

(This is a slightly revised version of my paper of the same title, which was published in Hawthorne (ed.) *Philosophical Perspectives* 2007, Blackwell)

Consciousness in a Space-Time Worldⁱ
Geoffrey Lee

In “The Analysis of Matter”, Russell remarked that traditional Cartesian dualism is in tension with the space-time view of the worldⁱⁱ. The dualist holds that there exist immaterial souls that exist in time but not in space. But in Relativity, there is not a sufficiently clear distinction between space and time for this to be possible. You can’t be just in time, and you can’t be just in space, since on this view there’s only one thing: space-time. Space-time is only artificially divisible into separate dimensions, and so you’re either completely in it, and therefore in space *and* time, or you’re not in it at allⁱⁱⁱ.

This is just one example of how Relativity can be relevant to debates about consciousness. In my view, there are other ways in which time is often tacitly appealed to in conventional pictures of what consciousness is like, generating assumptions that may turn out to be incompatible with the space-time perspective. One goal of this paper is to articulate these assumptions and investigate how to resolve any conflict between them and Relativity.

But first let me make a few remarks about Russell’s argument. It provides a simple and strong case for thinking that mental events occur in time *and* space^{iv}. Is this tantamount to giving up the game, and conceding that physicalism is true? I think that is unclear, involving as it does the notoriously unresolvable question of what it takes to be physical. Rather than pursue that question, I would note that even if physicalism is conceded on these grounds, much of what is at stake in contemporary discussions of the mind-body problem is still at stake. For much of that debate has been about whether some mental events or mental substances are ontologically fundamental^v, in the sense that their existence is not constituted by the existence of non-mental events or substances, rather than about whether they fall in a different *category* from physical events or substances. Thus even if it is conceded that there is only one basic category of being (call it the “physical”, if you like), views analogous to Substance Dualism and Property Dualism could be true. There could exist soul-like mental objects whose basic parts are not ordinary physical objects like fundamental particles, although like fundamental particles they exist in time *and* space. And the existence of mental events, despite their being spatio-temporal, might not be constituted by the existence of non-mental, microphysical events, as the conventional physicalist would maintain. The lesson of the Russellian argument might be that this is how we should view the debate.

In what other ways might Relativity impact on our understanding of consciousness? Russell’s argument tells us that conscious events must be spatio-temporal, but it doesn’t tell us *how* these events fit into space-time. There are hard questions about what kind of metaphysics a temporally extended stream might have in a Relativistic world. For example, it is natural to think of perceptual consciousness as a continuous stream that has

a certain qualitative character^{vi} at each moment that it is being enjoyed. Just as it is natural to regard a material object as having a continuous series of instantaneous overall physical states (mass, shape, etc.), so it is natural to regard consciousness as having a continuous series of overall phenomenal characters: to put it somewhat metaphorically, it is like a densely packed series of infinitely thin phenomenal slices. This picture comes into obvious tension with relativity, since in a Relativistic world, there are no moments of time, or at least, they only exist in a thin frame-relative sense. One question is how a continuous stream of phenomenology could nonetheless exist in space-time, in a way that is consistent with this fact (it is far from obvious that the stream actually is continuous, but we would still like to know what continuity would consist in, were it to obtain). Much of the paper is concerned with this question.

Another possible area of impact is the debate over the unity of consciousness. Many philosophers agree that there is a relation of “co-consciousness” that different parts of a person’s experience can stand in. It is in virtue of standing in such a relation that different experiences form a unified conscious field, and it is this relation that may fail to obtain between parts of the experience of a split-brain patient (if it is true that such a patient sometimes has two distinct conscious fields). It is standard in this debate to distinguish synchronic from diachronic unity. The kind of unity that holds at a time is often thought to be quite different from any relation of unity that might hold over time. For example, it has been argued by some that synchronic unity is transitive^{vii}. That is, it is supposedly not possible for two distinct unified conscious fields to share a proper part in common at any time, without sharing all their parts at that time – for example, two distinct people cannot share numerically the same headache. On the other hand, it appears that experiences that are sufficiently separated in time are not co-conscious, so if co-consciousness holds over time at all, it holds only over short periods of time. But if this is right then diachronic co-consciousness (if there is such a thing) cannot hold transitively: if it did, we could probably join distant mental events by chains of neighboring co-conscious events, and conclude that the distant mental events are themselves co-conscious.

In a Relativistic world, there is no objective relation of simultaneity that can hold between spatio-temporally distinct events. At best, there is only a relation of “space-like separation” between such events, which does not hold transitively. This raises the question of whether, in this context, we can reconstruct the distinction between synchronic and diachronic unity of experience that some have believed in. One might suspect that like the physical events described by Relativity, conscious mental events also cannot stand in an objective relation of simultaneity; if so, we may be forced to think about unity in a different way.

I will be focusing more on the metaphysics of the extended stream than on co-consciousness, although that issue will be relevant in what follows. Also, rather than approaching these matters directly, the paper will be organized around attempts to resolve an interesting puzzle that arises when you try to understand how experiences exist in a space-time continuum. The attempts at resolution involve different views on the spatio-temporal metaphysics of the stream, so assessment of these views will be crucial to resolving the puzzle.

Very briefly, the puzzle is this. It is well-known that in a Relativistic world, many apparently non-relational properties of objects, like their shape, are instantiated only relative to a frame of reference. But for reasons that I will explain, it would be intolerable if the phenomenology of experiences was similarly frame-relative. However, there are also arguments that can be given that lead to the strange conclusion that phenomenology *is* frame-relative. The puzzle is to explain how consciousness relates to the physical world in a way that avoids this conclusion.

A certain confused tendency in our thinking about temporal experience will be particularly salient in discussing these matters. When thinking about the phenomenology of temporal experience, there is a temptation to think that the temporal *properties* of experience – how experience plays out over time - play a special direct role in fixing what the experience feels like to its subject. One interesting example of this kind of thinking is the dubious “Cinematic” view of temporal perception^{viii} – the view that experience presents temporal phenomena in virtue of its own temporal layout, in a way analogous to the representation of time of a cinema screen (for example, on this view, a perception of a moving object consists in a series of perceptions as of the object being at different locations). If this view were right, there would be a direct link between the temporal properties of perception, and its temporal content - which is presumably in turn intimately related to the phenomenology of the experience.

Even if we reject the Cinematic view, as we probably should, it might still seem correct that how the stream is laid out in time has a special role to play in fixing what it’s like to undergo it. For example, think of the difference between the temporal properties of experience and its spatial properties. If Russell’s argument is right, experiences have both spatial and temporal properties. But the temporal properties are intuitively much more relevant to how experience strikes us from a first person perspective than the spatial properties^{ix}. For example, this is one reason why it seems plausible that experiences exist in time, but not in space. Allowing one’s thinking to be guided by such ideas is a significant source of confusion that will be especially relevant to thinking about experience in a Relativistic context.

The structure of the paper is as follows. In the first section, I will give a very brief overview of those aspects of Relativity that are relevant to the discussion, and explain one version of the puzzle I’d like to discuss. This version of the puzzle has certain problems, which will be avoided by giving a different formulation. I will then discuss what I take to be the most important strategies for resolving the puzzle.

A Puzzle about Consciousness and Relativity

I begin with an extremely brief run-down of some basic ideas from Relativity that will be relevant to the discussion (some may want to skip the next few paragraphs). I will focus here and throughout the paper on Special Relativity (I believe the most important points I will make all have analogues in General Relativity, although I won’t argue for that here). Now, it is well known that according to Relativity, it cannot be determinately the case

that two distinct events (i.e. two events that do not happen at the same time and place) are simultaneous, or more generally that they are separated by a certain amount of time. What exactly does this mean?

A useful analogy is with different ways of describing the layout of objects in space. I might draw a two-dimensional map of an island on a Cartesian grid, so that every place is given an x and y coordinate. Obviously, it is simply an arbitrary matter where to place the x and y axes – some choices may be more convenient than others, but there are no axes objectively imprinted on the island itself.




Fig.1: A map of an island, with two different Cartesian coordinate schemes marked. Obviously, two places may have the same x or y coordinate in one scheme but not in the other.

One surprising idea in Relativity is that if we are instead mapping the playing out of events over time in 3-D space using a 4-D coordinate system, our choice of coordinates is similarly arbitrary, including even our choice of the direction in which we place our temporal axis. In particular, just as different places on our island might have the same y -coordinate in one island-map, and different y -coordinates in another (see fig. 1), two distinct events in space and time may have the same time-coordinate in one space-time map, and different time-coordinates in another equally optimal map (see fig. 2). The temporal relation between events is therefore not objective, in the sense that it is relative to an arbitrary choice between equally optimal coordinate maps. Just as there are no objective x and y axes imprinted on our island, the universe does not come objectively marked out with a distinction between space and time: instead, there is just one thing, called “space-time”, on which we can impose different but equally optimal coordinate

maps, each of which gives a different rendering of the 4-D object into space and time dimensions.




Fig. 2: A space-time diagram with two alternative coordinate systems shown. The points in the upper and lower diagonal quadrants are time-like separated from the center, and those in the right and left quadrants are space-like separated from the center. Events e1 and e2 are simultaneous in one frame but not in the other.

