

This paper is currently undergoing significant revision;
proceed at your own risk...

Consciousness and Relativity

The relativistic structure of spacetime places constraints on the possible spatiotemporal structure of consciousness (or its neural correlates). These constraints put severe pressure on canonical physicalism, as well as the more naturalistic forms of dualism.

0. Introduction

Developments in modern science, especially modern physics, reputedly favor physicalism, if only by casting alternatives to physicalism in a poor light. In our present state of scientific knowledge, it seems very unlikely that there is any aspect of human behavior that physics can't explain, and therefore unlikely that conscious experience explains anything about our behavior, unless consciousness itself falls within the purview of physics. This compels scientific-minded anti-physicalists to deny experience any explanatory role (at least with respect to the physical domain), which is generally considered bad news for the anti-physicalists (how bad is a matter of ongoing debate).¹

¹Maybe "compels" is too strong. Another option for the anti-physicalist is to hypothesize a widespread overdetermination of behavior by physical and non-physical factors; see (Mills, 1996). For the closure argument for physicalism, see (Kirk, 1979), and, for the classic epiphenomenalist reply, (Jackson, 1982).

Let's grant that physics leaves nothing physical unexplained. Well, physics isn't a black box. It consists of detailed theories, including theories about the structure of time and space. Physics tells us that the physical world has a certain *shape*, to put it crudely, and it seems reasonable to ask whether we can fit consciousness into anything with that shape. To regard this question as settled by the fact that physics explains everything physical would be to assume the truth of physicalism.

The main message of this paper is that we cannot fit consciousness into spacetime, except on terms that are deeply unfriendly to physicalism or any other thoroughly naturalistic account of conscious experience.

Here is an overview of the paper.

§1 starts off with a neglected passage from Bertrand Russell, which points to an argument against the Identity Theory. The argument has three premises: (a) that one sometimes has a conscious experience that instantiates its phenomenal properties (its "qualia") absolutely simultaneously; (b) that if a conscious experience instantiates its qualia absolutely simultaneously, then the experience is confined to a single point of spacetime; and, (c) that if a conscious experience is confined to a single point of spacetime, the experience isn't a purely physical phenomenon.

§2 contains the basic argument for (c), which is that if an ordinary human experience were a physical point-event, the same kind of experience could occur in rocks and other things that are not, in fact, capable

of having experiences like ours (or any experiences). We briefly consider possible panpsychist resistance to this reasoning.

§3 addresses two possible challenges to (b). One rejects relativity theory, the other rejects the idea that physics is in the business of describing the world. I argue that in both cases, what we're asked to reject is more credible than that for which we're asked to reject it.

§4 is the heart of the paper. After addressing some tempting but weak objections to (a), we consider a stronger objection, which is that consciousness never occurs as a momentary state, but only as a temporally extended process. I develop this "process theory" in detail, clarify the threat it poses to the Russellian argument, and venture a guess as to why Russell himself failed to notice the threat. I then criticize the process theory for violating a highly plausible principle of natural supervenience.

§5 takes stock of the arguments to that point, locating the Identity Theory's vulnerability to them in its tacit assumption that spacetime has more structure than it actually does.

§6 develops the Russellian argument against the Identity Theory into a more general argument against any theory that locates consciousness in the same time as physical events, and maintains that the phenomenological facts metaphysically supervene on the physical facts.

§7 surveys possible solutions to the mind-body problem that survive the arguments of this paper.

1. An argument against the Identity Theory

In a neglected essay from 1926, Bertrand Russell writes:

Two events which are simultaneous in my experience may be spatially separate in psychical space, e.g. when I see two stars at once. But in physical space the two events are not separated, and indeed they occur in the same place in space-time. Thus in this respect relativity theory has complicated the relation between perception and physics.²

Russell doesn't elaborate on these remarks, but it seems likely that their target is a natural physicalist picture of how our conscious experiences relate to what goes on physically in our brains (see Fig. 1).

When you look at the constellation Orion, you perceive, among other things, the stars Alnitak, Alnilam, and Mintaka. The stars emit light, some of which enters your eyes. The light stimulates the rods and cones in your retinae, initiating a cascade of events in your optical nerves leading back to the visual centers of your brain. There, a lot of complicated and so far poorly-understood neural processing takes place, culminating in, among other things, conscious visual images of Alnitak, Alnilam, and Mintaka. On the physicalist assumption that these images are brain-events, it's natural to suppose that each of them occurs at a different location in the brain from the others—a supposition that receives support

²(Russell, 1926, 130). To my knowledge, the only others to discuss the implications of relativity theory for consciousness are Henri Poincaré (Poincaré, 1898) and

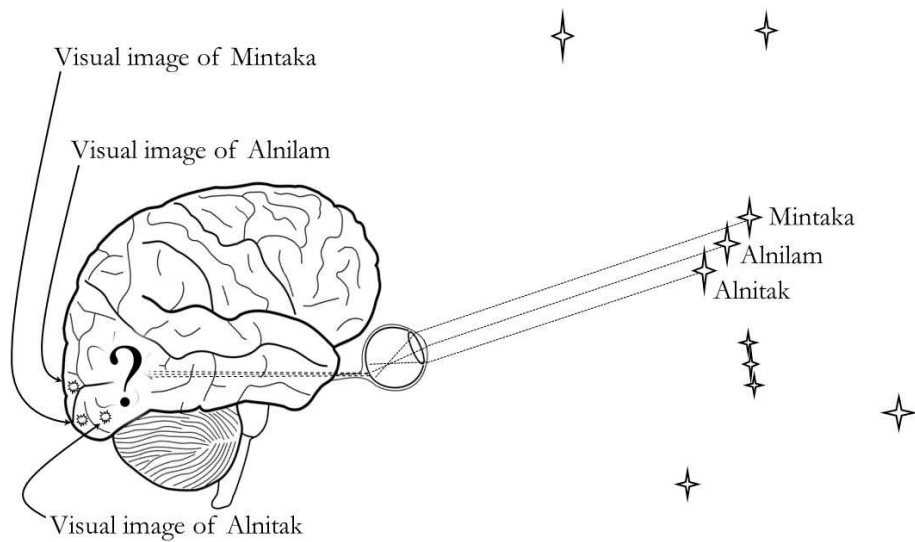


Figure 1: A Natural Physicalist Picture

from the known existence of retinotopic maps within the primary visual cortex.

It’s this last feature of the picture—its representation of the conscious impressions of the stars as occurring at different locations in the brain—that Russell calls into question.

In relativistic spacetime, there is no such thing as absolute distal simultaneity. This means that if we can correctly describe two spatially separated events as simultaneous, we can also correctly describe them

Geoffrey Lee (Lee, 2007); more on their discussions, below. The first edition of (Russell, 1926) also discusses relativity and mind, but the quote is from the second edition.

as occurring successively (by describing them in terms of an alternative frame of reference). Only if two events occur at the same point in both time and space do they occur absolutely simultaneously; so, for example, if a certain mass-property gets instantiated at the same point in space-time as a certain charge-property, then these events (the instantiation of the mass-property and the instantiation of the charge-property) are simultaneous in terms of every complete and accurate description of the world.³

What about your visual impressions of the stars in Orion's Belt? Presumably these occur at the same time; otherwise, how could you perceive the stars as forming a line? According to Russell, this simultaneity is absolute. Any complete and accurate description of the world must describe the experience you have when you look at Orion as involving the simultaneous occurrence of three conscious visual images of stars. A description of your experience that failed to describe you as experiencing the stars simultaneously would either misdescribe or underdescribe your experience.

