

# Temporal Experience and the Temporal Structure of Experience

Geoffrey Lee

*University of California, Berkeley*

© 2013 Geoffrey Lee

*This work is licensed under a Creative Commons  
Attribution-NonCommercial-NoDerivatives 3.0 License.  
<[www.philosophersimprint.org/014003/](http://www.philosophersimprint.org/014003/)>*

## 1. Introduction

The experience of time passing, a central part of human experience of the world, has a number of philosophically puzzling features. Perhaps the most salient is the way that our point of view seems to glide seamlessly from one moment to the next, somehow caught at the cusp of the present as it moves ever forward. How do we account for this sense of motion through time? But in addition to this elusive experience of time passing, we have temporal experience in a less mysterious sense: conscious experience is directed on temporal phenomena, such as the duration of sounds, or the temporal order of visually perceived occurrences. We perceive events playing out in time in just the same sense in which we perceive their more static features, such as shape, texture, color and illumination. Temporal experience in this latter sense is also philosophically problematic, and will be the subject of this paper.

More specifically, I will be focusing on a number of closely related intuitions about temporal experience that have informed philosophical theories in this area, but which I will argue are problematic. First, there is an intuition that experiences of extended processes, such as an experience of a segment of a melody, are themselves processes that unfold in time, rather than punctate events. For example, a perception of a melody is naturally thought of as itself unfolding through time, perhaps by having shorter experiences of individual sounds as components. Call this the "Process Intuition". Second, there is an intuition that the specific way the experience of the melody unfolds over time *mirrors* the way the melody itself appears to unfold over time. For example, if the melody is experienced as containing a low-pitched sound following a high-pitched sound, then intuition tell us that we experience the low-pitched sound and *then* experience the high pitched sound: an experienced change corresponds to a change in experience. Call this the "Mirroring Intuition". Finally, there is the idea that time as presented in experience is in some way represented by time itself. The idea is that part of the *reason* why, for example, one sound appears to follow another is that my experiences themselves

have a certain temporal order. Call this the “Representation by Resemblance Intuition”.

Against the Process Intuition, I’ll be arguing for an “Atomic View” on which there is an important sense in which temporal experiences are *not* temporally structured processes. Against the Mirroring Intuition, I’ll be arguing that temporal experiences are not constrained to have a temporal structure that in any way mirrors the temporal structure of the events they present (I will articulate several versions of this constraint). Against the Representation by Resemblance Intuition, I’ll be arguing that temporal experiences are underwritten by neural states that do not represent temporal features through resemblance, and that this implies that experience itself does not represent time by resemblance.

My view is that the most powerful considerations in this area have to do with the ways in which temporal information is processed in the brain, so that ultimately what is at stake depends on empirical considerations of a quite general kind. My methodology involves linking together claims about how temporal information is neurally coded and processed with claims about how experience itself is temporally structured, a methodology which requires assumptions about how the timing of neural and experiential events are connected, which I will make explicit below. It is worth noting at the offset that because I adopt this “bottom-up” methodology, my approach contrasts quite strongly with certain other authors such as Dainton (2000), who argues for views in this area more on the basis of phenomenological than empirical considerations.

The questions I’m considering all presuppose that we *do* have temporal experiences — experiences as of temporally extended states of affairs. Some theorists reject this assumption, endorsing a “Snapshot” view on which perceptual experience is really a series of static snapshots with no presentation of temporal features like duration and temporal order (for a recent defense, see Chuard (2012)). On this view, insofar as we are explicitly aware of temporal information, it is a cognitive phenomenon that is the result of post-experiential

processing, rather than a component of conscious experience itself. In support of the view, many people do have an intuition that certain kinds of experience have a “series of snapshots” structure: for example, it might seem as if an experience of a ball moving is really a series of snapshot experiences, each presenting the ball at a different location (we can think of this as a version of the “Process Intuition” I mentioned above). But despite the existence of these intuitions, I think it should be an uncontroversial starting point that we *do* have temporal experience. Consider auditory experience, for example. We can’t even begin to describe an auditory experience such as an experience of music without supposing that it presents the durations and temporal orders of sounds: it would not be an exaggeration to say that *all* aspects of auditory experience are temporal in the sense that I am interested in. More can be said to justify the attribution of temporal content in other cases (for example to visual experience), but here I will proceed on the assumption that the Snapshot View is wrong.

The discussion that follows will have two stages. In section 2 I explain in detail what I see as being at stake in this discussion, explaining how the three intuitions I began with are closely connected, how they can be developed into precise claims about the stream of consciousness, and how the opposing views should be understood. In section 3 I present what I take to be the best arguments against the Process/Mirroring/Resemblance views set up in section 2, arguments which can also be thought of as providing a positive case for my preferred Atomic View.