This said, the failure of space and time to be distinct is subject to an important qualification. In the case of mapping an island, you can choose to put your y-axis in any direction you like, but the same is not true for a choice of time-axis when mapping events in space and time. Although in relativistic space-time there is no privileged way of distinguishing space and time, there is nonetheless an objective distinction between “time-like” and “space-like” directions from an event in space-time. We want a choice of coordinate system to reflect this fact. Hence, if our coordinate scheme is to reflect the objective structure of space-time, its temporal axis must be placed in a time-like direction (this constraint still leaves an arbitrary choice between many different frames, however).

What distinguishes time-like and space-like directions? Two events are time-like separated if it is possible for a causal signal traveling slower than the speed of light to move from one to the other. Conversely, two events are space-like separated if no such connection is possible^x. Time-like separated events can be thought of as having an objective temporal order – in terms of choosing coordinate frames, this means that in every choice of frame they will be assigned the same order in time. Conversely, space-like separated events have no objective temporal order – in some frames they are given one temporal order, in others another, and they will be simultaneous in exactly one frame. This fact about space-like separated events will be important in what follows.

Not only are the temporal relations between events frame-relative, so too are many other features of the world, such as macroscopic properties of ordinary objects like shape and mass. This means that what we might call the “classical” picture of how objects instantiate physical properties is at best incomplete in the light of Relativity. On the classical picture, we think of a material object as having a sequence of momentary physical states during a certain period of time: at each instant, the object has a certain shape, mass, and so forth. But in a space-time world, we can only talk about how an object is at a time, relative to a coordinate frame. As a result, an object has a different sequence of momentary states for every frame that can be used to describe it, moments being frame-relative objects.

A helpful way to understand this is by thinking of the object as having at each time a momentary part, or “slice”, which is the real bearer of these properties. I will do this, assuming that what I say can be recast in terms acceptable to those who dislike this metaphysical gloss. In these terms, the important point is that from the space-time perspective there is no privileged way of slicing an object into instantaneous parts. An instantaneous slice is simply a set of simultaneous events belonging to the life of the object but, since simultaneity is frame-relative, we will get a different slicing of the object into instantaneous parts in every frame. And just as if we cut a sausage into parallel slices using different orientations for our sausage-cleaver we will get different sequences of slices with different properties (for example, the shape of each slice depends on the orientation of the cleaver), we will get different sequences of temporal slices with different properties, depending on which frame-relative slicing of an object we choose^{xi}. In particular, an object will have a different sequence of properties like mass and shape in different frames.

This is a strange result to end up with, but once you get used to the new conceptual framework, it is one that is easy to live with. However, macroscopic objects sometimes have more unusual macroscopic properties than shape and mass. In particular, let us now consider a very special kind of macroscopic object, a human being. Human beings have conscious mental properties like *feeling a pain*, and *enjoying a visual experience*. More generally, we seem to be subject to a continuously changing stream of consciousness: intuitively, at every moment of consciousness, there is something that it feels like to be conscious at that very moment. This is what we might call the “classical” picture of consciousness, since it is the analogue for consciousness of the classical picture of how material objects instantiate properties like mass and shape.

Now, it would obviously be very strange if, like shape and mass, the phenomenal properties of our stream of consciousness turned out to be also frame relative. But as I will now explain, there are a couple of lines of reasoning that suggest that this is indeed the case.

How Phenomenology Can be Argued to be Frame-Relative

Here's the first line of reasoning. Suppose at a certain precise moment of time you begin to feel a pain, and at roughly the same time you begin to have a visual experience of a blue patch. Let us make the assumption that the event of the pain starting, and the visual experience starting, are events located in space-time (as they must be if Russell's argument is correct), and also that they have space-like separation in the sense outlined above. Then according to Relativity, the onset of the pain and the visual experience will have different temporal orders in different frames, and will be precisely simultaneous in only one frame. But this is really quite strange! For it's plausible that *whether or not they happen at the same time makes a difference to what it's like for the subject*. Or to put the point slightly differently, if the pain starts before the visual experience, it appears that this will feel different from the case where the visual experience starts first. It looks like we are forced to say that the phenomenal character of these experiences is a frame relative phenomenon!

It's worth pausing for a moment to clarify just how different the timing of these events will be in different frames. One well-known effect of Relativity is that events can have different temporal *orders* in different frames, and, as the last paragraph suggests, this may apply as much to conscious mental events as other kinds of events, apparently implying that the character of experience is frame-relative. Another effect is *time-dilation*. Consider a conscious subject who is speeding away from the earth in a spaceship, close to the speed of light. If we describe this subject using a coordinate frame centered on Earth, then events on the spaceship will be extremely stretched out in time. In particular, a sequence of neural events that takes only a few seconds according to a clock on the spaceship may be stretched out over many centuries relative to clocks on earth. There is at least some pull to the idea that how long our experiences last makes a difference to what it feels like to undergo them: for example, normally, a pain that lasts an hour feels different from a pain that lasts a second. But if we were to apply that idea to the present case, we would have the conclusion that the space traveler's experiences are qualitatively quite different from the point of view of an Earthly frame than from a frame centered on his ship!

We can combine these points about ordering and time-dilation: there exist frames where the visual experience happens arbitrarily long *before* the arm pain, or arbitrarily long *after* it. (But let us also note that, even in a frame where there is a very long delay between the pain and the visual experience, since both are certainly parts of temporally extended pains and perceptions, the pain will probably still be going on millions of years later when the visual experience starts! This is because temporal order is only frame-relative for space-like separated events and two temporally extended and temporally overlapping events will inevitably both have parts that are time-like rather than space-like separated from one another, and whose order is therefore preserved across frames.)

To sum up: the natural extension of the classical picture of shape and mass into a Relativistic framework involves accepting that an object can have a different sequence of e.g. shapes in different frames: shape is a frame relative property. But attempting a similar extension in the case of the stream of consciousness creates a puzzle: it would appear to have the result that you are subject to different streams of consciousness in

different frames, so that the phenomenal character of experience, like shape and mass, turns out to be a frame-relative phenomenon. The question is how best to avoid this odd result.

There is a variation on this puzzle that can be framed in terms of the “co-consciousness” relation that I mentioned earlier. Suppose you believe that there exists a special unity relation between experiences that only holds strictly at a time (as would probably be required, for example, for it to be a *transitive* relation (see earlier discussion)). Suppose you also think that whether or not this relation holds between two experiences – whether they are unified - makes a difference to what it’s like for the subject of the experiences. Finally, suppose that some of the experiences the relation holds between are located in different regions of space-time that are space-like separated. Now the problem should be clear. Since simultaneity only holds between space-like separated events relative to a frame, such mental events could at best only be unified relative to a frame. But if unity is a relation that matters to phenomenal character, then we have the consequence that the overall phenomenology of experience is frame-relative.

My main focus will not be on this version of the problem. I mention it to illustrate how the unity of consciousness could potentially come into play in this context. Also, both this and the first line of reasoning are subject to a couple of *prima facie* problems, which will motivate introducing yet another version of the puzzle.

First, the most compelling case for frame-relative phenomenology has so far come from the idea that experiences might have different temporal orderings in different frames (and be simultaneous in only one frame). But to get this conclusion, we needed to make the assumption that the mental events in question are *space-like separated*. This may not be an assumption that would be accepted by all, particularly not by some physicalists. One reason is that the physical realizers of simultaneous parts of an experience, like the start of the pain and the start of the visual experience, may always spatially overlap. For example, this will be the case on an extremely holistic view of mental realization, according to which the most precise spatial location we can give to a mental event is the region of the whole brain. Overlap might also be motivated by thinking about unity – perhaps two mental events could not be co-conscious if located in completely distinct space-time regions. Anyway, if this *is* the case then mental events in the same stream may not ever, in any clear sense, have different temporal relations in different frames.

Second, the argument works by appealing to the idea that experiences have parts that stand in temporal relations. But it’s not in general an easy business saying how, if at all, experiences are composed of *parts* – what are the parts of my current visual experience, for example? It’s tempting to think it has parts corresponding to the parts of the space I’m perceiving, but on reflection that is a dubious assumption, to say the least. So even if we think that talk of the parts of experience is *sometimes* cogent, we should accept that if experience has parts, it does not in general do so in the same way that a spatially extended object does. This suggests that if we want a version of the puzzle that applies to all of our experience, and not just to special cases where arguably experience

has spatially separated parts, it would be best to set it up avoiding appeal to the idea of a part of experience.

I think this can be done – there is another line of reasoning to the same conclusion that does not make these assumptions.

This puzzle arises from trying to understand the idea of a continuous stream of experience in a space-time context. But before getting to that, it is better to start further back, with the simpler question of what it would be for a continuous phenomenal stream to exist in a *Newtonian* world.

One gloss on this is as follows. Take a subject who is conscious during a certain interval of time. If, for each of a continuous series of moments during the interval, there is a fact about what it is like for the subject at that precise moment, then the subject has a continuous phenomenal stream. Or to put it in terms of an ontology of temporal parts as bearers of phenomenal properties: a continuous stream consists in each of a continuous series of momentary parts of the subject bearing phenomenal properties.

There *are* other ways of understanding the idea of a continuous Newtonian stream, and there are objections to this “stream of momentary states” view that could be raised. In particular, note that this understanding of a continuous phenomenal stream requires us to believe in strictly momentary phenomenal events, which some may find objectionable. We should also be clear that there are alternatives to the idea that the stream is continuous. For example, experience might come in a series of discrete bursts. I’ll return to these matters later on. For now, let us take this understanding of continuity for granted, and ask how it might be extended to capture the idea of a continuous stream in the space-time domain.