³In general, to say that it's absolutely the case that p is to say that it is the case that p according to every complete and accurate description of the world; to say that it's relativistically the case that p is to say that it is the case that p in some but not all complete and accurate descriptions of the world. For example, it is absolutely the case that Jupiter has mass, but only relativistically the case that it is in translational motion at any given instant. (The distinction between absolute and relativistic states of affairs occurs throughout science, and isn't exclusive to so-called relativistic physics.)

But if your conscious experiences of Alnitak, Alnilam, and Mintaka occur absolutely simultaneously, then they occur at the same point in spacetime. It follows that your visual image of Orion's Belt isn't a spatially extended complex of events, as depicted in Fig. 1, but a single point-event, or a triplet of strictly co-located point-events.⁴

This is the result that Russell says complicates the relationship between perception and physics. Although Russell doesn't tell us what complication he has in mind, his observations pose an obvious threat to the Identity Theory of consciousness, according to which conscious experiences are identical to purely physical states (such as brain-states). Let's develop this threat as an argument, which I'll call *the chronometric argument against the Identity Theory*:

- P1 One sometimes has a conscious experience that instantiates its various qualia (visual, auditory, etc.) absolutely simultaneously.
- P2 If a conscious experience instantiates its qualia absolutely simultaneously, then the experience is confined to a single point of spacetime.
- P3 Some conscious experiences are confined to single points of spacetime. (from P1, P2)
- P4 If a conscious experience is confined to a single point of spacetime, it isn't purely physical.
- P5 Some conscious experiences aren't purely physical. (from P3, P4)

⁴The theoretical upper-bounds on the dimensions of a point-event are given by the Planck distance (10^{-35} meters) and the Planck moment (10^{-44} seconds). A collection of point-events occurring at the same point in spacetime is itself a point-event.

Call P1 *the phenomenological claim*, P2 *the relativistic claim*, and P4 *the microphysical claim*. In the sections that follow, we assess these claims in reverse order, starting with the microphysical claim.

2. The microphysical claim

At a rock concert, each of my experiences is characterized by visual and auditory qualia (among other phenomenal features). The microphysical claim is that none of these experiences is a purely physical point-event.

Why not? Because any purely physical point-event that occurs in me can also occur in a stone, or a fern, or a ham sandwich, none of which is capable of conscious experience. What distinguishes the physical natures of sentient beings, like us, from the physical natures of insentient beings, like sandwiches, is not that our bodies contain subatomic features that are absent from sandwiches. The distinctive thing about the body of a sentient being isn't that it's made up of distinctive subatomic events, but that the generic subatomic events that constitute it relate to one another in ways that yield distinctive structures and processes (such as neural processes).⁵

⁵According to some versions of the token-identity theory, a physical event *E* might be a phenomenal event, even if some events of the same physical type as *E* aren't phenomenal events. But even if we accept this, any token physical point-event *E* that occurs in a brain could have occurred in a rock (if, for example, the history of the world had developed so that *E* had been surrounded by events constituting a rock instead of a brain).

Some panpsychists maintain that conscious experience is found everywhere, including at the level of individual subatomic events. They might not be troubled by the suggestion that ham sandwiches (or their constituents) are conscious. But, to my knowledge, no panpsychist holds that sandwiches or their constituents are capable of having experiences with the same phenomenal character as ordinary human experiences (such as the ones I have at a rock concert). So, even a panpsychist should be troubled by the suggestion that a ham sandwich could contain experiences just like ours.⁶

In any case, if it turned out that the only way to resist the chronometric argument were by embracing panpsychism, that would be an important and surprising result in its own right.

3. The relativistic claim

The relativistic claim is a straightforward consequence of relativity theory (special and general) as standardly interpreted. To resist the claim, one must either find fault with relativity theory, or cast doubt on the standard interpretation of it.

Like any scientific theory, General Relativity is open to revision, or, in face of the right observations, rejection. Since not all parts of current physics are logically compatible with one another, we know that some

⁶For panpsychism, see (Eddington, 1928), (Chalmers, 1996, 293-301), and (Strawson, 2006).

part needs to be revised or replaced, and General Relativity might prove to be that part.

There are two reasons why this shouldn't concern us very much here.

First, even if some other theory is destined to supersede General Relativity, the superseding theory might retain the feature of General Relativity that drives the arguments of this paper, namely, a spacetime free of the structure required for absolute distal simultaneity.⁷

Second, the only known phenomena that might lead us to reinstate absolute distal simultaneity are quantum phenomena that have no obvious bearing on conscious experience. So, future developments in physics aren't likely to undermine the chronometric argument, unless quantum mechanics unexpectedly turns out to bear directly on the search for neural correlates of consciousness—a turn of events that would be at least as interesting as any refutation of the Identity Theory.⁸

A different way to resist the relativistic claim is by adopting a radical pragmatism about natural science.

We normally think of natural science, including physics, as a descriptive and explanatory enterprise. Physicists describe what they observe, and try to explain it, often successfully; or, so we normally think.

⁷For a glimpse of what a future physics that retains this feature might look like, see (Tumulka, 2006), and the discussion in (Maudlin, 2008).

⁸This is particularly so, considering that none of the phenomena that people point to as evidence of absolute simultaneity give us any basis for identifying any

One might reject this way of thinking. One might suggest that, despite appearances to the contrary, physics is not really in the business of describing and explaining the physical world. Really, one might suggest, physics is just a machine for generating predictions. What makes a physical theory good is not that it offers illuminating descriptions and satisfying explanations of physical phenomena. A good physical theory is just one that generates many accurate predictions, and few inaccurate ones.⁹

If we accept all this, we can agree that General Relativity is a good theory, but deny that it describes a world that conflicts with the Identity Theory, on the grounds that it doesn't describe the world at all.

There are three problems with this manoeuvre.

First, the proposition that the Identity Theory is false seems more credible than the proposition that the proposed pragmatism is true.

Second, it's not clear that the pragmatist proposal even helps, since it might turn out (and, if the arguments that follow are sound, does turn out) that, when combined with General Relativity, the Identity Theory generates many false predictions, and so counts as a bad theory, even by pragmatic standards.

particular pair of spatially separated events as absolutely simultaneous: see (Callender, 2008). It would be remarkable, to say the least, if introspection turned out to provide such a basis.

⁹Something like this seems to be William James's position in (James, 1907).

Third, if physics doesn't explain anything, it doesn't explain human behavior, in which case we're free to hold that our behavior is partly explained by our conscious experience, without committing ourselves to the view that consciousness is a physical phenomenon (or that our behavior gets systematically overdetermined by physical and non-physical factors). So, if a shift to pragmatism undermines the chronometric argument against the Identity Theory, it does so only at the cost of depriving the physicalist of his only real argument for physicalism.

