the physical dynamics and in bringing about spatial experience. This is going to be a kind of imperfect realism. There are no perfect spatial properties, primitive space exactly as presented in experience. But there are still imperfect spatial properties, complex properties that play the space role.

As in the case of color, I think we've moved from primitivism to functionalism. We started with a kind of intuitive spatial primitivism where space involves these primitive qualities that we're acquainted with, and everything is spread out in that primitive space. We've ended with a spatial functionalism. To use a familiar functionalist slogan: space is as space does. Or better: space is whatever plays the space role. As with color, spatial properties need not themselves be functional properties (as the first slogan suggests), but they are picked out as what plays the functional role. This is analogous to various other familiar sorts of functionalism where we pick out Xs as whatever play the X role. Colors are whatever plays the color role. It's very common to say for many theoretical terms in physics like mass, take mass to be whatever plays the mass role in the theory. For example, whatever properly resists acceleration and yields gravitational attraction in a certain way, that counts as mass.

On spatial functionalism, space is whatever plays the space role with respect to physical laws and/or experience. These are two distinct aspects to the role which I'll return to later on. Some varieties of spatial functionalism focus on the role in physical laws, and some focus on the connection to experience.

As in the case of color, I think there are at least two concepts of space. There's capital-S Space, Edenic space, for which primitivism and eliminativism are correct. We have this primitive concept of Edenic space, space as it was in the Garden of Eden. Relativity, quantum mechanics and so

on at least tend to strongly suggest that we don't live in that world. There is no absolute Edenic space. We've also got a non-Edenic concept, small-s space, for which functionalism and realism are correct. This concept picks out space as whatever plays the relevant role. Even in our post-fall world, something plays this role. As in the case of color, I think both concepts are in some sense perfectly coherent, but the functionalist concept is the most useful in making sense of the world we are confronted with. Again, this involves the move from intuitively having a categorical grip on the nature of space and spatiality to merely having a structural grasp of imperfect space, grasping it in virtue of the role it plays.

4 Spatial functionalism and explanatory gaps

To some extent, spatial functionalism can be motivated even in physical frameworks where space is fundamental. Some bits of it can be motivated even in Newtonian and relativistic physics. One motivation comes from spatial Twin Earth cases. I discuss a number of these in "Three Puzzles about Spatial Experience". There's Brad Thompson's (2010) case of Doubled Earth, suggesting that we don't have a categorical grasp of absolute size properties. Even in Newtonian physics, there's no absolute grasp of a length of one meter. That length is roughly picked out as whatever brings about certain kinds of experience. There's also the relativistic case of Lorentz Earth, suggesting that we don't have a categorical grasp of ordinary shape properties. Shape and relative size properties are also picked out in terms of the role they play.

Spatial functionalism becomes especially relevant and useful for making sense of frameworks where space is not fundamental, including some interpretations of quantum mechanics, quantum gravity, computational physics and so on. In quantum mechanics, wave-function fundamentalism (e.g. Albert 2015 and Ney 2015) says all that's fundamental in quantum mechanics is the wave function defined over a configuration space evolving according to the Schrödinger equation. A number of popular interpretations of quantum mechanics have this form. Everett-style interpretations typically go this way. Collapse interpretations introduce extra dynamics with of wave-function collapse, but still tend to say the wave function is all that's fundamental. The wave function is standardly taken to inhabit this high-dimensional configuration space. The question then arises, how do we recover familiar 3- or 4-space from high-dimensional configuration space?

The worry involves a kind of explanatory gap. How do you get from this massively high dimensional wave function to the familiar three-dimensional space or four-dimensional space-time that were confronted with? It's not easy to find space of that kind in the wave function. On closer

examination, I think spatial primitivism leaves an explanatory gap but spatial functionalism does not. This is an instance of a more general phenomenon involving functionalism and explanation.

In general, functionalism about a domain helps us to close explanatory gaps involving that domain. Familiarly, in the case of the mind–body problem, if you're a functionalist about the mental, if you think that what it is to be in pain, for example, is to be in a state that plays the relevant causal role and that's what it means to be in pain. That helps you a lot in closing the gap from physical to mental. You can see roughly how in a physical world something would play that role. So if this analysis was true, you could see how a physical system could be in pain.

I don't endorse functionalism about pain myself, but I like the general structure. This is a version of the Canberra Plan that is familiar from the work of people like Frank Jackson and David Lewis. To close an explanatory gap, say from the physical to the mental, or from the fundamental level to something non-fundamental, what you need is something like an a priori entailment. That way you can see why given that things were like this at the bottom level, they just have to be like this at the higher level. Somehow the relevant non-fundamental facts can be inferred from the fundamental facts. I don't want to put too much stress on the role of the a priori here but the rough idea is that we want some fairly transparent epistemic connection from the fundamental facts so the non-fundamental facts in order to close the explanatory gap.

Functionalism about the mental helps us see how that could be true. The picture from Lewis, Jackson, and others goes something like this.

1. A priori (pain is what plays the pain role).
2. A priori (if P, then C-fiber-firing plays the pain role)
3. A priori (if P, then x's C-fibers are firing).

---

4. A priori (if P, then x is in pain).

First, you hold it's a priori that pain is what plays the pain role. That's a functionalist analysis of pain. Then, let's say P represents relevant physical facts about the physical configuration of our world. The thought is that given these physical facts one can deduce functional properties of the brain, and in particular can deduce C-fiber firing is playing that role. It will also be deducible that in a certain individual, their C-fibers are firing. Putting all those together, from P it will be deducible that this individual is in pain.

This potentially closes the explanatory gap from physics to pain. It's controversial and I would reject the first step, but it's a nice general explanatory structure for locating pain in a physical world.

Likewise, functionalism about the spatial, if true, helps to locate space in a world where it is not fundamental. It closes the gap from the non-spatial to the spatial. For example, from a world that's fundamentally characterized at the wave function level independent of space. If we have a concept of space for which it's a priori that space is what plays the space role, then I think you can plausibly make the case for a similar structure.

1. A priori (space is what plays the space role).
2. A priori (if P, then Q plays the space role)
3. A priori (if P, then x has Q)

---

4. A priori (if P, then x is in space).

Here P is the relevant fundamental physical configuration characterized independently of space, perhaps in quantum mechanical configuration space. Then if P is the case, one can deduce that some complex quantum mechanical state Q plays the space role. One can also deduce that some particular object X is in this state Q. Putting all those together, it's a priori that if P is the case then I'm located in space, thereby closing the explanatory gap from the non-spatial to the spatial.