One natural strategy would be the following. On the Newtonian view just outlined, each momentary temporal part of a conscious subject is a bearer of phenomenal properties. But in a space-time world, there are only momentary temporal parts relative to a frame of reference; there are many ways of dividing a subject into a continuous series of temporal parts, not just one. One view is that every one of these continuous series is such that each momentary part of it is the bearer of phenomenal properties. Rather than there existing a *single* continuous series of phenomenal states belonging to a subject, we have *infinitely many* - a different phenomenal series for every way of dividing a subject into a continuous series of temporal parts.

It is this view that generates the second puzzle. If we assume the view gives a correct description of consciousness, the obvious question then is – what is the relationship between these different phenomenal series? If they turn out *not* to be the same, then it starts looking again like we have different streams of consciousness in different frames. More specifically, it seems to be true that if two subjects are subject to two qualitatively distinct temporal series of phenomenal states, then what’s it’s like for them must be

different. But then if an ordinary subject has different streams of phenomenal states in different frames, then what it's like for her seems to be different in different frames.

Here's a version of the problem set out as an argument in more explicit form. In this version, I assume that there is a determination relation between the physical and phenomenal levels, each phenomenal stream being generated by a physical stream. This may be a causal determination relation, as on a Property Dualist view, or a stronger relation of identity or constitution, as on a physicalist view.

- (1) A stream of neural activity has a different instant-by-instant sequence of physical states in different frames. (Relativistic Premise)
- (2) An instant-by-instant sequence of physical states of the brain of a conscious subject determines, step by step, an instant-by-instant sequence of phenomenal states. (Premise)
- (3) If we consider any two frame-relative sequences of neural states, they will each determine instant-by-instant phenomenal sequences. ((1) + (2)).
- (4) These sequences are different^{xii}. (Premise).
- (5) Therefore, the subject's experience has a different instant-by-instant phenomenal sequencing in different frames.
- (6) Therefore, the overall phenomenology of the subject's experience over time is different in different frames.

This, then, is the second version of the puzzle. Note that this version of the puzzle in a way embeds the first version(s). For if we make the assumption that our subject's stream contains phenomenal events that are space-like separated, then in some frames one event will occur in an earlier temporal part than the other, and vice-versa. Note also though, that even if we don't make the assumption of space-like separation, for example because we take a holistic view of mental realization, we can still accept that there are different series of phenomenal states in different frames, because there will be different series of total brain states – and that is enough to generate a puzzle. Thus this version of the puzzle succeeds in avoiding some of the assumptions of the first one.

I would also note that the relation between the different series of phenomenal states is far from obvious. It is wrong to assume that each series will contain the same set of momentary phenomenal events, perhaps ordered differently in time. (In fact, that could only be right if each phenomenal event occurred at a single space-time point, a point I will return to below). It seems that the relation between the different phenomenal series is much more opaque than that – each series might contain a quite different, although presumably related, set of momentary phenomenal events. This is a point that I will discuss again in more detail below.

Now I will consider how to respond to the two versions of the puzzle. The first kind of response is to simply accept the conclusion that phenomenology is frame-relative. I will try to persuade you that we cannot accept this idea. The next two responses both avoid the puzzle by arguing that experience is after all objectively organized into series of phenomenal events with a total order in time. They therefore adopt a different

metaphysical picture of what an extended space-time stream would be than the one that generated the puzzle. The first response of this type, and this is a response to the puzzle that I have often encountered, is to deny that all frames of reference have equal status in determining experience. In particular, it is tempting to believe that the subject's own frame of reference has a privileged role to play. I think this is wrong, as I will argue. The second response of this type is to claim that despite there being a different sequence of neural states in each frame, there is the *same* sequence of phenomenal states. I'll argue that whether this response is plausible, depends on certain features of the stream: for example, it may partly depend on whether the stream is discrete or continuous. Finally, I will explain what I think is a better approach to resolving the puzzle: we can try to explain how, even if the stream has different temporal specifications in different frames, that need not lead to the conclusion that phenomenology is frame relative. This approach itself raises some interesting puzzles about the stream of consciousness with which I'll end the paper.

Why You Shouldn't Think that Phenomenology is Frame-Relative

It may appear that claiming phenomenology to be frame-relative has the same problems as certain other similar-sounding views. For example, it might be that any arbitrary set of human temporal parts together form a material object. This view could have the consequence that there are infinitely many material objects co-located with you, each enjoying experiences just like your present ones^{xiii}. Or for another example, on a "no collapse" interpretation of quantum mechanics, on observing the outcome of an apparently indeterministic measurement, your brain enters a superposition of different neural states, each corresponding to the observation of one possible outcome of the measurement. According to some^{xiv}, this superposed state determines an infinite multiplicity of experiences, each an experience of one possible outcome.

However, what I take to be the most persuasive objection to frame-relative phenomenology does not also apply to these other views. The objection is that phenomenal facts are objective and so they cannot be frame-relative. I'll call this the "objectivity objection". The other important objection here I'll call the "parsimony objection". Postulating many phenomenally distinct streams of consciousness where we thought there was only one is to postulate more states of affairs than we normally think exist – the view is ontologically inflationary. Analogues of the parsimony objection have force against the views described above: in the present case, I think the objection should be seen as at best supplementary to the objectivity objection.

Baldly stated, the objectivity objection is that phenomenology cannot be frame-relative, because it is objective in a way that is incompatible with frame-relativity. But stated so plainly, the point can hardly be expected to have much force – why can't an opponent just deny that phenomenology is objective? The point does have force, though, once we make the effort to elucidate what sense of "objectivity" is in question here.

"Objective" is commonly used in philosophy to contrast with "subjective" or "mind-dependent". But this can hardly be the sense of "objective" that matters here. First,

frame-relative states of affairs *are*, in general, objective in this sense. For example, it will be a quite mind-independent matter that, relative to a certain frame of reference, an object has a certain velocity at a certain time. Second, phenomenal states of affairs are clearly *not* objective in this sense – in fact, they are paradigms of the mind-dependent!

Another thing that “objective” could mean here just *is* “frame-independent”. But of course, the problem with that is that the argument, so understood, would just beg the question. What is the sense of the term that matters here, then? Frame-relative facts are not objective in the sense that they are facts that depend not just on how things are objectively laid out but also on which arbitrary scheme of description we are using – the schemes of description in question being arbitrary in a way analogous to the way that a choice between describing the world in either French or German is arbitrary^{xv}. However, this explanation appeals to the idea of being “relative to an arbitrary scheme of description”, which is itself quite unclear. Let me therefore try to explain the sense of objectivity that matters here in another way – by appealing to an example where the kind of objectivity in question also doesn’t apply, but where it is also easier to understand what this amounts to.

The example was already mentioned earlier. There exist different ways of describing a 3-D Newtonian space at a time using an (x,y,z) coordinate map, because there are different choices for where to place the origin of the map, and how to orient its axes. Suppose that someone working in such a Newtonian framework claimed that which conscious experiences were happening in space at a certain time depended on which choice of (x,y,z) coordinate system we used to describe the space. This would obviously be an insane thing to think. The facts about a phenomenon like consciousness obviously couldn’t depend on an arbitrary choice of (x,y,z) coordinates. Why not? The reason seems to be that the different choices of coordinates are simply different choices of how to describe the same concrete, or objective states of affairs. Since the facts about consciousness are presumably among these objective states of affairs (or supervenient on them), they cannot vary with a choice of map. It’s not that there *aren’t* coordinate-scheme relative facts, like “Place p has coordinate (X,Y,Z) ”, or “Places p and q have the same x and y coordinates”. But it’s obvious that these are only facts of a thin “description dependent” kind, and could not plausibly be taken to include anything as robust as the phenomenal facts.

The important point now is that these different maps of a Newtonian 3-D space are different “frames of reference” in exactly the same sense in which the Relativistic frames are, so facts which depend on them are not objective in the same sense in both cases. This means that if phenomenology can’t vary with a choice of Newtonian coordinate frame, because it has a certain kind of objectivity, then the same thing is true of it’s relation with Relativistic coordinate frames.

Another way to look at this is that Relativistic frames of reference are not essential parts of the theory, but merely dispensable schemes of description. SR and GR can be formulated without reference either to frames of reference or the properties that are relative to them, like Euclidean shape (in a flat space-time), and absolute velocity. This

means that neither these properties nor the frames of reference are part of the fundamental ontology of the theory. It is in this sense that Relativistic frame-relative states of affairs fail to be objective.

To clarify, let me note that strictly speaking, this understanding of Relativity is not forced upon us. There is a deviant way of understanding Relativity that gives frames of reference, and the properties that depend on them, much more ontological import. According to Kit Fine's "many worlds" interpretation of Relativity^{xvi}, different frames of reference are not simply different ways of describing the same world, but rather descriptions of different, but equally actual worlds. On this interpretation, we have *different* objective facts corresponding to each frame-relative description, not the same objective facts – this is therefore an interpretation on which frames of reference and the facts that are relative to them *are* part of our fundamental ontology. Understood this way, frame-relative facts are perfectly objective, and the claim that phenomenology is frame-relative is not vulnerable to the objectivity objection – each of a set of distinct worlds could coherently contain different phenomenology. The many worlds interpretation rather has a status that closely resembles the status of the "many worlds" interpretation of quantum mechanics I mentioned earlier.