There is one objection to the relativistic claim that arises more from a confusion about the relativity of (distal) simultaneity than from doubts about relativity theory or the standard interpretation of it. You might think that the considerations Russell puts forward merely point to the need to describe a person's conscious mental life in terms of a suitable frame of reference: a frame in terms of which that person's body is at rest. The idea is that a person's rest-frame corresponds to that person's point of view, making this the uniquely appropriate frame in terms of which to characterize his conscious experiences. And of course it is trivially true that any pair of events (spatially separated or not) that are simultaneous in terms of this frame of reference are simultaneous in terms of this frame of reference; i.e., trivially true that it's an absolute fact that any given two events are (or, as the case may be, aren't) simultaneous with respect to that frame.

If all this sounds plausible as an objection to Russell, it's only because you are thinking of a frame of reference as a point of view. But that is a mistake. A frame of reference is just a map, e.g. (in special relativity) an assignment of 4-coordinates (x, y, z, t) to events. It makes no more sense to speak of observing something from a frame of reference than to speak of observing something from a map of Africa. The temptation to single out a particular frame of reference as uniquely suited to the description of a given person's conscious mental life arises from a tendency to think of reference frames as standpoints rather than maps.

Anyway, even if we accept (as we shouldn't) that one must describe a person's experience in terms of his rest-frame in order to describe his experience accurately, that poses no threat to Russell. At least, it poses no threat if we accept relativity theory (as standardly interpreted). Relativity theory (as standardly interpreted) implies that, for the purposes of giving a complete and accurate description of a person's body, it doesn't matter which frame of reference you use. So if you say that it *does* matter which frame you use when it comes to describing a person's conscious mental life—e.g., that only a description of the person's mental life in terms of that person's rest-frame can be accurate—then you are saying, in effect, that there's a difference between a person's conscious mental life and what's given with the physical goings-on in his body, which is precisely what the chronometric argument purports to show.

4. The phenomenological claim

The phenomenological claim is that we sometimes have conscious experiences that instantiate their respective qualia absolutely simultaneously. (This is compatible with saying that we also have experiences that instantiate qualia successively. For example, one might hold that we have many relatively simple experiences, each consisting of the simultaneous instantiation of multiple qualia, and further, more complex experiences consisting of temporal sequences of those simpler experiences.)¹⁰

Russell takes it to be just obvious that, whenever one has any conscious experience, one has experiences that instantiate their respective qualia at the same time, and absolutely so. For example, when I see a white horse standing in a green field next to a red barn, I have an experience that is absolutely simultaneously characterized by phenomenal whiteness, phenomenal greenness, and phenomenal redness; if I hold a crying baby, I have an experience that is absolutely simultaneously characterized by phenomenal redness, phenomenal softness, and phenomenal loudness; etc. Could anything be more obvious?

Well, yes. The phenomenological claim is not as uncontroversial as Russell seems to think. It *is* obvious that we often have experiences as of simultaneous states of affairs; the experiences of the horse in the field

¹⁰(Dainton, 2006) develops a theory roughly along these lines.

and the crying baby are cases in point. But from this it doesn't obviously follow that we have experiences, or instantiate distinct qualia, simultaneously. "I have an experience as of a baby who is simultaneously loud and red" doesn't obviously entail: "I simultaneously have an experience as of a baby who is loud and an experience as of a baby who is red." Maybe the latter follows from the former, but if so, this is not something we can just take for granted, as Russell seems to.

I said that it's obvious that we have experiences as of simultaneous states of affairs, such as an experience as of a baby who is at the same time loud and red. It saddens me to report that there are nevertheless people who profess to doubt this.

You might think that these people are just confusing the claim that the baby consciously *appears* red and loud with the claim that I simultaneously *sense* the baby's redness and loudness. If sensing is a matter of having sense-organs that are appropriately affected by what is sensed, it's conceivable that I don't sense the redness and loudness simultaneously—for instance, it could be that the receptors in my eyes are never stimulated at the same time as the receptors in my ears.

Alas, even after all such distinctions have been drawn, there are those who persist in maintaining that it might be an illusion that there is even a conscious appearance of something that is at once loud and red. But this kind of skepticism can never really get off the ground.

Certainly there can be an illusion by which an object that isn't simultaneously loud and red appears to be simultaneously loud and red. I might be subject to such an illusion while observing a beacon that alternately emits red flashes and loud clangs. If so, it's because the intermittent flashes and clangs cause me to have a conscious experience as of simultaneous color and sound. Possibly this is even how it is with the baby. But how can it be an illusion that I have a conscious *experience* as of simultaneous color and sound? Any such "illusion" would have to involve a conscious appearance of color and sound. But then the existence of such an appearance wouldn't be illusory after all. It would be a plain matter of fact.¹¹

Could one argue that the illusion is purely cognitive, so that there is no need to posit any troublesome illusory experience? Hardly. Here I am, confronting the baby. It's not as though my inclination to assert, "I'm having an experience as of a baby who is at once loud and red" comes as an alien impulse, or diminishes upon further consideration of the experience. I'm not like one of those deluded sightless people who claim to have full powers of vision—or if I am, that's only because, or only if,

¹¹The claim here is just that in order for there to be a conscious illusion of a seen and heard state of affairs, there must be a corresponding conscious audiovisual experience. It doesn't follow from this that we are infallible or omniscient about the phenomenal character of our experience, or that introspection reveals to us the essences of our experiences.

those people have pathological visual *phenomenology* that they mistake for veridical experience.

Not only should it be uncontroversial that we have experiences as of loud red babies: it should also be uncontroversial that it is absolutely the case that we have such experiences. It should be uncontroversial that any description of our conscious mental lives that fails to describe them as including multimodal (e.g., audiovisual) experiences is either inaccurate or incomplete. The fact that we have such experiences has at least as much claim to absoluteness as the fact that the Moon orbits the Earth.

Although there can be no reasonable doubt that we have experiences as of loud red babies (and absolutely so), the question remains whether this entitles us to say, with Russell, that we have experiences that simultaneously instantiate phenomenal redness and phenomenal loudness.

One way to resist Russell on this point is by saying that no experience has more than one phenomenal property. On this view, when I confront the baby, my experience doesn't instantiate a property of phenomenal redness and a distinct property of phenomenal loudness, but only a single, complex phenomenal property of redness-plus-loudness.

Of course, in a realistic case, my experience will, *prima facie*, instantiate more phenomenal properties than redness and loudness; it will include additional visual and auditory phenomenology, as well as phenomenology associated with other sensory modalities (such as phenom-

enal odor), and probably proprioceptive, nociceptive, and cognitive phenomenology too. The suggestion we're considering is that however rich the phenomenology of my experience may be, the experience has just one multimodal phenomenal property, rather than multiple properties within and across modalities.

But if that were so, there would be only two degrees of phenomenal similarity: complete similarity, and complete dissimilarity. If each conscious experience has one, and only one, quale, then any two experiences must have either all or none of their qualia in common. Since it is evident that there are many more than two degrees of phenomenal similarity (or dissimilarity), we should reject this proposal.

Anyway, even if the "uni-quale" hypothesis were true, it would pose no threat to Russell. He could just reply that each quale-instantiation is at least absolutely simultaneous with itself.