## 2. Setting up the Debate

The debate that I’m interested in has been discussed by a number of other authors, although they don’t set things up in quite the same way as I do. An important recent example is Dainton (2000, 2010), who describes a debate between “Extensionalists” and “Retensionalists”. I will not adopt this terminology here, because it does not perfectly capture any of the distinctions that I think are fundamental to the debate, *i.e.* the distinctions between views that honor the Process,

Mirroring, or Resemblance intuitions, and those that don't. Dainton's Extensionalism is clearly intended to be developed in such a way as to honor the Process Intuition and the Mirroring Intuition — at least in a weak version — but it is explicitly defined as the view that temporal experiences are themselves extended in time: a claim that (as I will argue) all parties to the debate in which I am interested can and should accept. His "Retentionalism" has much in common with my Atomism, although it involves a commitment to the idea that atomic experiences involve "Retention" — a special kind of memory experience — which I do not take to be a commitment of the Atomist, and which will not be a central focus in most of what follows (although I will say something about it in developing the Atomic View). Part of the aim of this paper, then, is to articulate what is at stake in the debate more perspicuously than previous authors. To this end, I will now develop in detail some views that are designed to either respect or reject the Process, Mirroring, and Resemblance intuitions.

Let's begin with the issue: are temporal experiences extended processes that unfold through time? To understand this issue, we need to know what it means to talk about a "temporal part" of an experience, or of the stream of experiences. A stream of consciousness is a collection of experiences. I will assume that experiences are events, and that these events are instantiations of experiential properties by subjects<sup>1</sup>. We can be fairly neutral here about what these experiential properties are, although I will assume that some involve the presentation of apparent objects or events and their properties to a subject of experience, including in particular the *temporal* properties of apparent objects and events, such as temporal order and duration.

On the property-instantiation view, talk about temporal parts of the stream of consciousness is just a way of talking about the experiential properties that a subject has at various times. For example, a momentary temporal part of the stream consists in the experiential properties that

1. Or perhaps by spatio-temporal parts of subjects, such as parts of a subject's brain or 4-D brain-worm — substituting these for the subject won't make a difference in this context.

a subject enjoys at a precise moment of time (there is an ambiguity here that I clarify below). An extended temporal part consists in the experiential properties a subject has over a period of time. (Notice that talk of temporal parts of experience is therefore a quite different matter from talk of temporal parts of material objects. There is no commitment here to an ontology of special objects, "experience parts". The only object we are committed to the existence of is the subject of the experiential properties.)

In virtue of what is an experiential property instantiated at one time rather than another? I will assume that experiences are realized by physical events in the brain, and that the timing of experiences is systematically related to the timing of their realizers. If physicalism is true, this realization is a minimal set of neural events that metaphysically necessitate the experience (its "total realization" in Shoemaker's (1981) sense), or the part of the total realization that differentiates the type of experience in question (*e.g.* the activity in V<sub>4</sub> that determines that an experience is of blue rather than of red): its "core realization". If dualism is true, I will take the neural realization to be the neural events that are the immediate physical causes of the experiences. What exactly is the relation between realizer timing and experience timing? Particularly important here will be the principle that if two experiences are realized over the same interval or moment, then they themselves occupy the same moment or interval. Call this the "Temporal Correlation Principle". The Correlation Principle follows from the stronger principle that experiences have the *same* timing as their realizers. Call this the "Temporal Identity Principle." (I think this is very plausible, but I won't argue for it here.) Note that the Correlation Principle is weaker: for example, it would be true if experiences happen at the *end point* of their realization. Note also that since experiences might have different kinds of realizers (*e.g.* total or core) that might have different timings, these principles are ambiguous. I suggest that at worst this means that "the timing of E" is ambiguous — we might need to distinguish between "core timing" and

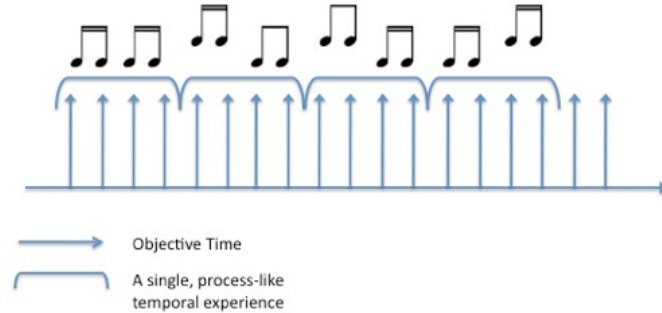


Fig. 1: The Experiential Process View

“total timing”, but there is no substantive issue about which is the “real” time at which an experience occurs.