But of course, it is much more plausible to interpret Relativistic frames of reference as being like different Newtonian coordinate-systems, than to adopt the kind of heavyweight interpretation suggested by Fine. Given this fact, it would be just as bad to think that phenomenology is frame-relative in the Relativistic case as it is in the Newtonian case. I hope this persuades you that the objectivity objection, correctly understood, has considerable force.

What about the second objection, that believing in frame-relative phenomenology is ontologically unparsimonious? This objection faces the following riposte. If this were a good objection to thinking that phenomenology is frame-relative, wouldn't it also be a good objection to thinking that other states that naively seem non-relative, like shape and velocity, are in fact frame-relative? You may have naively thought that an object has only one series of shapes or velocities, but in Relativity it has many – one relative to each frame of reference. That suggests that Relativity is *in general* guilty of an undesirable lack of parsimony.

However, that obviously wouldn't be a good objection. Why? Provided we interpret Relativity in a non-deviant way, frame-relative facts are not fundamental facts, and therefore there is nothing unparsimonious about a theory that countenances them. To put the point differently: if we adopted a "many worlds" interpretation of Relativity, according to which there is a different world of objective facts corresponding to each frame, then it really would be guilty quite generally of ontological bloatedness. But for precisely this reason, we should surely reject such an interpretation, and with it the idea that frame-relative facts are objective. On the other hand, as I have argued, it is hard to reject the idea that *phenomenal* facts in particular lack this kind of objectivity. But they could only be objective *and* frame-relative if something like a Many Worlds interpretation were correct. The parsimony objection is therefore best seen as supporting

the objectivity objection, by ruling out the Many Worlds interpretation, so that we can legitimately infer the frame-invariance of phenomenal facts from their objectivity. Having noted this, I move to consider the next line of response to the puzzles.

The Subject's Own Frame is Privileged in the Determination of Experience (Rejection of Premise 2)^{xvii}

I would anticipate that many people will be tempted to think: “There’s no real problem here! Experience is determined relative to a *privileged* frame – namely, the rest frame of the subject having the experiences. We should simply ignore all other frames and believe that the only real train of experiences that exists is the one the subject is having relative to her own frame.” For example, you might agree that whether the pain or visual experience starts first does make a difference to what it’s like, but think that only one hypothesis about their timing corresponds to the real facts about what it’s like: the hypothesis that gives their timing in the subject’s own frame.

Although this response is admittedly tempting, I think that it doesn’t stand up to scrutiny.

First, consider how odd it would be to hold an analogous view of other physical properties of an object. Suppose, for example, that someone balked at the idea that an object has a different sequence of shapes in different frames, and held that really the object only has one sequence of shapes – the one determined in the object’s own frame – and that the sequences of physical properties that the object has in other frames shouldn’t really be thought of as determining sequences of shapes. This response is very hard to make sense of, to say the least. The whole point of Relativity is that temporal relations between events are frame-relative, and that no frame is privileged. Therefore any property of an object that depends on such temporal relations is also frame-relative, and cannot be fixed by how things are in some supposedly privileged frame.

There seems to be no reason why this point shouldn’t apply to experience also, especially if we assume that physicalism is true. If experiential properties are just another example of instantaneous physical properties like shape, there is no justification for treating them any differently. To think otherwise would be to misunderstand the way Relativity applies across the board to every physical property^{xviii}.

The physicalist can also back this point up in the following way. Suppose there were two conscious perceivers, who have exactly the same moment-by-moment physical description. The only difference is that this description applies to the first perceiver as he is in his own frame, whereas the second only applies to the second observer as he is in a different frame from his own. According to the view we are mooting, this shared sequence of physical states determines a certain sequence of momentary experiences for the first subject, but does *not* determine the same sequence in the alternative frame for the second subject, despite his being a special kind of physical duplicate. This strikes me as just as weird as denying that the subjects share, relative to different frames, a sequence of overall shapes, despite being physically identical relative to these frames. There is a kind of relativistic supervenience principle that holds in every normal case^{xix}, but is being

challenged for the case of conscious properties, without sufficient explanation for how it could fail there.

Second, let us note that it is not actually determinate what the observer's frame *is*. The observer occupies a 4-D region of space-time, through which there pass infinitely many different world lines. Many of these world lines will naturally determine different partitionings of the subject's brain into momentary slices (Technical explanation : although these world lines are very unlikely to be inertial paths, one can still define a 'proper reference frame' with respect to each one, which is well-defined over a finite region of space. Many such frames will be well-defined in the region occupied by a subject's brain or body.) Each such partitioning may determine a different sequence of experiences, so we may just get the same problem as before, but just limited to a smaller group of experience sequences.

We could try to find a reason for picking out one world line in particular, e.g. "the world-line that tracks the observer's center of mass", on the grounds that it is more natural than some other choices. Now, it is hard to believe that the phenomenology of one's experience could depend on one's center of mass – it may be quite arbitrary where the line around one's body is drawn (or it could be that the *brain* is the relevant object), and it seems that exactly the same neural processing could occur in bodies with quite different centers of mass (could this highly extrinsic fact about the processing really result in different experiences?) More generally, I would speculate that any such suggestion will seem poorly defined and implausible in a similar way.

My next points are not really straightforward objections, but rather related suggestions about what makes this view seem intuitively promising which, when made explicit, can be seen to in fact provide no support for the view. First, some may be simply confusing two different senses of the phrase "the subject's frame of reference". It could (at a stretch) refer to a subject's experiences of the surrounding world, or it could refer to a coordinate frame centered on the subject. But of course, these are quite different things, not obviously connected with each other.

That said, the subject's coordinate frame does, on reflection, have an interesting relation to the subject's conscious point of view. In particular, the spatial-temporal layout of the world around the observer is seen as it is in (roughly) her own coordinate frame. For example, in some frames of reference the objects surrounding you have quite different shapes from those you see them to have, because of length contraction. You see the shapes they have in your own frame, not these strange distortions. So the property of being a veridical perception with a certain content really *is* one that depends on a privileged frame of reference. Whichever frame we are describing the observer from, we will need to switch to describing the world from the observer's *own* frame in order to specify which states of affairs she's perceiving.

This is the basis for another line of intuitive support for the view. For similarly, it might be tempting to think that if we want to know how the observer's own experience feels to her, we need to describe the experience in the observer's own frame. Just like the tables

and chairs around her, the experience is an object which she encounters as it is from *her* point of view. But this is a misleading analogy. An experience is not an object you encounter from a point of view, in the same way that you encounter a table or chair from a point of view. It is primarily a state that you are *in*, not something that you perceive like a material object lying in front of you. And if it is simply another state that you are in, there is no reason not to treat it like other states, which do *not* somehow depend on a privileged frame.

It might be objected that experiences *aren't* just states that we are in, but also states that we can be aware of in introspection. Indeed, everyone accepts that there is at least some sense in which we are aware of our experiences, although it is controversial whether the awareness in question is anything like perception of external objects. But for introspection to play a relevant role here, it would have to present experience exactly as it is in the subject's frame, and the phenomenology of experience would have to be somehow constituted by the way these higher-order states present it. In short, we would need a special strong form of higher-order representation theory of consciousness.

There are well-known objections to such theories that I will not rehearse here. It is also important to note that it is far from obvious that introspection presents experience as it is in one frame. For example, introspection doesn't reveal to us the spatial structure of experience, in the way that perception reveals to us the frame-relative shapes and spatial layout of external objects. So there is no parallel justification for thinking of introspection as revealing frame-relative properties. Further, although introspection may sometimes reveal to you the order of your experiences and other temporal properties of them, it is not obvious that the way it does this means that it presents experiences in one frame rather than another. It is plausible that our faculty of introspection is of sufficiently coarse sensitivity that the order of experiences could only be reliably detected when they are time-like separated, for example. More generally, it could be that the only temporal structure that you are aware of is frame-invariant structure (see the next section for more discussion of this important point). If that is right, then we need not admit that our experiences are presented as they are in one frame, even if we hold a strongly perceptual model of introspection.

Some people may argue that our perceiving the external world as it is in our own frame has a different significance. Naïve realists and intentionalists believe that the phenomenology of experience is given by how we perceive the world to be. So if what we perceive is the world in a privileged frame, then that frame plays a special role in fixing phenomenology, by being part of what we are intentionally related to in having the experience.

In response, I have no objection to one's own frame playing a privileged role in this sense. What I am objecting to is the idea that the timing of one's *experiences*, as they are in one's own frame, plays a privileged role, not the frame-relative timing of the external world. Even for the Naïve Realist, the timing of one's experiences is surely not given by the timing of the *objects* of those experiences, but rather by the timing of neural events that enable one to be perceptually related to those objects (for example, you don't see

something until after light from it hits your eyes). These neural events can be equally well described in any frame we like, and there is no reason to not think the same of the perceptions themselves, even if they involve a relation to one frame in particular.

You would need a special view of the temporal content of experience for these perceptual relations to frame-relative temporal facts to imply anything like the privileged frame view. If you thought temporal perception worked through *resemblance*, as it does on the Cinematic View, and more particularly if you thought that we perceive external temporal relations by first introspecting temporal relations between experiences, then the privileged frame view might follow. For then the perception of frame-relative relations between external events might seem to require us first to introspect certain frame-relative relations between experiences. These frame-relative experiential relations would play a special role in determining the temporal content of experience, and hence its phenomenology. But apart from involving the extremely dubious Cinematic View of temporal perception^{xx} (see endnote for a few remarks on this), this takes us back to a view of introspection as confronting us with experience as it is in our own frame, as discussed above.