However, the ease of this last reply should raise our suspicions. If the fact that an instantiation of phenomenal redness occurs absolutely at the same time as itself is enough to show that it's confined to a single point of spacetime, why doesn't the fact that an instantiation of *physical* redness (or softness, or triangularity) occurs absolutely at the same time as itself show that *it* is confined to a single point of spacetime?

Evidently, there is some ambiguity in the phrase, “at the same time as.” We can resolve the ambiguity with a distinction between *simultaneity* and *concurrency*.

Simultaneity is a relationship between *events*, as events are understood in relativistic physics (i.e., as the physical contents of, or goings-on at, individual points in spacetime). As noted already, only events that occur at the same point in spacetime occur absolutely simultaneously.

Concurrency is a relationship between *processes*, temporal sequences of events in which each event occurs absolutely earlier or later than each other event. In the standard jargon, the events that constitute a process are time- or light-like separated from one another; any causal sequence of events is a process, in this sense.

Just as we distinguish absolute from relative simultaneity, we can distinguish absolute from relative concurrency.

Two processes, A and B, are *absolutely concurrent* just in case the first event in A is absolutely simultaneous with the first event in B, and the last event in A is absolutely simultaneous with the last event in B. If A and B are absolutely concurrent, then every complete and accurate description of the world describes them as beginning at the same time and ending at the same time.

Processes A and B are *relativistically concurrent* just in case the first event in A is relativistically simultaneous with the first event in B, and the

last event in A is relativistically simultaneous with the last event in B. If A and B are relativistically concurrent, then there is a complete and accurate description of the world that describes the processes as beginning and ending at the same time, but also complete and accurate descriptions of the world that describe A as beginning before B begins and ending before B ends, as well as ones that describe A as beginning after B begins and ending after B ends.

These distinctions drawn, we can now consider a more interesting objection to the phenomenological claim.

Start by granting that in order to have an experience as of a baby that is at the same time red and loud, you must have an experience that instantiates phenomenal redness and phenomenal loudness at the same time, and absolutely so. But now insist that the experience instantiates the qualia absolutely “at the same time” only in the sense that it instantiates them absolutely *concurrently*, and not absolutely simultaneously. Since absolute concurrence, unlike absolute simultaneity, doesn’t entail confinement to a single point of spacetime, this move threatens to stop the chronometric argument in its tracks.

Call this the *process theory*. The theory holds that every conscious experience is a temporal sequence of purely physical point-events. None of these events is itself sufficient for consciousness, but they collectively

constitute a process that is, according to the process theory, sufficient for (because identical with) a conscious experience.¹²

Each of the experiences I have while holding the crying baby has multiple phenomenal properties. For example, I have (among others) an experience characterized by phenomenal redness and phenomenal loudness. The experience doesn't instantiate these properties successively (I don't perceive the baby as being first red, then loud, or first loud, then red), but neither, according to the process theory, does it instantiate them simultaneously (as Russell claims). Rather, the experience has the properties concurrently.

Since the audiovisual phenomenology of my experience is absolute, the process theory must say that the experience has its auditory and visual properties *absolutely* concurrently. Figs. 2 and 3 illustrate why.

Fig. 2 depicts my audiovisual experience as two coterminous processes. Here, the process comprising the *a*-events is the visual component of the

¹²Geoffrey Lee suggests that relativistic considerations favor the process view (or something very much like it); in particular, he suggests that the best way “to conceive of the stream [of conscious experience] is as a 4-D array of physical events in the brain,” which “could only be a conception of experience if physicalism is true.” (Lee, 2007, 371) Lee thus sees relativity as constraining the space of possible solutions to the mind-body problem in favor of physicalism, whereas I see it constraining that space in favor of anti-physicalism. Despite this disagreement, Lee and I agree with each other, and with Poincaré and Russell, that relativity does constrain the space of solutions in *some* important way.

experience (the instantiator of the experience's phenomenal redness), and the process comprising the b -events is the experience's auditory component (the instantiator of the experience's phenomenal loudness). 2a and 2b represent the same processes (consisting of the same events); the difference is that 2b represents the processes in terms of a frame of reference (x', t') that is boosted 0.85 times the speed of light, relative to the (x, t) frame of reference employed in 2a.

The key thing to note is that in both frames of reference, the processes begin at the same time and end at the same time: the first event in the a -process (event a_1) coincides with the first event in the b -process (b_1), and the last event in the a -process (a_n) coincides with the last event in the b -process (b_m). So, Fig. 2 succeeds in depicting my visual and auditory phenomenology as absolutely concurrent.

Fig. 3 depicts my audiovisual experience as two parallel processes. Again, the a -process is the visual component of my experience, the b -process its auditory component. 3a and 3b represent the same processes, the only difference being that 3b represents them in terms of a frame of reference boosted $0.85c$, relative to 3a. The key thing to note here is that in 3b, the process that instantiates my auditory phenomenology (the b -process) ends before the process that instantiates my visual phenomenology (the a -process): by the time I'm done having phenomenally




Fig. 2a




Fig. 2b

Figure 2: The Process View




Fig. 3a




Fig. 3b

Figure 3: Not the Process View

loud experience (at t'_3), nothing has happened in me that is sufficient for phenomenally red experience. Fig. 3 therefore wrongly depicts the audiovisual character of my experience as a representational artifact, rather than the absolute feature of my experience (and the world) that it is.

There are several spatiotemporal structures for experience that are compatible with the process theory.

One possibility is that the visual component of the experience (its phenomenal redness) and the auditory component (its phenomenal loudness) are spatiotemporally coincident processes. That is, it could be that the visual component is a sequence of physical point-events a_1, a_2, \dots, a_n , and the auditory component a sequence of physical point-events b_1, b_2, \dots, b_n , such that for all i , a_i occurs at the same point in spacetime as b_i . Call this the “single-process model.” (See Fig. 4, in which “ a_1b_1 ” represents the coincident occurrence of a_1 and b_1 , “ a_2b_2 ” the coincident occurrence of a_2 and b_2 , etc.)

Another possibility is that the visual and auditory components are non-coincident but coterminous processes. That is, it could be that the visual component is a sequence of physical point-events a_1, a_2, \dots, a_n , and the auditory component a sequence of physical point-events b_1, b_2, \dots, b_m , such that a_1 occurs at the same point in spacetime as b_1 , and a_n at the same point in spacetime as b_m , but not all of the intermediate




Figure 4: The Single-Process Model




Fig. 5a




Fig. 5b




Fig. 5c

Figure 5: The Bundled-Process Model

events in the *a*- and *b*-sequences coincide. Call this the “bundled-process model.” (In Fig. 5a, the auditory and visual processes do not coincide at any of their intermediate stages; in Fig. 5b they coincide at some but not all of their intermediate stages, but not at their second or penultimate stages; in Fig. 5c, the processes begin and end with a run of coincident events, but contain some non-coincident events in the middle.)