Given these preliminaries, we can now distinguish between an Experiential Process View and an Atomic View. On the Experiential Process View (fig. 1)<sup>2</sup>, experiences of temporally extended phenomena are always built up of shorter experiences as temporal parts. For example, an experience of an object’s extended motion might be built up out of temporally disjoint experiences of parts of the motion (are these parts also built up out of shorter experiential parts, ad infinitum? — a good question, to which we will return below). Or an experience of a sound’s long duration might be built up out of disjoint experiences of the durations of shorter parts of the sound. These parts might be qualitatively varied, but they may also be distinct instantiations of the same phenomenal property (*i.e.* they have different physical realizations occurring at different times), as when one experiences a constant, unchanging scene.

On the contrasting Atomic View, temporal experiences are *never* process-like in this way. That is, an experience of temporal duration or temporal structure does not itself have temporal structure, in the sense

2. See Foster (1982), Dainton (2000, 2010), and Phillips (2008) for examples of theories that have this view as a component.




Fig. 2: The Atomic View

of having *experiences* as disjoint temporal parts. Despite this fact, each experiential atom may provide an extended window onto a temporally structured state of affairs, such as series of sounds in a piece of music (see fig. 2). (An early defender of a version of this view was C.D. Broad (1925); Grush (2005) also defends a version.)

To properly understand the Atomic View, some clarifications are needed. First, atomic experiences, although they are not experiential processes, need not be instantaneous<sup>3</sup>. An atomic experience might be a property instantiation enjoyed most fundamentally by a subject as they are over a short interval of time. Related to this, an atomic experience may be realized by an extended physical process (as I would argue all experiences are — see below). What makes this coherent is that the proper temporal parts of the realizing process need not themselves realize any experiences; in particular, they might be simply too short-lived. For example, suppose that experiences of phenomenal red are partly realized by 50hz neural firing. The existence of this firing rate arguably depends on neural firing during an extended interval, or “time bin”. The neural activity within this minimal window does have temporal stages, but they are too short to determine any firing rates,

3. Although this is a part of the view defended by Broad (1925).

and therefore (on this view) too short to determine the existence of any experiences. Thus an experience might be atomic in the sense that it has no *experiences* as proper temporal parts, but still be realized by a complex physical process.

This last point suggests a further necessary clarification. If an experiential property is instantiated holistically over a period of time (in the sense that it is not instantiated in virtue of the existence of shorter, independently existing instantiations of the same property at sub-intervals within the interval, *e.g.* because it has an extended realizer like 50hz firing), we can still say that the subject *derivatively* enjoys the property at each moment during the interval (compare how a section of a beautiful musical performance might be beautiful in its own right, vs. derivatively beautiful by being part of a beautiful whole). An extended atomic experience does have an experiential part at each moment it exists in this derivative sense, but this is the *only* sense in which it has experiences as proper temporal parts. On the contrasting Process View, a temporal experience has experiences as temporal parts in the non-derivative sense (given the Correlation Principle, we can infer that the parts have distinct physical realizers with different timings). Even if the temporal parts of a process-experience are qualitatively identical, the idea is that they are distinct, independent instantiations of the same property, rather than merely derivative aspects of a single temporally extended property instantiation.

It is very plausible that *all* experiences are realized by extended physical processes. For example, the representational content of neural states (in any reasonable sense of “representational content” that is relevant to what experiences are happening) depends on a temporally extended pattern of neural firing of large populations of neurons; for example, in many cases *firing rates*, which involve averaging the number of firings over extended intervals, will probably be relevant<sup>4</sup>. Given

4. Even if we include temporal derivatives of changing physical quantities in our description of the brain at-a-time, such as the rate of change of properties of the fundamental particles in the brain, I believe this is not enough to

the Temporal Identity Principle (we need this stronger principle in this instance<sup>5</sup>), any experience whose content depends on the content of such extended neural states (and its plausible that *all* experiences are like this), will itself be extended in time. This suggests that if we want to defend Atomism, we should defend a version on which the atoms are extended: “Extended Atomism”. This is the view I am arguing for in this paper.

Note that even if the Atomic experiences are extended, they needn’t be discretely separated in time. They might overlap in time, *e.g.* because their realizers overlap in time. For example, your experience might get gradually more intense as a neuron or group of neurons gradually fires more rapidly. Each firing rate, and hence each level of intensity, is realized over a period of time, but these periods are temporally overlapping, not discretely separated. In this way, a series of atomic experiences might form a kind of messy continuum, rather than a discrete series.

In recent work, Phillips (2011, forthcoming) endorses a kind of Process View (in particular he endorses a Mirroring constraint – see below) but also holds a view on which the different stages of temporal experiences are holistically realized (both in the core and total sense) over the same temporal interval, so that the stages are derivative from the whole experience. Notice that this involves rejecting the Temporal Correlation Principle, because it requires the view that the parts of an experience can happen at different times, despite being physically realized over the same interval. I find this view hard to understand: what is meant by the claim that experiences happen at different

---

determine the relevant properties of the neurons that will be constitutively relevant to the existence of an experience; this point warrants further discussion.