An austere version of this spatial functionalism project includes just fundamental physics (P) in the base. A less austere version allows the base to also include facts about experience. This can make it easier to locate space. It allows us to characterize the space role as not just a physical role but also a role in causing experiences. On this way of doing things, instead of closing the gap between physics and space we close the gap between physics plus experience and space. That will of course leave open the residual mind–body question of reducing experience, but it's fairly common to kick that problem down the road. I will look at both versions in what follows.

In the mind–body case all this gives rise to a complex structure of philosophical positions. We can put the issue by saying that these five views are inconsistent, so one of them has to go. Here phenomenal properties are properties characterizing what it is like to be conscious beings, and phenomenal concepts are our introspective concepts of those properties.

1. Physicalism (only physics is fundamental)

2. Phenomenal realism (things have phenomenal properties)

3. Phenomenal primitivism (phenomenal concepts are primitive).

4. Scrutability (no epistemic gap between fundamental and nonfundamental).

5. Epistemic gap (epistemic gap between physics and primitive phenomenal properties).

The first item in the quintet is physicalism: Only physics is fundamental. The second is phenomenal realism: some things have phenomenal properties. That is, there are some beings (like us) for whom there is something it is like to be them. that is the qualitative properties of experience. Phenomenal primitivism says that phenomenal concepts are primitive, meaning that they are not analyzable in nonphenomenal terms—and certainly that they are not functional concepts. Scrutability is roughly the claim that you can't have these deep epistemic gaps between the fundamental and the non-fundamental domain. The version of it I like is that the non-fundamental facts should be a priori entailed by the fundamental facts. Finally, there's the claim that there's an epistemic gap between physics and primitive phenomenal properties. If phenomenal properties can't be analyzed in nonphenomenal terms, there's an epistemic gap (and no a priori entailment) between the physical and the phenomenal.

Different people reject different claims here. Analytic functionalists reject phenomenal primitivism by saying that phenomenal concepts are functional concepts. Eliminativists reject phenomenal realism, saying that no one has phenomenal properties. So-called Type-B materialists reject scrutability, saying that the phenomenal can be ontologically grounded in the physical even if there is an epistemic gap. Some panpsychist physicalists reject the epistemic gap thesis by saying that physics builds in consciousness. I'm inclined to reject physicalism.

In the spatial case, we can lay out an analogous inconsistent quintet.

1. Wave-function fundamentalism (only the wavefunction is fundamental).

2. Spatial realism (things have spatial properties)

3. Spatial primitivism (spatial concepts are primitive).

4. Scrutability (no epistemic gap between fundamental and nonfundamental).

5. Epistemic gap: (epistemic gap between wavefunction and primitive spatial properties).

The analog of physicalism here is wave-function fundamentalism. Then there's spatial realism, which says that things have spatial properties. Spatial primitivism says that spatial concepts are primitive—that is, that they are not analyzable in nonspatial terms, and certainly not in functional terms. Scrutability is the same principle as before, saying there is no epistemic gap between the fundamental and the non-fundamental. Finally, there is an epistemic gap between the wave function and spatial properties, at least as construed by the spatial primitivist.

Again, something has to go. The analog of my view in the case of the mind–body case would be to give up on wave-function fundamentalism and say, we need something else besides the wave function. Maybe we need some primitive particles in 3-space, maybe we need primitive space some other way, or something else.

In the mind–body case, my view is that phenomenal realism is more or less non-negotiable. When I say I am phenomenally conscious, in the ordinary sense of phenomenal consciousness, this is true. It also think it's extremely plausible that the ordinary phenomenal concepts used to express this claim are primitive concepts.

In the case of space, I don't think these things are non-negotiable. It's not a non-negotiable fact that things are in space in a primitive sense. We do have a primitive concept of Edenic space, analogous to our concept of phenomenal consciousness, but I think it's more palatable to deny that Edenic space exists than to deny that phenomenal consciousness exists. So I think the combination of spatial realism and spatial primitivism is something one can put much more pressure on than in the mind–body case. A spatial primitivist who takes our concept of space to be a concept of Edenic space should deny spatial realism. Likewise a spatial functionalism can deny spatal primitivism and accept spatial realism. That's the combination I prefer.

An epistemological motivation here is also worth mentioning. Spatial functionalism has the potential to save the appearances in a certain sense when it comes to space. One sense of saving the appearances is saving exactly how things appear to be. In that case, to save the color appearances, what we really have to do would be to vindicate the Edenic picture of colors in the world. That we don't get to do, given color science. But we do vindicate the picture whereby we have these color experiences. Arguably, our having these experiences is the most fundamental data. If you regard those as the appearances, we can then construe colors as the things that bring about these experiences. Arguably, this saves the most fundamental data which are the data of our experience. Likewise, for spatial functionalism, as long as we're in a world where something is playing the right space role, it's arguably going to be the case that, that will be a world that supports the same kind of spatial experiences.

Even though the external world may not be the Edenic world of our original conception, as long as things there function in the right way and bring about our experiences in the right way, then it will save the appearances. One thought here is that even if there's no Edenic space out there in the world, there'll be something out there which is playing least a very broadly functionally isomorphic role in a functionally isomorphic world. A functionally isomorphic world, even if it's not Edenic, will at least have functional isomorphs of ordinary perceivers, like you and me as we standardly take ourselves to be. If you can motivate the case that in a world with such functional isomorphs of ordinary perceivers there will also be experiential isomorphs, people having corresponding experiences, then you've saved the experiences.

Now, admittedly, that last step is a strong step and interacts with controversial issues about the mind–body problem. But at least if you're inclined to think that what really plays a central role as a physical basis of experience is something about our functional dynamics, rather than say, facts about our absolute spatial layout, then you'll be inclined to go this way. An opponent might try to make the case that you need Edenic space to have the kind of experiences that we have—for example, how could we have experiences as of Edenic space in a world without Edenic space? That would be an interesting place to put some pressure. Still, I think this is a good prima facie case for some epistemological virtues of the view.

An even more austere picture is suggested by developing David Lewis's austere picture of Humean supervenience. On Lewis's picture, fundamentally there are just primitively distinct natural properties throughout space-time. Lewis builds up the rest of the world from there by Ramsifying properties like mass and charge as what plays the mass role and what plays the charge role. He builds in space and time as fundmental, though. He never takes the further step of Ramsifying space and time via spatiotemporal functionalism, but it's a natural thing to do. That leads to an even more austere Humean supervenience that doesn't mention spacetime at the fundamental level. On this picture, a fundamental description just has logic, mathematics, and naturalness as its basic ideology. If one likes austere ontologies, I think this is an attractive picture of the world.