To sum up, there is no good reason to think that the timing of your experiences in your own frame has a special role to play in fixing phenomenology, and there are positive reasons for resisting this idea. I now consider the next important response to the puzzle.

Experience has the same Temporal Sequencing Relative to every Frame (Rejection of Premise 4)

The next view to consider is the following. Everyone can agree that a neural system will have different moment-by-moment sequences of physical properties in different frames. On this next view, each frame also involves a moment-by-moment sequence of phenomenal properties but, unlike the case of ordinary physical properties, it is the *same* sequence of phenomenal properties that appear in each frame. Despite the different physical sequences in different frames, the phenomenal sequence is always the same. This is therefore a view on which phenomenology is frame-invariant, and to that extent the view is promising. However, it faces a number of objections.

What does it mean to say that there is the *same* temporal sequence of phenomenal states in each frame? It can't mean that, relative to each frame, the subject is enjoying the same experiences at the same times, because there is no way to identify times across frames. A better attempt at explaining the idea would be as claiming that the same set of momentary experiential states exists for the subject in each frame, and they stand in the same temporal relations in each frame.

One objection to the claim, understood this way, is that neural activity will be massively stretched out in time in some frames, and less so in others. Therefore, at the very least, a sequence of phenomenal states lasts different amounts of time in different frames. This objection could be dealt with by understanding the criterion for identity between phenomenal sequences in a special, weak sense. If two sequences contain the same set of

momentary states standing in the same *relative* temporal relations (like “x is twice as far in time from y than z”), but not necessarily the same absolute temporal relations (like “x is happening one light meter after y) then we can say they are the same sequence, even though one may be stretched in time relative to the other. The claim in play perhaps should be that the phenomenal sequences are “the same” in this weaker sense. (It will have to be the case that absolute temporal relations between conscious events, like “happening one light meter after”, don’t matter to phenomenology, but that’s quite plausible). Even given this clarification, the view is subject to the following objections.

First, this is no reply to the first version of the argument, according to which experiences can have different *orders* in different frames. That is, if it is legitimate to believe in space-like separated mental events that are part of the same stream, their order will be frame-relative, which is inconsistent with the stream being the same in different frames, even in the weaker sense. So the proponent of this view may have to deny that mental events from the same stream can stand in a space-like relation.

Second, since your brain has a different sequence of physical properties in each frame, it seems like it would be a strange metaphysical coincidence if nonetheless the sequence of phenomenal states you have is the same. I assume that there is some kind of determination relation between the physical and phenomenal levels – whether it is a causal relation, as on a Property Dualist picture, or an identity or constitution relation, as on a physicalist view. Since we have different physical sequences in the different frames, we might expect different phenomenal sequences to be determined – this is just what we get with other high-level properties, like functional properties, for example (since the functional properties depend on your basic physical properties, and you have different sequences of physical properties given different ways of dividing your space-time worm into temporal parts).

Perhaps it will be replied that given the puzzle we are facing, it’s not so strange to think the different physical streams produce the same results. There could be a special kind of multiple realization or over-determination here, whereby each momentary experience is produced by a multitude of instantaneous physical states, one in each frame. However, this idea is subject to another objection, which should persuade us to reject it. Consider the relation between two different frame-relative divisions of the subject into temporal parts. Consider in the first frame, a momentary experience associated with a certain time, *t*. According to the current view, such a momentary experience also occurs in the other frame at some precise time – it is multiply realized in a special way. But at what time does it occur in the other frame? The realizer of the experience in the first frame is presumably extended in space, and therefore *extended in time* relative to the second frame. Therefore, there appears to be nothing about how the experience is located in the first frame that would allow us to single out one instant in particular as the time it happens in the other frame. So if there’s a fact about this, it may be a brute fact. The rules linking the physical and phenomenal levels would have to have a certain arbitrariness that is hard to believe in.

Finally, if we assume that each momentary experience is associated with a unique region of space-time (rather than multiply realized in a multitude of overlapping slices, one for each frame of reference), then the only way in which an experience could have a precise timing in every frame would be if that region were a *single space-time point* (this is a simple consequence of the geometry of space-time). A philosopher such as Chisholm, who believes he is identical to a point-like physical particle^{xxi}, or a Cartesian dualist who believes his soul is spatially point-like, might be fine with this. But on the reasonable assumption that the most precise location we can give to an experience is an extended space-time region, this will not be an acceptable consequence.

I think these objections provide more than enough reason to reject this view. However, there are a couple of variations on the idea that are more promising. The claim we just considered is that there is the same temporal sequence of phenomenal states in each frame, in the sense that there is the same *continuous* sequence of phenomenal states. If we reject the continuous view in favor of one of its alternatives, the idea that experience has the same temporal sequence in each frame is more promising.

Alternative 1: The Discrete Model

The clearest alternative to a continuous view is the view that consciousness comes in discrete bursts, each realized over a different region of space-time, and each with a certain phenomenal character. This view in turn comes in different flavors: perhaps there is just one temporal series of pulses, or different series in different modalities, or perhaps even collections of pulses in a single modality that do not form a neat sequence in time. Views on which the pulses have an objective ordering are particularly interesting in this context. If each burst were in the absolute future of previous bursts, the bursts would form a nice time-like series, whose order would be completely frame-invariant. This would be a view on which there really is the same temporal sequence of phenomenal states in each frame (the temporal *distance* between bursts will vary from frame to frame, but perhaps that is not a problem.) If this kind of view were correct, it would be a satisfactory way to escape puzzlement about how experience fits into a space-time world.

Are there any reasons for thinking that this, or any other discrete view, is true? I am skeptical about whether there could be any good *a priori* reasons for preferring a discrete to a continuous view – this is surely an empirical matter. Furthermore, let us be clear that *introspection* is surely no guide here – experience might not seem discrete and gappy, but it seems very unlikely that we have an inner monitoring system that would pick up on this feature, especially since on any realistic discrete model, there will be a large number of pulses every second.

Although I think the issue is empirical, there are some *a priori* objections to the idea of continuous stream that are worth considering, if only to explain why they fail. The gloss on the idea of a continuous stream I gave earlier was as a continuous series of instantaneous phenomenal states (as we have seen, that gloss may be extended into a space-time context in different ways, but all those we have considered have preserved an analogue of instantaneous states). Some may be uncomfortable accepting even the

possibility of such instantaneous phenomenal states. There are a couple of more precise objections they might raise, which I will respond to. I will also argue that it is not even clear that we need to understand continuity as involving instantaneous states.

The physicalist may worry that instantaneous phenomenal states would have to be either identical with, or at least realized by, instantaneous *physical* states. But the physicalist may also think that the instantaneous physical state of a neural system is not rich enough to be sufficient for the existence of any conscious states. Perhaps only temporally extended physical states are sufficient for consciousness.

I agree that instantaneous high-level states must be instantaneously realized, in some sense. Their timing is surely derivative from the timing of their lower-level realizers – so if looking at the realizers of the state doesn't point to one moment as privileged, it is hard to see how the state could nevertheless be instantaneously located. But admitting this is consistent with the momentary physical state not being *strictly* sufficient for the existence of the phenomenal state in question. Perhaps, in order for consciousness to exist at all, a certain kind of temporally extended physical system has to exist. But once this background condition is satisfied, it may be that each instantaneous physical part plays a privileged role in fixing which particular phenomenal properties are instantiated at the moment in question. That could be true even if in isolation each momentary physical state isn't sufficient more generally for the existence of consciousness. I'll call this "the conditional view".

The physicalist might still worry that an instantaneous physical state isn't sufficient to realize phenomenal properties even in this conditional sense. But it's hard to see how there could be any a priori reason for thinking this. The instantaneous physical state of a neural system can include both static and dynamic facts about its constituents^{xxii}, like both the positions and velocities of constituent particles and the forces acting on them. These facts, together with the laws governing the system, will determine a rich functional profile for the system at that time. Of course, it *might* be that all this is still not sufficient in any reasonable sense for phenomenology, but it seems that we do not have reason for believing this just from reflecting on the richness or complexity of an instantaneous neural state. At best, we could only discover such a fact by developing a precise empirical theory of the neural realization of experience, a theory that is not currently well developed enough to rule on this matter, as far as I know.

The conditional view offers an answer to another a priori objection. The existence of a continuous series of instantaneous phenomenal states might appear to entail the possible existence of a solitary instantaneous experience. But how could there be something it's like to be conscious for an infinitely small amount of time? On the conditional view, however, even though consciousness only comes in temporally extended chunks, it is still possible for it to have instantaneous parts. It's just that these parts will not be able exist except by being embedded in an extended series of states. The continuous view therefore need not be committed to the possibility of solitary instants of experience.

Finally, it is not even clear that the problem of instantaneous experiences need arise in the first place, because there are at least two ways for continuous consciousness to exist that don't entail the existence of any instantaneous states. First, the stream could be "gunky" – every (perhaps extended) temporal part of the stream could contain a shorter temporal part, entailing that it has no instantaneous parts, or parts of a minimum duration. The failure of the stream to have such parts serves to distinguish it from a discrete stream, and thus it seems right to call such a stream "continuous", despite its not consisting of a continuous series of instantaneous states. (As far as I can tell, the puzzles we have been discussing for a "stream of moments" views arise also on a gunky view, although they may require a more complicated set-up).