5. We cannot infer from the weaker Temporal Correlation Principle that if an experience has an extended realization, then it is itself extended in time; for example, it is consistent with the Correlation Principle that experiences happen at the end point of their realizations. The argument for Atomism given below only depends on the weaker Correlation Principle; the additional step to *Extended* Atomism requires a stronger link between experience timing and realizer timing, such as the Temporal Identity Principle.

times, if this is consistent with them having the same core and total realizer timing? Compare an analogous debate about the spatial properties of experience. Suppose someone claimed that experiences of different regions of space themselves occupy different regions of space, or that *e.g.* an experience of a banana is itself banana-shaped (“Spatial Mirroring”), but also held that the different components of a total experience are each holistically realized across the whole brain. It is unclear what it means to say that two experiences are located in different regions, or that an experience is banana-shaped, if this is consistent with them being core and totally realized in the same brain-shaped region. I will therefore assume in what follows, contra Phillips, that the disagreement between Atomists and Process theorists (including Mirroring theorists like him) about the timing of experiences implies a disagreement about the timing of the *realizers* of these experiences.

On the Atomic View, each atomic experience is associated with a “specious present”: an interval of time in the world that it presents (see fig. 2). It is easy to confuse the “specious present” in this sense with (1) the objective length of the experience itself or (2) the minimal duration between two events required for them to be perceived as non-simultaneous (sometimes referred to as the “perceptual moment” or “window of simultaneity”). For the Atomist, there is no requirement that the amount of time presented by an experience is the same as the length of the experience itself — indeed an atomic experience of duration might even be instantaneous. Furthermore, events in a “specious present” can be perceived *as* happening at different times, hence the contrast with events occurring within a single “perceptual moment”.

The content of an atomic experience need not divide events into past, present and future, as we find in Husserl’s (1964) position, which involves the well-known distinction between “retention” “perception” and “protention” (see also Miller (1985)). Atomic theorists could deny that there is any tensing in the content of experience at all. Or they could hold that everything that is perceived seems to be happening *roughly* in the present, in a way that is consistent with the relevant

events being non-simultaneous, but that no further differentiation of tense is involved.

Related to this, the Atomist can reject Husserl’s idea that temporal experience involves a special kind of “retentional” memory experience (it is because he is considering views that have this feature that Dainton (2000, 2010) calls the main class of views that contrast with his Extensional view “Retentional”). Atomic temporal experiences might involve just one kind of conscious perceptual experience, not differentiated between “retention” and “perception”. Moreover, an atomic temporal experience need not involve “retention” even in the weak sense that it involves retaining contents from immediately past experiences: a temporally extended content could include — perhaps exclusively — information about events that were not presented in any previous experiences<sup>6</sup>. We should also note the possible use of *prediction* by the visual system in some cases (see *e.g.* Nijhawan (2008)). For these reasons, I am reluctant to use the term “Retentionalist” to describe the Extended Atomic View, even though “Retentional” views are those commonly presented as the main alternative to “Extensional” process views<sup>7</sup>.

So much for clarifying what the Atomic View is. As I mentioned earlier, I think our *prima facie* intuitions on this matter may instead

6. For example, perhaps sometimes there is a kind of Dennettian “redrafting” that involves a sudden change in the brain’s view of what happened during a certain interval of time (see Dennet and Kinsbourne (1992)).
7. There may be a truth lying behind the “Retention” idea, which is that we may need to distinguish between temporally extended information that is immediately part of the content of phenomenal consciousness, and a less immediate sense of the surrounding temporal context of these events. For example, Pöppel (2004) cites various pieces of evidence in favor of the hypothesis that we have working memory representations in vision and audition that represent temporal information from the last 2–3 seconds of perceptual experience: for example, subjects are able to accurately reproduce visual or auditory information perceived within the last 2–3 second, but their performance rapidly drops off beyond this range. Perhaps the contents of these working memory states are part of a “wide” retentional specious present, a temporal analogue of our sense of how occluded parts of objects are arranged in space, despite not being immediately perceptually experienced. This idea warrants further discussion.

favor the Experiential Process View (see fig. 1). Experiences — especially temporal experiences — certainly seem intuitively to be experiential processes that play out over time rather than static atomic events, an idea that is probably reflected in the term “stream of consciousness”. Let’s now consider how this view might be fleshed out.

Fig. 1 depicts process-like temporal experiences as disjoint from one another, but the view is perfectly consistent with them overlapping by sharing temporal parts (as they do on Dainton’s (2000) “Overlap” model). In what sense, on the Experiential Process View, is a temporal experience “built up” out of shorter experiences as temporal parts? An experiential process is supposed to be more than a mere mereological sum of adjacent experiences. For example, if I experience a complete brain reset in between two experiences, presumably they will not form parts of a single temporal experience. Process theorists think that the parts of the experiential process have to be *unified* in the same way that simultaneous experiences belonging to a single subject are typically unified: they are components of a larger field of experience. Not any old pair of experiences are unified in this sense — for example, my experiences and your experiences are not, and neither are my experiences at sufficiently remote times.