5 What is the space role?

Now the question arises, what kind of functional analysis are we going to give of space? That is to say: when we analyze space as whatever plays the space role, what counts as the space role?

There are two basic options for functionally analyzing space (with many variants of each and also approaches that combine both). The first is a phenomenal analysis, where we pick out space

in virtue of its role in causing relevant spatial experiences, very much analogous to the case of color. I've discussed that sort of analysis in "Three Puzzles about Spatial Experience" and other work where it is motivated by the analogy with color and Twin Earth cases and the like. Here I want to focus on another option.

The second option is a nonphenomenal analysis where space is picked out whatever plays the physical space role in the dynamics of the physical world. That line that doesn't have so much of a direct analog in the case of color. It looks like colors don't play such an interesting dynamical role independently of us in the physical world. Their role is mostly interesting only relative to brains organized a certain way. But space is interesting and plays an important role independently of us, it seems. So it's possible to spell out key roles of space quite independently of conscious experience, in terms of what we might think of as the physical space role.

Why should we be concerned with that analysis? One reason is that this analysis might be more useful for materialists about the mind who want to reduce space to something more fundamental without presupposing any phenomenal or experiential notions. If we presuppose those notions, then that will leave your theory of space hostage to your theory of the mind, and at some point you may need to worry about the mind–body problem. More generally, it may be theoretically cleaner to make sense of the emergence of space in terms of physics alone. Many physicists and philosophers are naturally uncomfortable to giving minds a central role in making sense of space. Of course space itself need not be mental on the phenomenal version of spatial functionalism. It's just that minds play a role in fixing reference to what counts as spatial, the same way that water appearances help fix reference to water. But even that may be more of a role than some people want to give.

What are the options for nonphenomenal functional analysis of space? A natural approach is to use the Ramsey-Carnap-Lewis method for analyzing theoretical terms. We start with a theory of space, including our relatively pre-theoretical folk theory and/or our relatively sophisticated scientific theories. We list the central principles concerning space in the theory, or the platitudes as they're sometimes called. We regiment these principles concerning space into a space role. Then we identify space with whatever fills the role. For mass, the mass role might be to resist acceleration and undergo gravitational attraction in certain ways. Then mass is whatever plays that role. That is, it is whatever resists acceleration and is subject to gravitation and so on.

What are the relevant roles in the case of space? I think there are a few roles. First, there are what we might call structural roles characterizing the broad abstract structure of space. Pre-theoretically at least, space seems to be three-dimensional with one corresponding dimension of

time. There seem to be locations in space with distances between them. It seems that objects take up space, they're located at various locations in space at least at times, and they thereby have properties like shapes, sizes, and relative distances. One can go on to spell out broad structural characterizations of space in these terms.

Second, there are phenomenal roles, which I've talked about already. There are various things we might think of as the phenomenal platitudes: square things normally look square, small things normally look small, distant things normally look distant, and so on. There's some room for illusions where we get things wrong but the thought is, by and large, things are roughly the way they seem to be. This is basically a role connecting spatial properties of objects with certain kinds of spatial experiences which we can then perhaps use to characterize objects as what play that role with respect to experience.

Third, there are what we might call folk physics roles. There's a sort of pre-theoretical physics that's familiar from being an agent in the world before you do a lot of science. Psychologists have studied folk physics quite a lot. According to folk physics our bodies are located in space, there are many familiar objects like bodies, rocks, tables, trees with familiar shapes, sizes, relative locations. Their spatial properties evolve according to certain familiar principles. You throw them up in the air and it comes down. If something is moving in a certain direction, it keeps moving in that direction, although maybe it gradually slows down according to folk physics.

Fourth, there are scientific roles. Science tells us that space plays such and such roles. It roughly obeys principes of Newtonian mechanics and Newtonian gravitation in ordinary contexts. More precisely it obeys principles of relativity theory, quantum mechanics, and so on. There's a question of how much of this science one wants to and one is entitled to bring in to one's analysis of space here. The basic worry is that the fact that space plays these roles is only discovered by science and so seems to be a posteriori. So if one is analyzing our ordinary pre-science concept of space, these roles don't seem to be part of the concept.

At the same time, we have post-science concepts of space, such as the concept of Newtonian space, for which the roles of Newtonian mechanics are absolutely central. It's arguably a priori of Newtonian space that it plays the Newtonian roles. In a similar way there's a relativistic concept of spacetime, and it's arguably a priori of relativistic spacetime that it plays the relativistic roles. For explaining why there is something approximately like Newtonian space in a relativistic context, you want to apply the Newtonian concept. You show that in relativity, something approximately plays the Newtonian space role under certain conditions. To explain why there is something like relativistic spacetime in string theory, you want to apply the relativistic concept. You show that

in string theory there is something that at least approximately plays the relativistic spacetime role under certain conditions.

In general, to pick out space in some later theory (say string theory), we will typically want at least a relatively a priori analysis that does not presuppose that theory. But it might presuppose some earlier theory such as relativity. You're not exactly vindicating the manifest image in terms of the scientific image, but you're vindicating at least a relatively manifest image (say the image of Newtonian space) in light of a relatively scientific image (say the image of relativistic spacetime).

So we have different concepts of space with different analyses. I am especially interested in our ordinary concepts of space, so I'm interested in relatively pretheoretical roles. Now, you might say, all these functional analyses of space are bad analyses of the pretheoretical concept, just as the functional analysis of pain is a bad analysis. Many people including me think the functional analysis of pain is bad because you can conceive of all those roles being played without pain, and maybe of pain without those roles being played. Likewise, you might say, "I can conceive of a space which plays none of these roles and maybe I conceive of these roles being played without space".

I'm sympathetic with this point to a degree. I think yes, one can conceive of space that plays none of those roles. Perhaps there could be a really weird space that evolves by completely different dynamics, it's not connected to experience in these ways and so on. In these intuitions I think I'm basically invoking the primitive or Edenic conception of space, which I think is not a functionally analyzable concept. Maybe you think you can conceive of a non-primitive space that's not the Edenic conception and still doesn't play these dynamic or folk physics or scientific roles. Arguing over this will involve some litigation of cases. My guess is that one will be invoking either the primitive or the phenomenal conception at that point

Our pretheoretical conception of space involves these primitive elements on my view, but it also involves functional elements. These might help to ground a functional concept of space that's both pretheoretical and nonphenomenal. Certain dynamic roles from folk physics may be especially crucial, especially certain roles in motion and in interaction. A central principle of folk physics says that objects move continuously through space. Another says that objects interact with nearby objects in space. To a first approximation, there's no discontinuous motion and there's no action at a distance. These are commonly taken at least to be principles of folk physics. Maybe they are violated here and there in scientific physics. But even in scientific physics, action at a distance is heavily constrained and there are many forms of action and motion that are local. That is, there seems to be less action at a distance and less motion at a distance, so that distance still has

some connection to the possibility of interaction and the character of motion.