The bundled-process model has more neuroscientific plausibility than the single-process model. Even at the neuronal scale, the activity associated with any conscious experience appears to involve extensive parallel processing, and it’s well-established that the neural structures primarily associated with different sensory modalities (e.g., vision, audition, and olfaction) are located in different regions of the brain. When we descend to the sub-neuronal level, the relative plausibility of the bundled-process model is even more evident. To take just one example, it is unlikely that all of the billions of ion-transport events that occur in the cell-membranes of the neurons involved in an audiovisual experience while that experience is underway occur in a linear causal sequence.¹³

¹³Even if the ion-events occur at time-like or light-like separation from one another, it’s extremely unlikely that they have a linear ordering in terms of (relevant) causation, and so extremely unlikely that they empirically *must* occur in time- or light-like separation, in order for the neural processing of which they are a part to occur. (Analogously, suppose the servants open all the mansion’s windows as part of the Spring cleaning. Even if the servants all act within a millisecond of each other, the window-openings might still occur in a time-like sequence (assuming a mansion of re-

For present purposes, however, the exact neuroscientific details don't matter. For, no matter which model the process theorist chooses, his theory faces a serious objection.

One useful way to classify theories of consciousness is by the stances they take on various supervenience claims. For example, it is standard to classify as physicalist any theory according to which the phenomenological facts metaphysically supervene on the physical facts; that is, according to which any two metaphysically possible worlds that duplicate each other physically also duplicate each other in terms of what conscious experience they contain.¹⁴

A more modest claim is that the phenomenological facts *naturally* supervene on the physical facts. Here, the claim is that it is a consequence of the natural laws that govern our world that in it, beings that duplicate one another physically also duplicate one another phenomenologically.

While all dualists reject metaphysical supervenience, most of them accept natural supervenience. (Needless to say, physicalists also accept it.) Unlike metaphysical supervenience, which is highly controversial

alistic proportions). But the openings don't *have* to occur in a time-like sequence in order for the Spring cleaning of which they are a part to take place. The same cleaning gets done if the window-opening events occur relativistically simultaneously.)

¹⁴Here we set aside debates about how best to define supervenience; they are irrelevant to the arguments that follow.

and, in my opinion, false, natural supervenience is highly plausible, and, I assume without argument, true.¹⁵

Equally plausible is the slightly stronger claim that the phenomenological facts naturally supervene on the *molecular* facts—molecular in the chemical, rather than logical, sense. According to what we may call *natural chemical supervenience*, or “chemical supervenience” for short, it is a consequence of the natural laws that govern our world that in it, beings that duplicate one another chemically—molecule-for-molecule doppelgangers—duplicate one another phenomenologically.

Instead of trying to give an argument for molecular supervenience, I’ll just say that it would be amazing if it were false. If the real moral of the story that follows is that it’s empirically possible to replicate someone’s molecular structure without replicating his experience, that is a more momentous result than any that I’m actually trying to achieve here.

Two beings can duplicate one another chemically without duplicating one another at the Planck scale. In particular, two human beings could duplicate one another molecule-for-molecule over the course of their entire lives, despite the fact that the life of one of those human beings contained one less physical point-event than the life of the other.

At least, this seems to be an inevitable consequence of the vast gulf that separates the chemical scale from the Planck scale (which gives the

¹⁵For arguments in favor of natural supervenience, see (Chalmers, 1996, 247-75).

upper-bound for the dimensions of a point-event). The molecules that make up a brain-cell are measured in nanometers (10^{-9} meters); the Planck length is 10^{-35} meters. A conservative lower-bound on the duration of the neural processing involved in a typical conscious experience is 10^{-4} seconds; the Planck moment is 10^{-44} seconds. To put this in perspective, 10^{44} is approximately the number of molecules in the Earth's atmosphere. To suggest that two humans might differ chemically despite differing physically only by a single point-event is like suggesting that two planets might differ in their weather despite differing physically only by a single atmospheric molecule.¹⁶

Chemical supervenience receives further support from the fact that all of the most plausible candidates for neural correlates of consciousness—the neural realizers or instantiators of qualia—identify the correlates with states of or relations among neurons (or groups of neurons). In the most plausible accounts of the neural correlates of consciousness, sub-neuronal structure and function are irrelevant, except to the extent that they bear on the input-output architecture and signalling behavior of the neuron as a whole.¹⁷

¹⁶Or, for a more exact analogy, consider two Earth-sized planets whose atmospheres differ only by one tenth of one billionth of a snowflake (a snowflake contains about 10^{14} molecules).

¹⁷See (Chalmers, 2000) for a survey of options in the search for neural correlates of consciousness. The literature contains one proposal that, if correct, would make

In such accounts, the neuron is a black box whose detailed inner workings are irrelevant to the existence or quality of a subject's experience. We happen to know that in us, neurons are made of proteins configured in various ways, but a brain composed of neurons made of silicon circuits or metal gears would realize the same phenomenal states as a protein brain, provided that the silicon or metallic neurons performed the same functional roles as the corresponding protein neurons, and that the connections among corresponding neurons were the same.¹⁸

Seen in this light, the suggestion that the occurrence or non-occurrence of a single subatomic event might make the difference between having and not having a certain conscious experience is no more plausible than the suggestion that the occurrence or not of a single subatomic event might make the difference between a certain gear's turning or not turning. If conscious phenomenology naturally supervenes on neuron-scale functionality, then chemical supervenience must be true, since differences that preserve molecular structure also preserve neuron-scale functionality.

subatomic phenomena relevant to consciousness, namely the proposal of (Hameroff & Penrose, 1996) that consciousness has its basis in quantum processes that ostensibly occur certain sub-cellular structures. This proposal has few defenders and many critics, but this isn't the place to wade into that debate. Suffice it to say that if the only way to escape the present arguments is by embracing a quantum theory of consciousness, that would be an important and surprising result.

¹⁸This claim, plausible in its own right, receives further support from the arguments of (Chalmers, 1996, 247-75).

The problem with the process theory is that it violates chemical supervenience. We establish this first for the single-process model of the theory, then for the bundled-process model.

Single-process model

Let Fig. 4 depict a process that the identity theorist identifies with an audiovisual experience of the sort one might have while observing a crying baby. Suppose that such a process occurs in Morton, and suppose that Lester duplicates Morton physically, except for that Lester doesn't contain a duplicate of event a_nb_n (the coincident occurrence of physical point-events a_n and b_n). This is an empirically possible scenario, assuming, with the process theorist, that the Morton process is itself empirically possible. (If you like, you can imagine that Lester gets vaporized by an atomic blast just before a_nb_n has a chance to occur.)

According to the process theory, Morton's audiovisual experience just is the $\langle a_1b_1 \dots a_nb_n \rangle$ process (the sequence of events starting with a_1b_1 and ending with a_nb_n). Since this process doesn't occur in Lester, he doesn't have the corresponding audiovisual phenomenology, unless the $\langle a_1b_1 \dots a_{n-1}b_{n-1} \rangle$ process is sufficient on its own for that phenomenology.