It’s controversial what this unity relation is, or even whether there is such a thing. I will grant the process theorist that there is such a relation, and will not take a stand on what it is (for a summary of some different views, see Bayne and Chalmers (2000) or Prinz (2012, chapter 8). For skepticism about unity, see Tye (2003)). Fortunately, a fairly uncontroversial *sufficient* condition for unity obtains when two experiences contribute their contents to the complex content of an experience. For example, an experience of one edge of an object and a second experience of a different edge may contribute to an experience as of the object’s shape. Since the case that interests us is an experience with a complex *temporal* content, such as an experience of a segment of a melody, we can apply this sufficient condition. On the Process View, an experience of a melodic segment has simpler experiences (like experiences of individual notes) as disjoint temporal

parts, which are unified diachronically, contributing their contents to an overall experience of the sounds being temporally related a certain way. The Atomist, by contrast, denies that temporal experiences ever involve diachronic unity; for example, they think that an experience of different sounds and their relations in a chunk of music happens *all at once*<sup>8</sup>.

Note that it would be possible in principle to hold that *some* temporal experiences are experiential processes, and that others are atomic. For example, we can imagine a version of the Process View on which very short temporal parts of process-experiences are atomic experiences that present very short durations. We might refer to such a mixed view as a “Weak Atomistic View” or a “Weak Experiential Process View”. The considerations that I will bring to bear later on tell against *any* temporal experiences being experiential processes, and so will tell equally against this view.<sup>9</sup>

So much for what the Experiential Process View is. The view is very naturally developed in such a way that it does justice to the “Mirroring Intuition” that I mentioned at the beginning. On a “Mirroring View”, there is a direct correspondence between certain aspects of the apparent temporal structure of events presented by an experience, and the temporal structure of experience itself. For example, on a “Topological Mirroring View”, an experience of event A as happening before event B has as temporal parts an experience A followed by an experience of

8. The Atomist might allow that something like the experience of a whole movie is, in a sense, a single “experience” with different stages, “unified” by belonging to a single subject, but these stages would not be unified in the sense we are interested in here; for example, you don’t directly experience relations between later and earlier events in the movie. For the Atomist, it would be more perspicuous to describe this as a mere series of experiences.
9. The possibility of such a mixed view points to an issue that I won’t discuss here: the modal status of Atomism and the Experiential Process View. Are these intended as necessary truths about the kinds of things temporal experience are, or are they merely contingent truths about the form that temporal experience takes in us? (Thanks to Adam Pautz for asking me this question.) The mixed theorist clearly thinks the second option is correct, holding that such varied realization holds not just across possible worlds but within our world.

event B. Metaphorically, on the Mirroring View, experience is a bit like a mirror held up to the world: the temporal pattern of the “experiential image” directly reflects the temporal pattern of events in the world. (For a recent defense of Mirroring see Phillips (2008), which is also recently endorsed in some form by Dainton (2000, 2010) and Mellor (1985). The view is, in my experience, a constant temptation for everyone who reflects on the experience of time.)

Mirroring comes in different strengths. On *Metrical Mirroring*, the ordering *and* duration relations between the temporal parts of a process-experience match those of the apparent perceived scene. So, for example, an experience of a 1-second gap between two sounds is mirrored by a 1-second gap between the experiences of the sounds themselves. *Topological Mirroring* is weaker, only requiring mirroring at the level of temporal order, not duration. *Structural Mirroring* is even weaker, only requiring that distinct temporal stages of the perceived scene are presented by distinct temporal stages of experience.

An important point of clarification: the Mirroring theorist will obviously want to allow for illusion and hallucination. In those cases, we may not get mirroring between the *actual* temporal arrangement of a scene and the corresponding temporal arrangement of an experience, but we might have a match between the *apparent* arrangement of the scene and that of experience itself. I will assume that it is this latter kind of mirroring that is a commitment of the theory.

If you reflect on your experience for a moment, these principles can all seem very plausible, especially the weaker ones. Take Topological Mirroring: it is hard to imagine how you can experience A as happening before B without first experiencing A and then experiencing B. This is indeed how we instinctively think about the timing of experience. Nonetheless, I will argue that all these principles are wrong. More specifically, I will be arguing that we *never* experience temporal features in a way that requires mirroring.

Notice that each version of the Mirroring requirement implies the Experiential Process View. We can therefore argue against Mirroring by arguing against the Process View. In the other direction, although

it is coherent to hold the Experiential Process View but to deny even Structural Mirroring, I think it would be extremely odd to do so; presumably different stages of an experiential process present different stages of the apparent scene. So, I will assume the process theorist at least believes in Structural Mirroring. If this assumption is right, then we can also argue in favor of an Atomic View by arguing that there are no Mirroring constraints on temporal experience (see section 3.1 below).