These principles are a central part of the folk conception of space. If we apply the Ramsey method, what we have to do is turn the principles around and identify distance with what plays the distance role. Instead of saying there's no action at a distance, we turn this around and we say, "Distance is what there's no action at." This is a nice slogan that goes back at least to Brian Cantwell Smith in his 1996 book *On the Origin of Objects* where he gives something like this functional analysis of distance in a very different context.

In a similar way, we can use the principles about continuous motion to say "Distance is what there's no motion at", or equivalently, that location is something that changes continuously. Or if you want to do this in light of the science of distance, we can say that distance is what there's less action at and less motion at.

There's a very nice statement of this kind of distance functionalism in Gregg Rosenberg's 2004 book *A Place for Consciousness*, which advocated panpsychism about the mind–body problem as well as a revisionary view of causation in the physical world. Along the way Rosenberg also gave what he called a kind of causal analysis of space-time. He put this point wonderfully by saying "There's a causality condition on locality, not a locality condition on causality." A locality condition on causality says, it turns out that causation only works in a local way. A causality condition on locality says, what it is to be close is to have the possibility of certain sorts of causal interaction. These claims may have to be weakened a little to accomodate a certain amount of nonlocal causation, but this is consiatent with an underlying conceptual connection between the ideas.

We can think of this as the interactive conception of space. It's perhaps just one strand in our ordinary conception of space, but it can play a key role in trying to make sense of one aspect of the manifest image of space. This is a really central aspect in light of the scientific image of space.

It's arguable that this interactive notion of space and its role in motion and interaction and so on corresponds to a more generic notion of space. We've got this notion of a space in a general sense, which transcends physics. People talk about social spaces and virtual spaces and administrative spaces and teaching spaces and cyberspace. There are spaces for everything these days. Something in this generic notion of space seems to be tied to thinking of space as a generic realm in which a sort of motion takes place and in which interaction takes place. There are locations in these spaces and locations are close in a space when things move and interact between those locations.

Perhaps at least when combined with some further constraints, this pretheoretical notion of interactive space can at least help to ground a reduction of physical space. Say for example, if in

the quantum mechanical wave function one can find a three or four-dimensional metric of relations so that motion is relatively continuous and interaction is relatively constrained by that particular metric, that could at least help to ground a reduction of 3-space or 4-space in the world of quantum mechanics.

One may well need further constraints. Just what constraints on a functional analysis of space does one need to ground a reduction in quantum mechanics and string theory and computational physics? The details I take to be very much an open question.

One route that may help is a phenomenal analysis of space. Here we can pick out spatial properties in a quantum-mechanical or string-theoretic world as what causes spatial experiences in that world. As long as the wave function or the string-theoretic structure does being about our experiences, then there ought to be something systematic to say about what quantum-mechanical or string-theoretic properties cause these experiences. Of course there's the worry about giving the mind a role in fixing reference to physics, but I think this happens a lot. One still needs to worry about how spatial experiences can arise in a fundamentally nonspatial world. That problem will become part of the mind–body problem, perhaps making the problem a little harder. For a nonreductionist about the phenomenal, the solution will come down to finding laws that connect certain nonspatial properties in brains and the like to spatial experiences. Anyway, there's at least a program for finding the causes of our spatial experiences in a fundamentally nonspatial world, and I think it's reasonable to be optimistic about finding things that can play that role.

If you want a completely nonphenomenal functional analysis of space, things are trickier. It's perhaps easier for theoretical concepts of space where theories give a fairly precise role. As before one can try to find something that plays roughly the Newtonian role or the relativistic role for space in a quantum mechanical world or a string-theoretic world without space at the fundamental level. This is highly nontrivial but can be pulled off if there's the right sort of fit between the theories. In delivering something akin to Newtonian space or relativistic spacetime, it may not quite deliver exactly the original pre-theoretical manifest image, but it still delivers some relatively pre-theoretical image.

The hardest case is grounding a reduction of space with a nonphenomenal analysis of the pretheoretical concept. It may be that the roles in interaction and motion that I've been gesturing towards can be used to deliver something even closer to the manifest image of space, picking out underlying properties in a quantum-mechanical or string-theoretic world as the grounds of motion and interaction. Or perhaps one can bring in other roles of space from pretheoretical folk physics and see what plays them in the relevant physical theories. I take that to be an open question which

is going to involve a lot of detailed work in the philosophy of physics combined with philosophical and/or psychological analyses of our concepts.

6 Virtual reality and Cartesian skepticism

Another interesting case is finding space within virtual reality. In other work I've argued that spatial functionalism applies there as well.[1] In virtual reality systems, there is typically something playing the space role. So if we're inhabiting virtual reality, we can identify space with those digital properties that play the space role.

This has some bearing on Cartesian skepticism. A Cartesian skeptic typically suggests that for all we know we're in something like a Matrix scenario where nothing that we seem to experience is real. But if spatial functionalism is right, then if we're in a Matrix scenario, space is quite real. At least, space is no more illusory than it is in quantum gravity. And once the space around us is real, it's a short step to the objects in it being real too.

On this way of looking at things, if we turn out to be living in the Matrix, that's just another stage in the fall from Eden analogous to the Galilean and Einsteinian falls from Eden. After Galileo, red was identified as a reflective property. After Einstein, squareness was identified as a relative property. After the matrix, squareness was identified as a virtual property. We've re-made our conception of it. We don't need to say things aren't square. We just say that squareness is not exactly what we took it to be.

There's still an intuition that the matrix scenario is an error scenario, one where people get things wrong. I have the intuition. But I think this is best explained by our using Edenic concepts in our intuition. In a matrix scenario, there is no Edenic color or Edenic space. The Edenic content of our experiences is incorrect, and our experiences are not perfectly veridical. But if this makes the Matrix a skeptical scenario, then the same goes for quantum mechanics, relativity, and so on. I've argued that this is too high a standard. By the lower standard where quantum mechanics comes out looking okay, so does the matrix. If so, spatial functionalism has some nontrivial purchase in addressing this sort of external-world skepticism.