If the $\langle a_1b_1 \dots a_{n-1}b_{n-1} \rangle$ process *is* sufficient for the audiovisual phenomenology, the natural thing for the process theorist to do would be to identify Morton's audiovisual experience with the $\langle a_1b_1 \dots a_{n-1}b_{n-1} \rangle$

process that occurs in him, instead of the process ending with $a_n b_n$. In that case, we consider a new character, Morton', identical to Lester, and a new character, Lester', who duplicates Morton' physically, except that Lester' doesn't contain a duplicate of event $a_{n-1} b_{n-1}$. Again, according to the process theory, Morton' and Lester' differ phenomenologically, unless the $\langle a_1 b_1 \dots a_{n-2} b_{n-2} \rangle$ process is sufficient on its own for the relevant audiovisual phenomenology. In the latter case, we reiterate the foregoing reasoning until we reach a process that does not contain any subprocess sufficient for the audiovisual phenomenology, and point out that the corresponding versions of Morton and Lester differ phenomenologically, according to the process theory, despite the fact that they are chemical doppelgangers.

But what if the process theorist doesn't do the natural thing? Suppose he says that Morton's audiovisual experience really is the $\langle a_1 b_1 \dots a_n b_n \rangle$ process, even though $\langle a_1 b_1 \dots a_{n-1} b_{n-1} \rangle$ process on its own would be sufficient for that experience. Then we just consider a version of Morton, Morton', who doesn't contain $a_n b_n$ (but is otherwise indistinguishable from Morton), and a version of Lester, Lester', who doesn't contain $a_{n-1} b_{n-1}$ (but is otherwise indistinguishable from Lester). According to the process theory, Morton' and Lester' differ phenomenologically, despite duplicating one another at the molecular level. Or, if the process theorist denies this by claiming that $\langle a_1 b_1 \dots a_{n-2} b_{n-2} \rangle$ would be enough

for the audiovisual phenomenology, we just reiterate the reasoning until he can no longer make such a claim. (There must come a point at which he can no longer make it, since otherwise he'd have to admit that there's an empirically possible case in which an audiovisual experience is identical to a single physical point-event.)

Objection: “In the logical sequence of processes $\langle a_1 b_1 \dots a_n b_n \rangle$, $\langle a_1 b_1 \dots a_{n-1} b_{n-1} \rangle$, $\langle a_1 b_1 \dots a_{n-2} b_{n-2} \rangle$, ... $\langle a_1 b_1 \dots a_2 b_2 \rangle$, there's no pair of adjacent processes such that the first is sufficient for audiovisual experience but the second isn't. Instead, there's a middle stretch in which it's vague whether the processes are sufficient for audiovisual experience. So, you can't point to any single event and say: ‘the process theory implies that two people who differ physically only with respect to this event differ phenomenologically.’”

Reply: According to the single-process model, process $\langle a_1 b_1 \dots a_n b_n \rangle$ is identical with a specific conscious experience. This experience may have all kinds of phenomenal qualities, including some to which it's indeterminate whether predicates like “phenomenally loud” or “phenomenally red” apply. But the phenomenology is what it is, regardless of any such semantic indeterminacies. If you have an experience that isn't clearly phenomenally red or clearly phenomenally non-red, but only “vaguely” phenomenally red, that's because you have an experience that clearly and non-vaguely possesses some other phenomenal quality—the one due to

which it's uncertain whether the predicate "phenomenally red" applies to the experience. So, even if we accept that the series includes adjacent processes that are vague with respect to various phenomenal qualities, the identity theorist has to admit that the series includes *some* adjacent processes that differ sharply with respect to *some* phenomenal quality; namely, the processes that mark the transition from experience that has a certain quale non-vaguely to experience that has that quale only vaguely (or vice versa).¹⁹

Bundled-process model

Now we turn to the bundled-process model, which we can divide into three cases: (a) the case in which there is no intermediate coincidence of the coterminous processes (as in Fig. 5a), (b) the case in which there is some intermediate coincidence, but no run of coincident events at the beginning or end of the processing (as in Fig. 5b), and, (c) the case in which there is some intermediate coincidence, and a run of coincident events at the beginning or end of the processing (as in Fig. 5c).

Case (a): no intermediate coincidence

Suppose Morton's life includes the two coterminous processes depicted in Fig. 5a, and Lester duplicates Morton except for that no counterpart


¹⁹We can guarantee the existence of such a transition by setting n to a finite number, e.g., by letting the events constituting the processes be Planck-scale events. For a

of $a_n b_m$ occurs in Lester. And suppose, with the process theorist, that the pair of coterminous processes in Morton is identical to an audiovisual experience, with the left-hand process (from a_1 via a_2 , a_3 , etc., to a_n) being his experience's visual component, and the right-hand process (from b_1 via b_2 , b_3 , etc., to b_m) being its auditory component.


In that case, Lester lacks the audiovisual phenomenology that Morton has, unless the occurrence in Lester of the $\langle a_1 \dots a_{n-1} \rangle$ process is sufficient for the relevant visual phenomenology, and the occurrence in him of $\langle b_1 \dots b_{m-1} \rangle$ sufficient for the relevant auditory phenomenology. But then we can correctly describe Lester as ceasing to have the relevant auditory phenomenology, before anything has occurred in him that's sufficient for the relevant visual phenomenology (see Fig. 6). It follows that Lester lacks the kind of audiovisual experience that Morton has, since if he had it, he'd have it absolutely. This is a violation of chemical supervenience.²⁰

detailed discussion of vagueness in the phenomenological sphere, see (Simon, Forthcoming).

²⁰We get much the same result if we consider a case in which Lester differs physically from Morton only by lacking a counterpart of event $a_1 b_1$; in this case, as in the case we've been considering, the process theory requires us to say that if Lester has the same visual phenomenology and auditory phenomenology as Morton, there is a period of time in which Lester has the visual phenomenology, but in which nothing happens in him that is sufficient for the auditory phenomenology (or vice versa).



Lester in terms of x, t coordinates



Lester in terms of x', t' coordinates

Figure 6: Lester, Case (a)

Case (b): some intermediate coincidence, no terminal runs

Suppose that Morton contains the coterminous processes depicted in Fig. 5b, and Lester duplicates Morton except that no counterpart of event $a_j b_k$ (the coincident occurrence of a_j and b_k) takes place in Lester. The argument for violation of chemical supervenience is just as in the preceding case: if Lester has the same auditory and the same visual phenomenology as Morton, it's only if Lester does *not* having the same audiovisual phenomenology as Morton, for the reasons given above.

Case (c): some intermediate coincidence with terminal runs

Suppose Morton contains the coterminous processes depicted in Fig. 5c, and Lester duplicates Morton except for that no counterpart of $a_{j+i} b_{k+i}$ occurs in Lester. Down to $a_j b_k$, it's the same story here as in the single-process case: either Lester's lack of $a_{j+i} b_{k+i}$ leaves him without any visual or auditory phenomenology like Morton's, or not. If not, we consider a new case, with Morton containing the processing that ends with $a_{j+(i-1)} b_{k+(i-1)}$, and Lester physically identical to Morton but for the non-occurrence in Lester of any counterpart of $a_{j+(i-1)} b_{k+(i-1)}$, and so on.

If the process theorist says that the processing extending from $a_1 b_1$ to $a_j b_k$ is enough for Morton's audiovisual experience, we consider a version of Morton, Morton', whose processing ends with $a_j b_k$, but is otherwise just like the original case (c) Morton, and a version of Lester,

Lester', who is physically identical to Morton' except that Lester' lacks $a_j b_k$. Now it's the same story as in the preceding bundled-process cases, with the same violation of chemical supervenience.