Second, perhaps our phenomenal states fail to be instantaneously realized, so that there is no way to locate them at precise instants but only over extended intervals of time. Nonetheless, we can make sense of the experiences forming a continuous series, provided we can make sense of a continuous series of *overlapping* temporal intervals, each being the interval occupied by a distinct "momentary" phenomenal state. In this way, we could have a continuous series of quasi-momentary phenomenal states, despite none of them being objectively instantaneous.

The view however has a significant disadvantage compared to views where each state in the series is instantaneous. It implies that you could be undergoing infinitely many incompatible experiences all at once. For example, if you are undergoing a pain increasing continuously in intensity, then on this view you will be undergoing infinitely many stages of the pain all at once. Perhaps we could be persuaded to live with this, but it is an odd result. (On the other hand, the view may have some advantages of its own that make it worth developing more than I will here.)

To sum up, there is no compelling a priori reason to reject continuous consciousness and opt for a discrete model. Is there any empirical reason for preferring a discrete model? There is some evidence that the visual system detects changes by sampling discretely in time (for example, this is one possible explanation of the "wagon wheel" illusion) (see Koch 2004). But it doesn't clearly follow that the experiences which are the upshot of this process come in discrete bursts – at best we get a result about the kind of information carried by the experiences that result.

I should also note that it may well turn out that either space-time itself is discrete at a very fine-grained scale, or that the matter or field that permeates it varies discretely rather than continuously. If either of these possibilities obtained, then strictly speaking everything would be discrete rather than continuous. However, although consciousness would be in a sense discrete in this world, the subject would still be divisible into different series of temporal parts in different frames (although there may now be only finitely many frames), and these different series could correspond to different phenomenal series. Such discreteness would also not imply that consciousness was discrete in the sense that there would be a unique way of dividing it into a finite number of distinct pulses, as on the kind of discrete view that I'm discussing here. Therefore, the problems with continuous consciousness could arise just as much in such a world.

To sum up, a discrete model might provide a way to resolve the issues we are concerned with, but I know of no good a priori or empirical reason for preferring such a model over a continuous one.

Alternative 2: Indeterminacy in the Timing of Experience

It turns out that, if there is a difference in the sequencing of your experience between different frames, it is only a very *small* difference, in a certain sense. As I mentioned, an experience that lasts a second in one frame may last a million years in another frame, and in this sense the different phenomenal sequences are not at all similar. But they are very similar, *modulo* time dilation, i.e. once we ignore their absolute temporal scale. The fact is that in order for two events occurring in the brain to be space-like separated, and therefore capable of having different orders in different frames, they have to occur at almost exactly the same time, at least relative to a frame where the brain is close to being at rest. For example, consider your own brain in a frame centered roughly on the earth. Since your brain is only about 25cm across, and the speed of light in a vacuum is 300 million meters per second (!), it will only take about 1 nanosecond for a beam of light to traverse the length of your brain, relative to this frame. Thus, two events in your brain must occur within about 1 nanosecond, relative to an earthly clock, in order to be space-like separated, and therefore capable of having different orders in different frames. Thus, variations in your stream of consciousness between frames occur only at an extremely fine temporal scale, once we factor out temporal stretching.

What difference does that make? Even if there are only extremely fine-grained differences in the sequencing of your stream between frames, these are still differences. But postulating such differences involves presupposing that there exist facts about how the stream plays out at this fine scale. It might be that at this scale, the timing of conscious events is sufficiently *indeterminate* that no such facts exists, so we can't correctly claim that two phenomenal sequences differ at this scale. Since the only possible difference between different frame-relative phenomenal sequences is at this fine scale, the effect of indeterminacy might be that the timing of the stream does not differ from frame to frame. An analogy might be with two extremely sharp images that differ only in very fine details. If we blur the images, removing the high-frequency detail from both of them, we can make them the same image.

The kind of indeterminacy in timing that is relevant here should be carefully distinguished from another kind that may have already occurred to some readers. If two events are space-like separated, then there is no objective fact about their temporal order. We might describe this situation by saying that it is *indeterminate* what their temporal order is. Since a subject's mental states could only be space-like separated at a fine temporal scale, we might conclude that at a fine temporal scale, the ordering of experiences can be indeterminate. But although this might be a reasonable way to understand the situation we get when mental events are space-like separated, it is not the kind of indeterminacy I am talking about here. The kind I'm talking about is supposed to imply that experience has the *same* (indeterminate) timing in every frame, whereas this

latter indeterminacy is a consequence of experiential timing being *different* in each frame. The indeterminacy I'm interested in here is *frame-relative* indeterminacy.

Let me also note that strictly speaking, this isn't an alternative to the continuous view as stated, but a specific version of it. When discussing the continuous view, I have until now been assuming that part of that view is that the facts about how the stream is laid out over time are precisely determinate. The current view is an alternative to this in the sense that it denies determinacy – however, it is not an alternative in the sense that it need be construed as denying that the stream is continuous. How to think about a continuous stream that is indeterminate in timing is not clear, and would need to be spelled out much more to make the view convincing – but continuity and temporal indeterminacy are not obviously incompatible.

What could independently motivate believing in such indeterminacy? Dennett has argued that the timing of experience is indeterminate at a fine temporal scale^{xxiii}. I am uncertain exactly what his argument is, but if I were to give an argument in the spirit of Dennett, it would be something like the following. The first step is to claim that there can't be more to the subjective character of experience than is available in introspection (what exactly is meant by "introspection" would need to be spelled out). The second step is to argue that if there were facts about how experience played out at an arbitrarily fine temporal scale, this would allow for introspectively inaccessible facts about its character. For example, if there were a slight oscillation in the character of my experience with a period of 1 nanosecond, I probably wouldn't be aware of it. Put these two steps together, and we get the result that "what's it like" isn't something that could vary at a very fine temporal scale.

In the present context, we might supplement the argument with the following observation. Given the incredibly fine-grained nature of any possible differences in the temporal layout of the stream in different frames, the following is probably true: the relatively coarse temporal properties of experience that I *can* be introspectively aware of include only frame-invariant properties. For example, whenever I am aware of a change in experience, the events involved are most likely time-like separated – the experiential order we can introspect is almost certainly a frame-invariant one. But then, if the only facts about the timing of experience that *exist* are introspectively accessible, there cannot be the kind of fine-grained facts about it that would vary from frame to frame. And this is just the conclusion we were trying to establish.

The problem with this argument is that it relies on a dubious *verificationist* principle about the character of experience. You might protest that verificationism is no better in this context than elsewhere (although admittedly it has special plausibility in this specific domain). Also, to think that everything about the subjective character of experience should be introspectively accessible is, ironically, an extremely *Cartesian* assumption to make. More importantly, there are positive reasons for believing that there can be facts about how the phenomenal properties of experience change over time that are *not* introspectively accessible. For example, consider what happens when you are looking at a uniform patch of color. The quality of your experience, depending as it does on a large

number of noisy continuous parameters in the brain, itself varies noisily around a mean value, in a highly complicated pattern. The fact that there is some such noisy variation is even introspectively accessible to you (try looking at a white surface close up), though the exact pattern of variability is completely *inaccessible* to you – for example, even if it were possible to completely replicate the experience for you, you couldn't tell from the inside that it was the same. Ironically then, there is *introspective* evidence for the existence of phenomenal facts that are not themselves introspectively accessible.

Cases such as these provide a strong case for rejecting Dennettian anti-realism. But they do not positively establish that the timing of experience is determinate at the fine level we are talking about. Perhaps there is some other reason why the timing of experience is indeterminate. However, it is hard at this point to come up with other convincing reasons why such indeterminacy would exist. For example, I doubt that any semantic indeterminacy in the vocabulary we use to describe consciousness is of much relevance here. The most promising strategy I can think of is to argue that such indeterminacy would be a consequence of a failure of instantaneous realization, as discussed in the previous section. However, I will not pursue this idea here, noting that there is certainly further discussion to have about the issue of indeterminacy in experiential timing. (Let me emphasize here that considerations to do with Relativity are not in themselves a good reason for believing in such indeterminacy: for example, the kind of indeterminacy in question (frame-relative indeterminacy) is not a simple consequence of mental events being space-like separated – these events can have a completely determinate *frame-relative* temporal relation even if, in some sense, their temporal relation is objectively indeterminate.)

The ideas considered in the previous sections, that consciousness depends on a privileged frame (that of the subject of experience), or that the temporal sequencing of your experiences is the same in every frame, do succeed in viewing experience as frame-independent in a certain way. But they artificially try to preserve the idea that experience has an objective frame-independent temporal sequence. The right thing to say is that even if this is not the case, we should not conclude that phenomenology is frame-relative. This is the idea that I turn to in the next section.

Different Phenomenal Sequences Doesn't Entail Different Phenomenology

The next view accepts that a continuous stream of consciousness has a different sequence of phenomenal properties in different frames, but denies that we need to conclude that the phenomenology of the experience is frame-relative. On this view, there is a difference between the sequence of phenomenal properties that a subject enjoys, and what might be called the “overall” phenomenology of the subject's experience over time. This overall phenomenology is something that has different frame-relative *descriptions*, but it is nevertheless the same phenomenology that is described each time.