A natural question about Mirroring concerns what it tells us about the fine-grained temporal structure of experience. For example, if you have an experience as of continuous motion, does your experience reflect the motion by being itself arranged into a continuous series of “snapshots”, each presenting the object as being at a different location?<sup>10</sup> The Mirroring theorist *could* think of experience as a continuum of momentary snapshots, but this is not compulsory. For example, Dainton (2010) thinks experience is gunky: it has no (non-derivative) momentary temporal parts, but it does have arbitrarily short temporal parts (see also Pelzcar (2011)). One could also hold that it is simply *indeterminate* how experience is arranged beyond a certain fine-grained temporal scale, or adopt a mixed view on which very short temporal parts of experience are atomic rather than experiential processes.

A Mirroring Theory on its own doesn’t include an account of *why* Mirroring holds. The factors that determine the temporal *content* of experience might be expected to be part of the story. If that’s right, then a *Resemblance theory*, inspired by the third intuition we started with, seems to fit the bill. On this theory, experiences present certain kinds of temporal features partly by themselves having the very same features. For example, on a *Topological Resemblance View*, it is partly in virtue of having experiential parts with a certain temporal order that an experience presents external events as having that order. Or on a *Metrical Resemblance View*, it is partly in virtue of having a certain

10. The assumption that motion breaks down into momentary parts does require that time is pointy and not gunky, which would be questioned by some theorists.

duration that an experience is an experience of an event having that very same duration. (Mirroring Theorists don't necessarily come out as Resemblance Theorists. For example, Phillips (2008) offers a motivation for Mirroring, but doesn't consider the question of *why* it is true; he therefore doesn't commit to any kind of Resemblance theory).

Note that Resemblance theorists are not committing the fallacy of thinking that *e.g.* a change in experience is *sufficient* for an experience of change. Their view is that it is *necessary*, but other constraints will have to be in place as well for an experience of temporal order, the most salient of which is that the relevant experiences have to be unified into a single experience, as discussed above (other constraints may be required too<sup>11</sup>).

A predictable objection is that the plausibility of Resemblance (and also Mirroring) derives from a content/vehicle confusion (this is Dennett and Kinsbourne's (1992) main objection to these views). A content/vehicle confusion occurs when an aspect of how a representation depicts the world as being is mistaken for an intrinsic property of the representational vehicle. An example of this would be thinking that your experiences of color are themselves colored. In the present case, thinking that the temporal structure of one's experience must mirror how it represents the temporal structure of the world is supposed to involve a similar confusion.

Although I agree that in at least some cases this confusion may motivate people to believe in a Mirroring/Resemblance view, I think that there are much more significant objections, which I will spell out below. The reason is that the complaint at best shows that the views are poorly motivated, not that they are false. In the case of a color content/vehicle confusion, we have independent reason for thinking

11. In fact, the Resemblance Theorist would do well to say that other constraints are required too. Having unified experience with a certain temporal structure is probably not sufficient for having experience of that temporal structure, as can be seen by considering the fact that a diachronically unified experience may have esoteric temporal features that the subject has no awareness of, such as containing parts whose durations form a geometric sequence (thanks to Ned Block for this example).

that experiences are *not* colored<sup>12</sup>. But the stream of consciousness *does* have some temporal structure or other, and for all we have said it may be guaranteed to have the temporal structure imputed by a given Mirroring constraint. Furthermore, there is no general prohibition against a representational vehicle having properties that mirror its content, and which furthermore partly explain the content through resemblance. Consider, for example, the Language of Thought hypothesis. According to the LOT hypothesis, we have mental states that represent propositions with a sentence-like structure, and these states *themselves* have a mirroring sentence-like internal structure, in virtue of which they represent a proposition with this structure. Proponents of this view have been accused of making a content/vehicle confusion, but they are quite clear that their reason for believing the view is not such a confusion, but that it is well motivated by empirical evidence. Why couldn't temporal structure work like syntactic structure in this way?

The Mirroring/Resemblance theorist might therefore claim that whether or not they are suffering from a content/vehicle confusion, an independent motivation for their view can be given. A full discussion would consider in detail all the positive motivations that could be given, including those that exist in the literature, such as Phillips (2010) and Dainton (2000). However, rather than trying to counter these positive motivations (see Lee (forthcoming) for discussion), I will present some positive arguments *against* Mirroring/Resemblance constraints, which I take to show that no such positive case can be compelling.

Let's take stock. If what I have said is correct, then the Experiential Process View goes hand in hand with at least a weak version of the Mirroring View, which in turn is best explained in terms of a Resemblance view. In what follows, I will present my case against the Process View (and a fortiori against Mirroring and Resemblance views)

12. Although your brain is colored, states of your brain are not even the kind of thing that *could* be colored. (Could the mass of your brain have a color, for example?)

and in favor of the Atomic View, starting with some counterexamples to Mirroring in section 3.1, before presenting my central argument — the Trace Integration Argument — in 3.2 and 3.3.