The upshot is that there is no spatiotemporal structure that a conscious experience can have, consistent with the claims (1) that the experience is something purely physical, (2) that the experience has its phenomenal character absolutely (as opposed to relativistically), and, (3) that the phenomenological facts naturally supervene on the chemical facts.

5. Taking stock

Before expanding the scope the foregoing arguments, let's pause to get clear about what, exactly, has gone wrong for the Identity Theory.

Prior to any consideration of the nature of time and space, the identity theorist was working with a picture of individual conscious experiences as complexes of contemporaneous physical events (e.g., neuron-firings) distributed across a spatially extended network (e.g., a brain). In this picture, a conscious experience is a simultaneity-plane, or perhaps a simultaneity-slab, of neural activity (see Fig. 7).

Suppose that the left half of the slab in Fig. 7 contains the events that ostensibly constitute the phenomenal redness of my experience of the baby, and the right half contains the events that ostensibly constitute

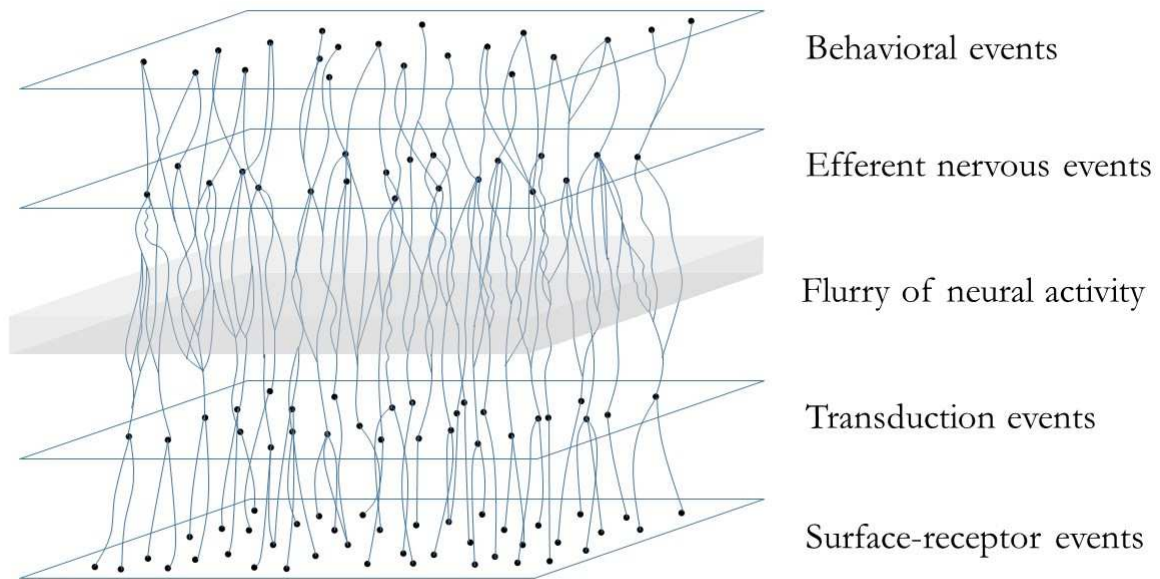


Figure 7: What the Identity Theorist Wanted

the experience's phenomenal loudness. Well, since these events occur in relativistic spacetime, it's possible to represent them (correctly) in such a way that all the events in the left half of the slab occur before some of the events in the right half (and also possible to represent them so that all the events in the right half occur before some of those in the left half).

But if, like the identity theorist, we identify my visual experience (the instantiation of phenomenal redness) with the event-structure in the left half of the slab, and my auditory experience (the instantiation of phenomenal loudness) with the event-structure in the right half, we have to say, wrongly, that my experience is correctly describable as failing to combine visual and auditory phenomenology in the way it does.²¹

In a Newtonian universe, we could identify a conscious experience with a collection of physical processes, without violating any plausible principle of natural supervenience. That's because in a Newtonian universe, spatially separated events and processes corresponding to distinct qualia can occur absolutely simultaneously (or absolutely concurrently). This gives the identity theorist considerable freedom in his choice of physical structures with which to identify the phenomenal features of our mental lives.

²¹We get the same result on any other vertical partitioning of the slab, including Swiss-cheese partitionings in which, e.g., the partition containing the visual processing comprises multiple disconnected "core samples" of the slab.

The problem for the identity theorist is that we don't live in a Newtonian universe. Instead, we live in a universe in which distal simultaneity is a feature of the maps we use to represent spacetime, rather than a natural cleavage-plane in spacetime itself.

In effect, the Identity Theory assumes that spacetime has more structure than it has, according to relativistic physics. That is what makes the theory vulnerable to the chronometric argument.

6. Extending the argument

We now consider a version of the chronometric argument that applies to all solutions to the mind-body problem that entail that the phenomenological facts metaphysically supervene on the physical facts. Call this the *generalized chronometric argument*:

- P1 One sometimes has a conscious experience that instantiates its various qualia (visual, auditory, etc.) absolutely simultaneously.
- P2 If a conscious experience instantiates its qualia absolutely simultaneously, then the experience is confined to a single point of spacetime.
- P3 Some conscious experiences are confined to single points of spacetime. (from P1, P2)
- Q4 If a conscious experience is confined to a single point of spacetime, then the phenomenal facts don't metaphysically supervene on the physical facts.
- Q5 The phenomenal facts don't metaphysically supervene on the physical facts. (from P3, P4)

We've already seen the arguments for P1 and P2. What about Q4?

Suppose E is a phenomenal point-event (a conscious experience that is confined to a single point of spacetime). The same arguments by which we have already established that E can't be a physical point-event also show that no physical point-event is such that its occurrence metaphysically entails the occurrence of E . (Otherwise, ham sandwiches could have the same kinds of experiences that we have.)

So the phenomenal contents (if any) of a spacetime point don't supervise on that point's physical contents.

But if the phenomenal contents of a spacetime point don't metaphysically supervene on its physical contents, yet do metaphysically supervene on the physical contents of some points of spacetime, then how things are at one point in spacetime can metaphysically necessitate how things are at other points in spacetime.

For example, take a sequence of physical point-events $\epsilon_1, \epsilon_2, \epsilon_3, \dots, \epsilon_n$, occurring at spacetime points $\langle x_1, y_1, z_1, t_1 \rangle, \langle x_2, y_2, z_2, t_2 \rangle, \langle x_3, y_3, z_3, t_3 \rangle, \dots, \langle x_n, y_n, z_n, t_n \rangle$. If the existence of this event-sequence metaphysically entails the occurrence of a phenomenal point-event E at $\langle x_n, y_n, z_n, t_n \rangle$, despite the fact that the occurrence of ϵ_n alone doesn't entail the occurrence of E , then "At $\langle x_n, y_n, z_n, t_n \rangle$, ϵ_n is occurring but E isn't" metaphysically entails: " ϵ_n wasn't preceded by the sequence $\epsilon_1, \epsilon_2, \epsilon_3, \dots, \epsilon_{n-1}$."