This is an instance of a general phenomenon: primitivism about Xs tends to open the door to skepticism about Xs, while functionalism about Xs makes skepticism about Xs harder to get off the ground. That applies to consciousness, to color, to space, and to the physical world in general.

[1] See Chalmers 2017 for virtual reality and Chalmers 2018 for applications to skepticism.

7 Spatial functionalism and spacetime functionalism

In the last few years, there has been a small explosion of work on "spacetime functionalism" in the philosophy of physics. Eleanor Knox (2014, forthcoming) has put forward spacetime functionalism as an interpretation of spacetime in general relativity. Vincent Lam and Christian Wüthrich (2018) and others have applied spacetime functionalism to analyze the emergence of spacetime in theories of quantum gravity with a fundamental nonspatiotemporal base. David Albert (2015, 2019) has developed his earlier work that takes a functionalist approach to space in quantum mechanics. Marco Dees (2015) developed a "causal theory of spacetime" on which all geometric properties are nonfundamental and grounded in causal facts at the fundamental level.[2]

How is spacetime functionalism related to the spatial functionalism that I discuss in this article?[3] Both are versions of spatiotemporal functionalism: the broad family of functionalist views of space, time, and spacetime. There are many different views within this class, however. In this penultimate section I will map out the relationship between my spatial functionalism (especially my version) and spacetime functionalism (especially Knox's version), considering other functionalist views along the way. In the final section I will connect all this to issues about the emergence of spacetime.

First, the similarities. Both spatial functionalism (as I present it) and spacetime functionalism (as Knox presents it) are theories of spatiotemporal *concepts*. I say that the concept of space is the concept of whatever plays the space role. Knox says the concept of spacetime is the concept of whatever plays the spacetime role. So both are forms of *realizer functionalism*: space or spacetime is identified with what plays (or realizes) the space role or the spacetime. This contrasts with *role functionalism* which identifies space or spacetime with the playing of the role. Correspondingly, these are not forms of metaphysical functionalism, on which spatial properties or spacetime properties are identified with functional properties (such as the property of playing a certain role).

[2] Dees died tragically in a rock-climbing accident in July 2018. His ideas on the causal approach to spacetime remain unpublished, but his 2015 Rutgers Ph.D. dissertation with a chapter on the topic is available online.

[3] In my own work, I first discussed space and time in a broadly functionalist spirit in "The Matrix as Metaphysics" (2003) and "Perception and the Fall from Eden" (2006). I first talked of "spatial functionalism" and "spatiotemporal functionalism" in *Constructing the World* (2012). These ideas were developed in four Jean Nicod lectures in 2015, which roughly correspond to "The Virtual and The Real" (2017), "Three Puzzles about Spatial Experience (2019), "Structuralism as a Response to Skepticism" (2018), and this article. Of course one can find functionalist ideas about space in a good deal of earlier work. In *Constructing the World* I argue that there are elements of spatial functionalism in Berkeley and Hume. I was also greatly influenced by the work of Gregg Rosenberg (1997/2004) and Brad Thompson (2003, 2010).

Rather, they are forms of conceptual functionalism, where the concepts of space and spacetime are analyzed as the concept of whatever plays the role, so that space and spacetime themselves can be metaphysically identified with what plays the role.

Second, the differences. The most obvious difference is that spatial functionalism is a theory of space while spacetime functionalism is a theory of spacetime. One might think this is just a matter of considering time as well as space, in which case there would be little substantial difference, especially as I discuss "spatiotemporal functionalism" in *Constructing the World* and also advocate functionalism about time. However, there is a more significant underlying difference. Spacetime as understood here is an essentially theoretical concept, one that emerges especially from the general theory of relativity. This concept can be applied to other theories, including Newtonian theories and theories of quantum gravity, but the concept itself did not exist (at least in the relevant sense) prior to relativity. Space is not an essentially theoretical concept. It can be used theoretically, of course. But there is also an ordinary concept of space used by ordinary people before and after the advent of theories of space.

The biggest difference between spacetime functionalism and my spatial functionalism is that the former is primarily an account of the theoretical concept of spacetime, while the latter is primarily an account of the ordinary concept of space. To a first approximation, spacetime functionalism is concerned with the status of spacetime as understood by physicists. I am concerned with the status of space as understood by ordinary people—though perhaps in light of developments in physics. That is the point of starting with Sellars' two images. The project is to see how ordinary space can be located within our scientific conception of the world.

This difference mirrors a familiar distinction in the philosophy of mind. This distinction is sometimes cast as the distinction between *analytic functionalism*, which offers armchair conceptual analyses of mental concepts in functional terms, and *psychofunctionalism*, which offers scientific analysis of mental concepts in functional terms. I don't think this is the best way to put things, at least for our purposes. This model suggests that analytic functionalism and psychofunctionalism are rival theories of ordinary mental concepts (and perhaps of their referents, ordinary mental properties). What we are concerned with here are compatible theories of two different things: ordinary concepts and theoretical concepts (and perhaps of their referents, ordinary properties and theoretical properties).

The distinction between *commonsense functionalism* and *theoretical functionalism* is more apt for our purposes.[4] The former analyzes the ordinary commonsense term of belief, memory, or whatever, as used by ordinary people. The latter analyzes the theoretical term, as used by theo-

rists. Both may well be forms of analytic functionalism, and they need not be in conflict with one another. One gives a conceptual analysis of an ordinary concept, the other gives a conceptual analysis of a theoretical concept. I think that much of what goes under the label psychofunctionalism in the philosophy of mind is best understood as theoretical functionalism.

Using this framework, spacetime functionalism is clearly a form of theoretical functionalism, as the concept of spacetime is a theoretical concept. Spatial functionalism comes in both commonsense and theoretical varieties, as there are both ordinary and theoretical concepts of space. I have touched on both of these in my work on spatial functionalism (for functionalism about the theoretical concept, see pp. 332 and 418-9 of *Constructing the World*, as well as section 5 above), but my primary focus has been on the ordinary concept and on commonsense functionalism.

As a result, it is no surprise that the analyses delivered by my spatial functionalism and spacetime functionalism look very different. Spacetime functionalism uses a spacetime role tied heavily to the role of space within theories such as general relativity. My spatial functionalism has focused on a space role tied more to folk physics, including the role of space in appearance, motion, and interaction.