A useful analogy here might be with a continuously changing 2-D image being projected onto a wall. If you imagine a description of how the image is changing over time, you naturally imagine a moment-by-moment specification of what image is being projected at

each time (although this description will be infinite, since there are continuously many times to describe.) But of course, if Relativity is true, then there's no objective way to divide the changing image up into a series of time-slices. Rather, the changing 2-D image is best thought of as a 3-D object that can be described relative to many different ways of dividing it into a series of time-slices (or indeed not relative to such a division at all). Furthermore, just because the projection will be described as consisting in a different series of images in different frames, that doesn't mean that we are forced to think of the pattern of images as fundamentally frame-relative. Rather, we can think of each frame-relative description as specifying the same facts about how the projection is objectively laid out over space and time, as something like a 3-D pattern of color.

The view we are now considering takes a similar line about phenomenology: different frame relative descriptions of the conscious stream are just different descriptions of the same thing – the same underlying phenomenal reality – not different descriptions of some phenomenon that is fundamentally frame-relative. They are different ways of cutting up the same spatio-temporally distributed phenomenal reality. Or to get at the idea from a different angle: if the sequence of phenomenal states in one frame is completely *independent* of the sequence in other frames, that really would suggest that phenomenology is frame-relative in a deep way. But far from being independent, the different sequences are really just different partitionings of the same phenomenal facts. Since the sequences are in this way strongly co-dependent, there is no need to worry that phenomenology is frame-relative.

Why, according to this view, did we get involved in the puzzles? The first puzzle arose because if two parts of a subject's experience have different orders in different frames then, intuitively, things feel different for her. The current view can accept that such frame-relative orderings are possible. It solves the puzzle by denying that their order matters for the phenomenology. The order of these experiences is simply not among the phenomenal facts; therefore, the fact that it is frame-relative doesn't entail that the phenomenology is *itself* frame-relative.

The solution to the second puzzle is similar. It is tempting to think that two different sequences of phenomenal properties must correspond to streams that feel different to their subjects. On this view, that is a mistake: two different sequences can involve the same phenomenology (just as two different series of projected images can constitute the same spatio-temporal pattern of light on a wall). That is, it is a mistake to infer the conclusion of the argument from premise (5).

Both this diagnosis and the previous one view the source of the puzzle as a tendency we have to view the temporal properties of experience as more directly relevant to the phenomenology of temporal experience than they really are. Such a tendency is most blatantly present in the Cinematic view – on that view, temporal experience consists in experience being laid out in a certain way *over* time (perhaps also in conjunction with our inner awareness of that layout, as on Locke's view). On such a view, it seems completely inevitable that different temporal sequences of phenomenal states mean different phenomenology; the difference in the sequences means there's a difference in the

Cinematic content of the experience and so there's presumably also a difference in phenomenology (compare – two reels of film tape that differ in the sequence of frames they have printed on them.) However, on the present diagnosis, this only seems right because we are inclined to believe in a closer connection between the temporal layout of experience and its phenomenology than really exists.

Despite its appeal, there is an interesting problem this view faces. If you consider different frame-relative descriptions of the projection on a wall, then not only can we say that they are descriptions of the same underlying reality, but we can also give a positive description of what that underlying reality consists in. We know that the series of 2-D images form a 3-dimensional block, whose nature can be specified by saying what color is present at each spatio-temporal point inside the block (I'm pretending here that the color is continuously distributed over each 2-D image – of course, in reality, things would be more complicated than that.) Once we know the layout of the 3-D block of color, we know the objective facts about how the projection plays out over time, and we can derive any frame-relative description.

A theorist who believes that the conscious stream has different frame-relative descriptions might reasonably think that there exists a phenomenal analogue of the 3-D block of color – a frame-independent phenomenal reality which is simply being divided up differently by the different descriptions. However, it appears to be not so easy to conceive of the stream of consciousness in such an objective, frame-independent way.

To be more specific, such understanding is hard to obtain if we assume that experience is *continuous*. If experience comes in discrete pulses, there is no such trouble: a frame-independent description of a discrete stream would just give the space-time region associated with each discrete pulse, and say what phenomenal properties it has. Notice that we can do this regardless of whether or not the pulses form a nicely regimented time-like series. Thus the discrete view has virtues other than just suggesting a way for the ordering of experiences to be invariant. Discrete streams that aren't time-like ordered still have the nice feature that we can easily see how they are fundamentally laid out in space-time.

As I say, the problem is in obtaining a frame-independent understanding of what it would be for a continuous stream to exist. The source of the trouble is something like the following. In the case of the projected image, we can easily picture what is objectively there, because we can conceive of each 2-D part in a *spatial* way, which allows us to understand how a series of these 2-D images can accumulate together into a 3-D block. However, at least when we conceive of experience from a first person point of view, we do not conceive of it as a phenomenon that literally has spatial structure. It is not as if an experience is made up of an array of colored phenomenal pixels, each of which has a precise point-like location in space-time. If that were true then, just as in the case of the projection, we could conceive of the experiences accumulating into a single 3-D (or more likely 4-D) spatio-temporal object, which could then be divided in different ways relative to different frames. That is, we could conceive of the conscious stream as being

fundamentally something like a 4-D tube of phenomenology. Unfortunately though, due to our apparent lack of a spatial grip on experience, we cannot do this.

Such difficulties may not seem so troubling to a physicalist. For the physicalist, the stream is likely to be viewed as identical with a stream of neural occurrences in the brain (or something along those lines). There is no difficulty in forming a spatial conception of a momentary 3-D part of this neural stream, and seeing how a series of these parts could constitute a 4-dimensional neural object. The physicalist can say – if there is problem conceiving of the stream in fundamental terms, it is that we cannot do this using a first personal, or phenomenal mode of presentation. But that doesn't mean that we can't do this using another way of thinking of the experiences – an objective, physical way. The physicalist could therefore argue that, far from providing us with a kind of canonical acquaintance with experience, the first person conception actually blocks us from conceiving of the conscious stream as it fundamentally is – a 4-D physical phenomenon laid out over space-time.

There is another interesting problem for conceiving of the stream that's worth pointing out. When you consider a physical system, like the projection on a wall, or a set of fundamental particles interacting with each other, there is a strong relation between different frame-relative descriptions. Not only does one such description (provided it is complete) supervene on any other (since they are describing the same states of affairs), it is also possible to *derive a priori* one description from another. This is related to the fact that we can derive a priori a coordinate-independent description of the system in question from a particular frame-relative one: once we have such a description it is a trivial matter to derive *any* frame-relative description.

If there are different frame-relative descriptions of a continuous conscious stream, can one be derived from the other a priori? At least if we are conceiving of the stream using phenomenal concepts, the answer appears to be “no”. In the case of an ordinary physical system, such a derivation exploits our spatial conception of each frame-relative momentary state, for example in deriving a coordinate-independent description of the system. But for the reasons already mentioned, we have no such spatial grip on experience (at least 1st personally), which means that we cannot carry out such a derivation. We may have the result that although one frame-relative description *supervenes* on any other, there is no a priori derivation relation between them.

As above, this problem won't seem so pressing for a physicalist, because if the stream is a stream of neural activity which can be 3rd personally conceived, then at least under that guise such an a priori derivation should be possible (although see endnote for an important qualification)^{xxiv}. The Property Dualist however, may have to accept that the relation between different frame-relative descriptions involves an epistemic gap without a corresponding metaphysical gap: the different descriptions supervene on each other, even though there is no a priori relation between them. As many readers will surmise, this would be especially embarrassing for the kind of Property Dualist who supports their position by inferring a metaphysical gap between the mental and the physical from an epistemic gap between them.

In summary, on the current view, it is difficult to form a positive coordinate-independent conception of a continuous experiential stream, and we may not be able to derive one frame-relative description of a continuous stream from another a priori. These difficulties are more worrying for a dualist than for a physicalist, since these are only problems that apply to experience conceived of under a first person mode of presentation not under a physical mode of presentation.

How can the dualist respond? There are two possible strategies. First, to argue that even though the view we are considering is correct, there is no epistemic problem of the sort I'm suggesting. Second, and perhaps more promisingly, they could adopt a different metaphysical conception of the stream that doesn't involve them in such difficulties. As for the first strategy, I don't know how this would be carried through. One idea is for the dualist to argue that there exists a mode of presentation of experience other than the phenomenal, under which it is possible to conceive of it in a more fundamental way. However, I am uncertain what kind of presentation this would be. What about the second strategy? I think we have considered the most important metaphysical views of the stream, although let me finish this section by mentioning a couple of more outré positions (it is unclear how helpful these views are to the dualist, however).

First, we have been assuming throughout the paper that the stream is composed of adjacent temporal parts. One could deny this, for example by holding that the minimal blocks of experience last from when you wake up to when you are in a deep sleep. Each such experience would be located over many hours, with nothing more specific to be said about its temporal arrangement^{xxv}.

Another idea would be to take experience out of the space-time realm altogether, and argue that there exists a new fundamental dimension called "subjective time" in which experiences exist and are objectively ordered (of course their relationship with space-time, and their mode of interaction with it's contents would need to be explained). Or perhaps experiences could exist both in space-time and in a primitive subjective time. I won't attempt an assessment of these ideas here, leaving you to ponder their implications.

Conclusions

To sum up, in response to the puzzles I set out, we have considered a number of different views of the space-time metaphysics of experience. We rejected the idea that phenomenology is fundamentally frame-relative. We then considered two views on which experience consists in a sequence of phenomenal states objectively ordered in time. First, the view that the timing of experience in one's own frame is privileged, and second, the view that experience has the same timing in every frame. The first view was subject to a number of objections, and we also explained how its intuitive appeal is likely to be the result of confusions or implausible assumptions. The second view was seen to be only plausible if experience is a time-like series of discrete experiential pulses, or perhaps if its timing is indeterminate in a certain way. We found no positive reason for

endorsing either of these possibilities, but they could be established through further empirical evidence or philosophical argument.