### 3. The Case for Atomism

#### 3.1 *Against Mirroring: Incompatibility with Examples*

In this section, I will briefly discuss some actual and hypothetical cases that might provide counterexamples to different Mirroring constraints. I will not pursue every possible response to the cases, as they are not my main argument against Mirroring, which I lay out in 3.2. The cases are helpful not only in supporting Atomism, but also because they are suggestive of the reasons *why* it is true, which will be discussed in 3.2.

Let us start with the experience of duration and rates of change (*i. e.* temporal properties that involve the temporal metric as well as the topology of time). Recall that if we have Metrical Mirroring in play, if you experience a sound as lasting 2 seconds, then your experience of the sound itself lasts 2 seconds, and more generally experienced temporal relations always are mirrored by corresponding temporal relations between parts of your experience itself. This makes for a kind of “surveillance screen” model of perceptual experience, on which changes in the world over time are correctly depicted provided perception itself changes in the right way over time. Such a theory makes predictions that appear to be in tension with certain empirically documented phenomena that I will now describe.

One prediction of Metrical Mirroring is that if you perceive one of a series of stimuli to be longer than it really is, then in order for your elongated experience not to create a *lag* between experience and reality, perception will have to compensate by either omitting a subsequent event from the stream, or presenting a subsequent event (which may be an inter-stimulus gap) as shorter than it really is<sup>13</sup>. Even though this

13. In certain cases like the well-known “stopped clock” illusion (Yarrow et al. (2004), Morrone et al. (2005)) a metrical mirroring theorist might also postulate a kind of retrospective elongation of experience *backwards* in time, which would not require subsequent compensation. If there is a long enough delay between

is a testable prediction, studies of transient duration illusions (of which there are many: see Eagleman (2008) for a review) unfortunately do not typically provide direct evidence against the existence of such compensation. However, one phenomenon which would provide a very clear counterexample if it existed would be sustained “slow motion” experience: experience that systematically and sustainedly presented durations and rates of change as stretched in time, without any compensation. This is problematic on the Metrical Mirroring View, because it would appear to require slow-motion experience to lag further and further behind reality.

Slow-motion experience is certainly conceivable, and this fact alone presents a challenge to explain why this isn’t evidence that it is possible. There are also theoretical reasons for thinking that slow-motion experience is possible. Arguably, we are only perceptually sensitive to duration and rates of change *relative* to the rate of physical processes inside our brains, just as we are only sensitive to the size of objects relative to the size of bodies or body parts. If this is right (and it is a point that deserves far more discussion), then it suggests that a being with structurally similar, but uniformly faster or slower internal processing would experience durations and rates of change as having a different magnitude or rate.

Perhaps more convincing than these abstract considerations would be real cases. There is much anecdotal evidence for the existence of cases where “time seems to slow down”, for example in cases involving extreme emotional duress or the influence of psychoactive drugs (see Sacks (2004)). Studies of the effects of dopaminergic stimulants like methamphetamines on duration discrimination partially confirm the drug anecdotes (see Meck (1996) for a review); however, the distortion effects documented in the literature are typically in the range of seconds, not milliseconds<sup>14</sup>. So, for example, these stimulants would

stimulus detection and conscious experience, this could happen without any mysterious backwards causation (see Dennett and Kinsbourne (1992)).

14. Many theorists think, partly on the basis of the selective effects found in pharmacological studies, that there are separate timing systems over these

not make music seem to play at a slower tempo, but would cause you to overestimate how long it's been since the song you are listening to began. Even if we have experiences of these longer durations, it's implausible to think mirroring applies here, so that *e.g.* a retrospective experience of 10 seconds having passed must itself last 10 seconds! This suggests that the process theory probably ought to be restricted to millisecond timing, and the Atomic View is correct for longer-range experiences of duration, if there are such things.

This suggests that the important question is: are there documented examples of sustained stretching of perception of duration *in the sub-second range*? Although there are many examples of transient duration illusions in the sub-second range (Eagleman (2008)), sustained cases are harder to come by. However, one surprising result suggests that ordinary human perception might, in a sense, be an example of such a phenomenon. It has been known for a long time that auditory stimuli are systematically judged to last longer than visual stimuli of the same length (for references see *e.g.* Droit-Volet (2007)). Penney et al. (2000) argue that the effect is due to the "internal clock" for audition running at a slower rate than for vision, not due to a faster onset or slower offset of an auditory "stopwatch". Critically for our present purposes, Wearden et al. (2006) found that the effect occurs for duration perception in the millisecond range (many of the studies of this effect in the literature look at durations in the range of multiple seconds). This suggests that consciously perceived sub-second durations systematically appear longer in audition than they do in vision, even though audition does not lag behind vision! If this is the right interpretation (and admittedly establishing this would require more discussion (see footnote for a different interpretation)<sup>15</sup>), this is a particularly compelling counterexample to metrical mirroring.

different ranges (Rammsayer (1999)).