This leads to one striking difference, already touched on in the last section: spatial functionalism often involves phenomenal roles (roles in causing spatial experience) whereas spacetime functionalism typically abjures those roles. I don't think this is an essential difference. As I discussed in the last section, there are versions of ordinary spatial functionalism that invoke only nonphenomenal roles such as roles in motion and interaction. In principle there are also versions of theoretical spacetime functionalism that invoke phenomenal roles, since a connection to observation can be seen as part of the relevant theory. Still, it is easy to see why phenomenal roles may be more central to the ordinary concept than to the theoretical concept.

Another difference between the projects is that the literature on spatial functionalism has focused mainly on specific spatial concepts such as distance, shape, and size. By contrast, as David Baker (2019) and David Yates (this volume) observe, the spacetime functionalism literature has focused largely on the overarching concept of *being a spacetime* and not on the specific concepts. Baker suggests that an analysis of the overarching concept might be leveraged into an analysis of specific concepts if we appeal to principles of general relativity whereby a spacetime structure

[4]The labels "commonsense" and "theoretical" are not perfectly apt, since analyses of an ordinary concept can be theoretical (through a folk theory) and need not obviously be tied to commonsense. "Quotidian" vs "scientific" functionalism might be better. But there are enough labels and distinctions floating around that I do not want to introduce more.

fixes an underlying metric. However, these principles will then have to be built into the analysis of the specific concepts, and it is not clear that the resulting concepts will be applicable outside the context of general relativity. So at least more work needs to be done to analyze specific spatiotemporal concepts in the framework of spacetime functionalism.

What is the correct analysis of the overarching theoretical concept of spacetime? In work on spacetime functionalism, there has been some debate over just what counts as the spacetime role. Knox focuses on an inertial role where spacetime is whatever fixes a structure of local inertial frames. James Read and Tushar Menon focus on an operational role tied to the way spacetime affects rods and clocks. David Baker suggests a cluster concept with many different components of the role. The argument between them can resemble old-school conceptual analysis. For example, Baker and Read and Menon call Knox's analysis into question by offering counterexamples in the form of theories that seem to have a spacetime but that do not meet her analysis. It is hard to see what settles the question of which analysis is correct. Does this turn on our intuitions about which theories and scenarios are best described as involving "spacetime", or on psychological and linguistic facts about use of the key terms by physicists? It is not obvious why we should care deeply about these facts except as a matter of the sociology of science. And in practice, even participants in these debates (including Knox) are happy to depart from these intuitions when it suits their purposes.

My own metaphilosophy here dictates *conceptual pluralism* (see Chalmers 2011). There are many interesting concepts in the vicinity of most important terms, and what matters most in choosing among them is not linguistic usage but explanatory roles. 'Space' and 'spacetime' are no exceptions here. In particular, there may be more than one theoretical concept of spacetime. The reason is that there are many different theoretical roles that spacetime plays, and for each role there is in principle a concept of spacetime as what plays that role. Some of these may come closer to actual usage of "spacetime" than others, but this may just reflect contingent facts about usage. More important, perhaps, is that different concepts of space wil be important for different purposes. For example, Read and Menon argue that Knox's inertial concept is more useful than Baker's cluster concept precisely because it "has the capacity to be put to novel interpretative work" in classifying what counts as spacetime under various theories. On the other hand, the operational concept may more closely track the way that spacetime is measured. And a phenomenal concept of spacetime, which connects spacetime in the theory not just to rods and clocks but to conscious observation, may help bridge the gap between the scientific image and the manifest image.

Some degree of conceptual pluralism may also be true for the ordinary concept of space. One

strand in it involves a primitive concept of Edenic space. Another strand involves the role of space in beinging about spatial experience. Still another involves the role of space in motion and interaction. I think all of these strands are present in the ordinary conception of space. In pre-scientific days, the primitivist concept may have been the most useful, but even then the functionalist strands were present. After developments in science, the functionalist concepts became more useful, because they manage to pick out something important in the external world. As always, all of the concepts are available, but some are more useful than others. It is the functional concepts that best serve to locate the key elements of our pretheoretical conception of space (the manifest image) in light of our best scientific theories (the scientific image).

8 The emergence of spacetime

One of the central motivations for spatiotemporal functionalism is the one suggested by my title: to help locate space in a world in which space is not fundamental. In "The Matrix as Metaphysics" (2003, section 3 and note 14), I used the possibility that space emerges from a deeper level as part of my argument for the broadly functionalist view of space articulated there. In recent years, this motivation has received increased attention. Discussion of "emergent spacetime" has exploded, driven largely by theories of quantum gravity—including versions of string theory, loop quantum gravity, and causal set theory – in which spacetime may not appear on the fundamental level. The relevant notion of emergence is not the philosopher's metaphysically demanding notion of strong emergence but the scientist's less demanding notion of weak emergence (see Le Bihan 2018, Baron 2019 and Wilson, this volume, for discussion). The key thesis is that spacetime exists at a nonfundamental level and is grounded in a fundamental level which is nonspatiotemporal. A number of philosophers have suggested that spacetime functionalism can play a role in this grounding. In this final section I will draw some connections between my approach and this recent literature, aiming to get clearer about the prospects for functionalist approaches in vindicating emergent spacetime.

The idea of emergent spacetime has sometimes been received with skepticism.[5] Underlying much of this skepticism is the issue I discussed in section 4—that there seems to be an explanatory gap between the nonspatiotemporal and the spatiotemporal. Whatever emerges from the nonspatiotemporal could not genuinely be *spacetime*. A number of theorists have used analogies with explanatory gaps in the philosophy of mind to flesh out this point.

[5]For articulations of this skepticism, see Le Bihan 2018 andforthcoming, Maudlin 2007, and others.

Baptiste Le Bihan (forthcoming) articulates the worry by suggesting that there is a "hard problem of spacetime" that is not addressed by standard approaches. Perhaps these approaches can solve the "easy problems" of recovering some of the mathematical structure of spacetime theories from a nonspatiotemporal basis. But they do not solve the "hard problem" of explaining "spacetime qualia". By this, Le Bihan means certain special qualities of spatiality and temporality that we associated with spacetime. The word "qualia" is perhaps misleading, as it strongly suggests mental qualities of conscious experience, and Le Bihan is clear that he does not mean this. So I will speak simply of *spacetime qualities* instead.