Finally, we considered a view that did not assume that experience consists of an objectively ordered series of states. This view gives a diagnosis of the puzzles with which I am sympathetic – that the relation between the temporal organization of experience and its phenomenology is sufficiently flexible that such phenomena as two experiences having a frame-relative ordering does not imply that phenomenology itself is frame-relative.

We also noticed that if experience is continuous, and is constituted by a different series of temporal parts in different frames, then it is hard to conceive of what it is like in frame-independent terms. Perhaps the stream is best thought of as a 4-D tube of phenomenology, but such a tube is only conceivable in spatio-temporal terms. However, a first person mode of presentation of experience is non-spatial and therefore prevents us from picturing such a phenomenal tube. It may be that most the objective way to conceive of the stream is as a 4-D array of physical events in the brain, but that could only be a conception of experience if physicalism is true. The view therefore suggests a surprising problem for dualists, which is that if their view is correct, it may be impossible for us to conceive of what experience is really like.

References

Bayne, Tim and Chalmers, David (2005). “What is the Unity of Consciousness”. In A. Cleermans (ed.), *The Unity of Consciousness : Binding, Integration, Dissociation*. New York : Oxford University Press.

Chalmers, David (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Chisholm, Roderick (1998). “Which Physical Thing Am I? An Excerpt from “Is There and Mind-Body Problem?””. In P. Van Inwagen and D. Zimmerman (eds.) *Metaphysics: The Big Questions*. Malden, Mass. : Blackwell Publishers.

Dennett, Daniel (1991). *Consciousness Explained*. Boston: Little Brown.

Dennett, Daniel and Kinsbourne, Marcel (1992). “Time and the Observer : the Where and When of Consciousness in the Brain”. *Behavioral and Brain Sciences*, 15, 183-247.

Fine, Kit (2005). “Tense and Reality”. In *Modality and Tense: Philosophical Papers*. Oxford University Press.

Gibson, Ian and Pooley, Oliver (2006). “Relativistic Persistence”. *Philosophical Perspectives 20: Metaphysics*, 157-198.

Locke, John (1979). *An Essay Concerning Human Understanding*. New York : Oxford University Press.

- Lockwood, Michael (1989). *Mind, Brain and the Quantum*. Oxford : Blackwell.
- McGinn, Colin (1995). "Consciousness and Space". *Journal of Consciousness Studies*, Vol. 2, No.3, 220-230.
- Mellor, D.H. (1998). *Real Time 2*. London : Routledge.
- Olson, Eric (1997). *The Human Animal*. New York : Oxford University Press.
- Russell, Bertrand (1954). *The Analysis of Matter*. New York : Dover Publications.
- Tye, Michael (2004). *Consciousness and Persons : Unity and Identity*. Cambridge, MA: MIT press.
- VanRullen, Rufin and Koch, Christof (2003). "Is perception discrete or continuous?". *Trends in Cognitive Sciences*, 7, 207-213.
- Weingard, Robert (1977). 'Relativity and the Spatiality of Mental Events', *Philosophical Studies*, 31, 279-84.

ⁱ I would like to thank the following people for helpful comments and discussion: Ned Block, David Chalmers, Kit Fine, John Hawthorne, Rory Madden, John Morrison, Tom Nagel, Carlos K. Schafer, Nico Silins, Jon Simon, Michael Strevens, and Sebastian Watzl. Special thanks to Helena Wright for proof reading the paper, and correcting my terrible use of commas.

ⁱⁱ Russell (1954), p.384 "The fact that mental events admittedly have temporal relations has much force, now that time and space are so much less distinct than they were. It has become difficult to hold that mental events, though in time, are not in space." Similar arguments are presented in Weingard (1977) and Lockwood (1989), although I think their versions are unnecessarily complicated.

ⁱⁱⁱ I think this point applies to other types of entities, like numbers and divine beings, although I won't argue for that here.

^{iv} The only way to avoid this argument would be to claim that they are neither in space nor in time, but that is an extreme position – for example, it would make the problem of mental and physical interaction particularly intractable.

^v What if there is *no* fundamental level, either in the sense that all objects have objects of a more fundamental kind as parts, or in the sense that all events are constituted by the existence of other, more fundamental events? We can still ask whether the world contains mental entities all the way (infinitely) down, or whether the existence of mental items is always in the end realized by the existence of non-mental items. Acknowledging this possibility doesn't therefore affect the point I'm making.

^{vi} Throughout the paper, I make liberal use of the notion of a "phenomenal property" of experience, or the "qualitative character" of experience. I am not intending to use them in

such a way that it would be controversial whether they actually apply to experience. I believe the points I make are neutral between a wide range of views of what “phenomenal character” is. It may entirely consist in the subject’s perceptually representing certain apparent properties of the environment, at the instant in question (intentionalism), or being directly acquainted with its features (naïve realism). It may also consist in the apprehension by the mind of certain properties of mind-dependent sense-data, or may include purely qualitative or “sensational” properties that do not strictly necessitate anything about how the subject sees the world to be.

^{vii} See e.g. Bayne and Chalmers (2006).

^{viii} This view was held by Locke (see Locke (1979)), and has been defended more recently in Mellor (1998).

^{ix} Your experience might make you aware of spatial relations between external items, but experiences don’t seem subjectively to be themselves spatially structured.

^x Strictly speaking, this definition is wrong because there are events that are “light-like” separated – they lie along the path that would be taken by a light ray traveling through a vacuum, or equivalently, along the boundary between points that are time-like and space-like separated from an event. The separation between these events is conventionally thought of as neither time-like nor space-like.

^{xii} See the section entitled “Experience has the same Temporal Sequencing Relative to every Frame” for further clarification of how “sameness of phenomenal sequences” should be understood.

^{xiii} See e.g. Olson (1997), for discussion of this idea.

^{xiv} See e.g. Lockwood (1989), Chalmers (1992).

^{xv} This elucidation suggests how the meaning of “objective” here is related to the meaning of “objective” in the sense of “not mind-dependent” or “not human-dependent”. Facts that are dependent in these ways, are sometimes taken to be so because they depend on some scheme of representation that would not exist without us or our minds. Relativistic frames are not dependent for their existence on us, but they *are* arbitrary modes of description, and that helps explain the relationship in meaning between the terms.

^{xvi} See Fine (2005).

^{xvii} Why is this the denial of premise 2? According to this view, only the series of physical states of a subject relative to the subject’s own frame determines a series of phenomenal states. Therefore, on this view, premise (2) is false if it is construed as being about *any* frame-relative series of physical states. But it must be construed this way in order to entail premise (3).

^{xviii} The situation is more complicated for the Property Dualist: if God had to do the extra work of giving objects phenomenal properties once he had laid down the microphysical facts, he could presumably choose to distribute these properties however he liked, and one possible arrangement would involve only giving phenomenal properties to 3-D parts of an object that are instantaneous in the object’s rest frame. Even though this makes sense (I think), it would still involve believing in causal laws linking the physical and mental domains, which mentioned specific frames or foliations. Normal physical laws do not make any such reference, so these laws would be unique in this respect – although

psychophysical laws would be pretty unique anyway, so maybe this isn't too serious an objection!

^{xix} Stated more explicitly, the principle is something like this: If o1 and o2 are objects, and o1 has property p relative to frame f1, and o2 physically duplicates in frame f2 the state of o1 in f1, then in virtue of this duplication, o2 has property p relative to frame f2.

^{xx} To lay out the case against the Cinematic View properly would require a substantial discussion, but the basic problem is that it is making a “content/vehicle confusion”. It is a kind of resemblance theory of mental representation, in the same camp as the theory that experiences have spatial and color properties that match those of external objects, as if your eyes were firing brightly colored phenomenal rays into the environment, or you had a full-color scale model of the environment in your head. A resemblance view of temporal representation is admittedly more reasonable than these analogous views – experiences do often temporally resemble what they represent, and we can sometimes be aware of these temporal properties. But although much more could be said on this matter, I think most would agree that the model is *prima facie* very dubious, and that is enough to confidently reject any intuitions that rest on it.

^{xxi} See Chisholm (1996).

^{xxii} On some views, properties that are rates of change of other properties, like velocity, are not strictly instantaneous even though they take values at instants. For example, velocity might be thought to depend on the positions of the particle at times close to the time of the velocity, not just on the state of the particle at the time of the velocity. If experience depends on properties like velocity, then it could be argued to not be strictly instantaneous in a derivative way. However, even on this view, phenomenal variables would still be able to take values at instants. Moreover this could be no more of a good objection to consciousness being continuous than to velocity being continuous, which it clearly can be.

^{xxiii} See Dennett (1991), Dennett and Kinsbourne (1992).

^{xxiv} The physicalist could coherently deny that phenomenal descriptions have this feature. Descriptions of some high-level physical properties will not behave like this: for example, if we specify the mass of each temporal part of an object in one foliation, we cannot just read off the masses of temporal parts in other foliations. These other masses will depend also on how the first set of masses are realized – they will depend on the specific distribution of mass-density over the region occupied by the object. The distribution of mental properties over temporal slices could similarly depend on how they are realized at a lower level.

^{xxv} Tye (2004) suggests a view along these lines.