Likewise, if the physical event-sequence $\epsilon_1, \epsilon_2, \epsilon_3, \dots \epsilon_n$ metaphysically entails the occurrence of E at $\langle x_1, y_1, z_1, t_1 \rangle$, despite the fact that no physical event occurring at $\langle x_1, y_1, z_1, t_1 \rangle$ necessitates the occurrence of E , then “At $\langle x_1, y_1, z_1, t_1 \rangle$, ϵ_1 is occurring but E isn’t” metaphysically entails: “ ϵ_1 will not be followed by the sequence $\epsilon_2, \epsilon_3, \dots \epsilon_n$.” (Similar implications follow from supposing that E is a phenomenal point-event occurring before t_1 , after t_n , or between t_1 and t_n .)

However, it is false that how things are at one point in spacetime can metaphysically entail how things are at other points in spacetime. How things are at one point in spacetime might *empirically* entail how things are at other points in spacetime; e.g., it might be a consequence of the laws of nature that if a photon gets emitted at one point in spacetime, a photon gets absorbed at another point in spacetime. But metaphysically, the contents of separate points or regions of spacetime are all “loose and separate”: there is no outright contradiction in the idea of an emission for which there is no corresponding absorption.

From all this, it follows that if a conscious experience is confined to a single point of spacetime, then the phenomenal facts don’t metaphysically supervene on the physical facts. That gives us Q4, and completes the argument against metaphysical psychophysical supervenience.

7. The state of play

The considerations brought forward in this paper may appear to rule-out any theory of consciousness according to which the phenomenological facts metaphysically supervene on the physical facts, and any theory that assigns spatial or temporal extent to our most basic conscious experiences.

But this isn't quite right. All that we've really ruled-out are theories *that locate conscious experience in time* and either affirm that the phenomenological metaphysically supervenes on the physical or deny that our most basic experiences are point-events. The qualification is necessary, since if conscious experience doesn't occur in time, P1 is false, and the chronometric arguments collapse.

So, which theories of consciousness have we *not* ruled-out?

First, and most obviously, there is Russellian dualism, according to which our most basic experiences are non-physical point-events occurring in the same spacetime as physical events. This is consistent with the view that the phenomenological facts supervene naturally (though not metaphysically) on the physical facts.

We also haven't ruled-out a non-standard physicalism(?) which denies conscious experience any location in spacetime, but insists that the phenomenological facts metaphysically supervene on the physical facts in spite of that. For example, nothing said so far rules-out a view according

to which the following is metaphysically sufficient for the existence of a non-spatiotemporally-located audiovisual experience: one physical process happens in your auditory cortex and another in your visual cortex, and these processes jointly contribute to your physical behavior in a certain way. As long as the audiovisual experience does not exist in time (or spacetime), its phenomenal character doesn't constrain the spatiotemporal structure of its neural correlate, and the arguments above get no traction.

Nor have we ruled-out a non-standard dualism which denies consciousness location in spacetime, but holds that the phenomenological facts supervene naturally (though not metaphysically) on the physical facts.

Finally, we haven't ruled-out an idealism that denies conscious experience any location in spacetime, and takes this as an opportunity to construe consciousness as the subvenient base of spacetime and its contents. While this may seem to be the least plausible theory to survive the arguments, it is, in some ways, the most intriguing. It offers all the benefits of monism, with the bonus of a reductive account of time and space.

To take the idealist suggestion seriously, we have to take seriously the idea that consciousness might not exist in time. If that's a bridge too far for some, it's at least *less* far than we might have thought, given the arguments above. If those arguments show anything, it's that we

can't fit consciousness into spacetime except on terms that confute all mainstream solutions to the mind-body problem. Any new solution is therefore bound to look very different from the old ones, especially in how it fits together what we know about space, time, and consciousness. For now, it seems that the best policy is to maintain an open mind.²²

8. Conclusion

Physics explains everything physical, and at first that seems to favor a physicalist approach to consciousness. But what the large print giveth, the small print taketh away: physical explanations occur against the backdrop of a certain conception of time and space, and this conception turns out to preclude any completely naturalistic account of consciousness. Ironically, the most persuasive argument against twentieth century physicalism might be: twentieth century physics.

Michael W. Pelczar
National University of Singapore
phimwp@nus.edu.sg

²²Henri Poincaré suggests that a radical idealism holds the best prospects for reconciling the absoluteness of conscious experience with the non-absoluteness of physical duration and simultaneity; see (Poincaré, 1898, 223-34). [Redacted] is a recent attempt to develop Poincaré's suggestion into a proper theory.

References

- Callender, Craig. 2008. Finding “real” time in quantum mechanics. *Pages 50–72 of: Craig, William Lane, & Smith, Quentin (eds), Einstein, Relativity, and Absolute Simultaneity*. New York: Routledge.
- Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chalmers, David. 2000. What is a neural correlated of consciousness? *Pages 17–39 of: Metzinger, Thomas (ed), Neural Correlates of Consciousness: Empirical and Conceptual Questions*. Cambridge: MIT Press.
- Dainton, Barry. 2006. *Stream of Consciousness: Unity and Continuity in Conscious Experience*. New York: Routledge.
- Eddington, A.S. 1928. *The Nature of the Physical World*. New York: Macmillan.
- Hameroff, Stuart, & Penrose, Roger. 1996. Conscious events as orchestrated spacetime selections. *Journal of Consciousness Studies*, 3(1), 36–53.
- Jackson, Frank. 1982. Epiphenomenal qualia. *The Philosophical Quarterly*, 32(127), 23–36.

- James, William. 1907. *Pragmatism: A New Name for some Old Ways of Thinking*. New York: Longmans, Green, & Co.
- Kirk, Robert. 1979. From physical explicability to full-blooded materialism. *Philosophical Quarterly*, **29**(116), 229–237.
- Lee, Geoffrey. 2007. Consciousness in a space-time world. *Philosophical Perspectives*, **21**, 341–374.
- Maudlin, Tim. 2008. Non-local correlations in quantum theory: how the trick might be done. *Pages 156–179 of: Craig, William Lane, & Smith, Quentin (eds), Einstein, Relativity, and Absolute Simultaneity*. New York: Routledge.
- Mills, Eugene. 1996. Interactionism and overdetermination. *American Philosophical Quarterly*, **33**(1), 105–117.
- Poincaré, Henri. 1898. La mesure du temps. *Revue de Métaphysique et de Morale*, **6**, 1–13.
- Russell, Bertrand. 1926. The world of physics and the world of sense. *Pages 106–134 of: Our Knowledge of the External World as a Field for Scientific Method in Philosophy*, 2nd edn. London: George Allen & Unwin.
- Simon, Jonathan. Forthcoming. Vagueness and zombies: why ‘phenomenally conscious’ has no borderline cases. *Philosophical Studies*.

Strawson, Galen. 2006. Realistic monism: why physicalism entails panpsychism. *Pages 3–29 of:* Freeman, Anthony (ed), *Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism?* Exeter: Imprint Academic.

Tumulka, Roderich. 2006. A relativistic version of the Ghirardi-Rimini-Weber model. *Journal of Statistical Physics*, **125**(4), 821–840.