In my framework, the hard problem of spacetime is the task of explaining Edenic space and time: space and time as they appear to be in our experience. There is a problem of explaining how objects in the world could be laid out in a genuine space of the sort that there seems to be in our spatial experience of the world, with time genuinely passing as it seems to be in our temporal experience. All that can seem to require special qualities of Spatiality and Temporality. I would identify these qualities with Edenic spatiality and temporality.[6]

The hard problem of spacetime is analogous to what Alex Byrne has called "the hard problem of color": explain how things in the world could have genuine colors like red or green, as they appear to be in our experience, in a world where colors are not fundamental. In my framework, Byrne's problem is the problem of explaining how physical objects could have Edenic colors: the primitive sensory qualities such as Redness and Greenness that color experience seems to present objects as having. There seems to be a gap between the non-color-involving properties of the world that physics gives us and genuine Color. Byrne thinks that the hard problem of color is at least as important as the hard problem of consciousness.

The most common reaction to the hard problem of color is to deny that objects have the special qualities such as Edenic redness and greenness that they seem to have. Those qualities do not exist in ordinary physical objects. At best they exist or are represented in the mind. At some level perception may present apples as having these special qualities, but to this extent perception involves an illusion. We solve the hard problem of color by removing the special qualities from the world and putting them into the mind, where they become an aspect of the hard problem of

[6]One difference is that Le Bihan does not question whether general relativity can accomodate his spacetime qualities (perhaps because spacetime is fundamental in general relativity) while I say that it cannot accommodate Edenic space and time (because space and time are not absolute in general relativity). On my picture, relativistic spacetime is fundamental but not Edenic. One could perhaps reconcile these pictures by holding that the relativistic fall from Eden is only partial and that there are some aspects of Edenic spacetime are compatible with a relativistic world.

consciousness. In the external world, all we have are small-c colors, understood functionally as properties that play the color roles.

If my analysis of the explanatory gap for space (in section 4) is correct, we should respond to le Bihan's hard problem of spacetime in the same way. That is, we should deny that the physical world has the special qualities of Edenic spatiality or temporality. At best these qualities exist or are represented in the mind. At some level perception may present the physical world as having these special qualities, but to this extent perception involves an illusion. We solve the hard problem of spacetime by removing the special qualities from the world and putting them into the mind, where they become an aspect of the hard problem of consciousness. In the physical world, all we have is small-s space and small-t time, understood functionally as properties that play the space role and the time role.

In principle, an analogous view is available for the hard problem of consciousness. We could deny that consciousness has the special qualities it seems to have, and say this is a sort of introspective illusion. This is the move made by illusionists about consciousness such as Daniel Dennett and Keith Frankish. Most people find this view much harder to believe than the corresponding view for color and space, in part because the existence of the relevant qualities of consciousness seem to be something of an introspective datum. This apparent datum is one of the things that makes the hard problem of consciousness hard. By contrast, the idea that physical objects have special qualities of color, space, and time does not seem to be a datum in the same way. As a result, the hard problems of color and of spacetime are easier to dismiss than the problem of consciousness.

If I am right, then as long as quantum gravity theories can plausibly say there is spacetime but not Spacetime, they can avoid the hard problem of spacetime. And as I argued in section 5, it is arguable that functionally construed spacetime is enough to save the appearances of our sensory experience, at least if we accept that functional isomorphs of ordinary perceivers will have the same sensory experiences. If so, then spatiotemporal functionalism holds out the promise of an empirically and philosophically adequate theory of spacetime.

In a similar vein, Lam and Wüthrich (2018) have raised the problem of "spacetime zombies" for spacetime functionalism. It seems conceivable that something could play the spacetime roles perfectly, without there being genuine spacetime. If so, being spacetime is not just a matter of playing the spacetime role.

One can treat this suggestion as I suggested in section 5, saying that the zombie intuitions arise primarily for Edenic or primitive space. That is, Spacetime zombies may well be possible, but not spacetime zombies. If we can plausibly deny that our world has Spacetime, this is no problem.

It is also worth noting that realizer functionalism can accommodate a certain sort of spacetime zombie even without invoking Edenic space. We might imagine a perfect functional simulation of our universe in a computer—but the simulation will not have the spatial properties of the universe (e.g. it may be much smaller!). A realizer functionalist will explain this by noting that spacetime is identified with the *actual* realizer of the spacetime role in our normal environment. The simulation has a different realizer, which is not spacetime. This sort of spacetime zombie is quite consistent with the functionalist approach.

Sam Baron (2019) and David Yates (this volume) present related dilemmas for the realizer functionalist approach to spacetime emergence. The first horn of the dilemma focuses on the possibility that the realizers are themselves fundamental properties in a quantum gravity theory. In this case, these spatiotemporal properties will themselves be fundamental properties and we will not have a genuine form of spacetime emergence. We can escape this problem by holding that the realizers are nonfundamental properties. But then we need a substantial account of how the nonfundamental realizer properties are grounded in fundamental nonspatiotemporal properties. Spacetime functionalism tells us how the realizer properties ground spacetime, but it tells us nothing about how these properties are grounded themselves.

How might the realizer properties be grounded in the nonspatiotemporal? One strategy appeals to mereological composition: fundamental nonspatiotemporal entities stand in a part-whole relation to spatiotemporal realizer properties. Baron argues that this cannot work, as there are conceptual connections between parthood and spatiotemporal location. A realizer functionalist might question this and appeal to a more general notion of composition that is not tied to location. Or perhaps better, they may appeal to non-mereological grounding relations. For example, fundamental causal powers in a nonspatiotemporal ontology may ground higher-level causal powers which may themselves serve as realizers for the functional roles associated with spacetime. Or high-level structural properties may be mathematically abstracted from low-level structural properties. The precise grounding of spacetime in the nonspatiotemporal remains an important open question, but there are many options available here.

This leaves open the question: what sort of functionalist analysis of space or spacetime is most promising for vindicating the emergence of spacetime in quantum gravity theories? This depends on large part on the details of the theories. The general shape of the project is clear: focus on one or more roles played by spacetime in general relativity (say), and make the case that under a certain theory of quantum gravity, a certain nonfundamental structure S will play that role. If that role is our criterion for counting as spacetime, then S will count as spacetime by this criterion.

Ideally, under quantum gravity, there will be some nonfundamental structure S that plays *all* of the roles of spacetime in general relativity, at least under appropriately restricted conditions. General relativity has been so successful that it would be surprising if its roles for spacetime were vindicated only in part. One can reasonably hope to find some S that at least approximates the full relativistic role of spacetime. To do this, we can take the approach suggested by David Yates in this volume: Ramsify the full theory of general relativity, identifying spacetime with whatever plays the full role of spacetime in this theory, and then find a structure S that plays this full role under quantum gravity. On this approach, we need not settle which functional roles of spacetime are essential for spacetime in advance. We use something close to the full functional role to find spacetime in quantum gravity.

Still, what if it turns out (perhaps surprisingly) that under quantum gravity, no structure plays all aspects of the spacetime role? Suppose that a nonfundamental structure S plays some aspects of the standard spacetime role but not others. Then whether we count S as spacetime or not will depend on which functional analysis of spacetime we endorse. For example, S might count as spacetime on an operational analysis but not on an inertial analysis. In practice this may end up as a largely verbal dispute, which we can resolve by distinguishing multiple notions of spacetime. For example we may distinguish inertial spacetime and operational spacetime, and make the case that the theory vindicates one but not the other.

Among specific functional analyses, it is not obvious that Knox's inertial spacetime functionalism is well-suited to finding spacetime with quantum gravity theories. The reason is that Knox's analyses of the spacetime role use some spatiotemporal notions, for example in defining an inertial reference frame. But then, to find spacetime in quantum gravity by finding a structure S that plays the spacetime role, we will already need to have found some spatiotemporal properties within the framework. So we will at least need some further analyses of these spatiotemporal properties to support a reduction.

We can also use commonsense functionalism about space to help locate ordinary space and time under quantum gravity theories. Presumably there will be certain nonfundamental properties that tend to bring about ordinary spatial and temporal experience, at least in relevant contexts. On a phenomenal analysis, these can then be identified with ordinary space and time. One can also use a nonphenomenal analysis, finding nonfundamental properties that play the appropriate folk-physics roles in motion and interaction. One can even combine elements of theoretical and commonsense approaches. For example, we could first make the case that in quantum gravity theories, certain nonfundamental properties play the key theoretical roles from general relativity or

even from Newtonian physics (in relevant contexts). From prior analysis, we know it is plausible that whatever plays these theoretical roles also plays the commonsense roles. In this way, the intratheoretical project of finding one theory's spacetime within a later theory's picture of the world can help with the Sellarsian project of finding the manifest image within the scientific image.

References

Albert, D. 1996. Elementary quantum metaphysics. In (J.T. Cushing, A. Fine A., S. Goldstein, eds) *Bohmian Mechanics and Quantum Theory: An Appraisal*. Boston Studies in the Philosophy of Science, vol 184. Springer.

Albert, D. 2015. *After Physics*. Harvard University Press.

Albert, D. 2019. Preliminary considerations on the emergence of space and time. In (A. Cordero, ed.) *Philosophers Look at Quantum Mechanics*. Springer.

Baker, D. 2019. On spacetime functionalism. http://philsci-archive.pitt.edu/15860/

Baron, S. 2019. The curious case of spacetime emergence. *Philosophical Studies*.

Byrne, A. 2006. Color and the mind–body problem. *Dialectica* 60:223-44.

Chalmers, D.J. 2003. The Matrix as metaphysics. Whatisthematrix.com. Reprinted in *Philosophers Explore the Matrix* (Oxford University Press, 2005).

Chalmers, D.J. 2006. Perception and the fall from Eden. In (T. Gendler and J. Hawthorne, eds.) *Perceptual Experience*. Oxford University Press.

Chalmers, D.J. 2011. Verbal disputes. *Philosophical Review* 120:515-66.

Chalmers, D.J. 2012. *Constructing the World*. Oxford University Press.

Chalmers, D.J. 2017. The virtual and the real. *Disputatio*.

Chalmers, D.J. 2018. Structuralism as a response to skepticism. *Journal of Philosophy* 15:625-60.

Chalmers, D.J. 2019. Three puzzles about spatial experience. In (A. Pautz and D. Stoljar, eds.) *Blockheads!*. MIT Press.

Dees, M. 2015. The fundamental structure of the world: Physical magnitudes, space and time, and the laws of nature. Ph.D. dissertation, Rutgers University. https://rucore.libraries.rutgers.edu/rutgers-lib/47350/PDF/1/

Fredkin, E. 2000. *Digital Mechanics*. Draft of December 2000.

Jackson, F. 1998. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford University Press.

Knox, E. 2014. Spacetime structuralism or spacetime functionalism? Draft of August 10, 2014.

Knox, E. forthcoming. Physical relativity from a functionalist perspective. *Studies in History and Philosophy of Modern Physics*.

Lam, V. and Wüthrich, C. 2018. Spacetime is as spacetime does. *Studies in History and Philosophy of Modern Physics* 64: 39-51.

Le Bihan, B. 2018. Space emergence in contemporary physics: Why we do not need fundamentality, levels of reality and emergence. *Disputatio* 49:71-85, 2018.

Le Bihan, B. forthcoming. Spacetime emergence in quantum gravity: Functionalism and the hard problem. *Synthese*.

Lewis, D. 1994. Reduction of mind. In (S. Guttenplan, ed.) *Companion to the Philosophy of Mind*. Blackwell.

Maudlin, T. 2007. Completeness, supervenience and ontology. *Journal of Physics A* 40:3151-71.

Ney, A. 2015. Fundamental physical ontologies and the constraint of empirical coherence: a defense of wave function realism. *Synthese* 192:310524.

Read, J. and Menon, T. 2019. The limitations of inertial frame spacetime functionalism. *Synthese*.

Rosenberg, G.H. 1997/2004. *A Place for Consciousness: Probing the Deep Structure of the Natural World*. Ph.D. Dissertation, Indiana University, 1997. Revised version published by Oxford University Press, 2004.

Seiberg, N. 2006. Emergent spacetime. https://arxiv.org/abs/hep-th/0601234.

Sellars, W. 1961. Philosophy and the scientific image of man. In *Science, Perception, and Reality*. Routledge.

Smith, B.C. 1996. *On the Origin of Objects*. MIT Press.

Thompson, B. 2003. *The Nature of Phenomenal Content*. Ph.D. dissertation, University of Arizona.

Thompson, B. 2010. The spatial content of experience. Philosophy and Phenomenological Research 81: 14684.

Wheeler, J.A. 1989. Information, physics, quantum: The search for links. In Proceedings III International Symposium on Foundations of Quantum Mechanics.

Wilson, A. this volume. Explanations of and in time. In (B. Le Bihan, N. Huggett, and C. Wüthrich, eds.) *Philosophy beyond Spacetime*. Oxford University Press.

Wolfram, S. 2002. *A New Kind of Science*.

Yates, D. this volume. Thinking about spacetime. In (B. Le Bihan, N. Huggett, and C. Wüthrich, eds.) *Philosophy beyond Spacetime*. Oxford University Press.