15. A different interpretation is that this effect involves a discrepancy in post-perceptual judgments of duration in vision and audition, but no discrepancy in the conscious experience of duration itself. Once we accept that the effect is caused by an auditory clock running faster than a visual clock, this is an implausible idea, because it is implausible to interpret these clocks as operating

A study by Johnson et al. (2006) provides further evidence against Metrical Mirroring. They showed that adaptation to viewing a drifting grating causes a subsequent visual stimulus presented in the same region to appear with shorter subjective duration than normal. They also made the target stimulus flicker at a certain rate, and on certain trials probed for the effect of adaptation on apparent flicker frequency rather than subjective duration. The Metrical Mirroring theory would appear to predict that if the flickering stimulus looks shorter after adaptation, it would appear to flicker at a higher frequency, because more oscillations are presented in a shorter period of time. However, Johnson et al. found the opposite effect — a *lower* perceived frequency. This kind of inconsistency is hard to explain on the Mirroring theory<sup>16</sup>.

The Metrical Mirroring theory also implies that in this study the subjectively contracted experience is *objectively* shorter than it would have been in the non-adapted condition. Allowing for this seems to require a shift in the onset or offset time of the experience, and so, on the Mirroring view, a shift in the apparent onset or offset time of the stimulus itself. Johnson et al. explicitly tried to rule out such a shift by asking subjects to judge whether the onset or offset happened before or after an auditory burst timed to be close to the onset or offset of the stimulus: if there is a shift, it should change the apparent temporal relation between the onset/offset and the burst. They found *no* shift in perceived onset or offset relative to the auditory burst in adapted cases, despite a contraction of subjective duration.

In response, the Mirroring theorist could postulate that the apparent onset of the auditory burst was *also* affected by visual adaptation; but this is empirically implausible, especially given the evidence Johnson et al. present that the adaptation effect is not only modality-specific

---

downstream of, or independently of, experience, especially if we want to attribute duration content to experience. Post-perceptual interpretations of temporal processing of this form are treated in detail in section 3.3 below, in response to replies A and B to the Trace Integration Argument.

16. Watzl (2012) discusses in detail a related counterexample to metrical mirroring, involving experiences (allegedly) presenting different rates of change, despite themselves changing at the same rate.

but specific to a certain *region* of the visual field. These results are therefore also hard for the Metrical Mirroring theory to accommodate.

From the point of view of cognitive theories on which temporal frequency, temporal order, visual duration and auditory duration are detected by *different mechanisms*, it should come as no surprise that such examples are possible. Furthermore, on theories of duration perception on which duration is measured by a counter, accumulator, or other clock-like neural process (see Grondin (2010) for a review of different models), there is no reason why, if the neural clock gives a faulty readout one moment, it need “compensate” by giving another faulty readout a moment later, nor why a neural clock could not operate at a uniformly slower rate for an extended period of time. That duration experience is the *output* of such a measuring process is an instance of the general point I will discuss in the next section.

One way for a Mirroring Theorist to avoid these counterexamples would be to adopt a weaker form of Mirroring, for example by endorsing only Topological Mirroring. Are there counterexamples to these weaker mirroring constraints?

One interesting example is the fact that one can perceive two events as non-simultaneous, without having a perception of their temporal order. In particular, there is evidence (see *e.g.* Hirsch and Sherrick (1961), Mitrani et al. (1985)) that events have to be further apart in time for you to perceive their ordering than they do for you to perceive them as non-simultaneous (*e.g.* you might be able to tell that two light flashes at different locations happened non-simultaneously, without being able to tell what order they have). Suppose I experience two events as non-simultaneous without being presented with their temporal order: how does the Mirroring theorist allow for this? To literally mirror this content, the experiences would have to be themselves non-simultaneous without having some particular temporal order. Assuming they do have some particular temporal order, is this consistent with Topological Mirroring? Why don't I experience this temporal order? We know the experiences are unified, so this cannot be what prevents us from experiencing the relationship.

Perhaps the order is experienced, but nonetheless inaccessible to me? This does not seem plausible. The example at least demands further explanation from the Topological Mirroring theorist.

Finally, as Dennett and Kinsbourne (1992) emphasize, to understand the timing of temporal order experience it is important to realize that signals from different events may take different amounts of time to arrive at the sensory periphery, and different amounts of time to process to a point where temporal order comparisons are possible. To avoid temporal order illusions, the brain must take into account these discrepancies, which it is able to do to some extent, either by recalibrating (see *e.g.* Vroomen et al. (2004)), or by inferring temporal relations from assumptions about event-relations like causal relations, such as hold between actions and their perceived effects (Morein-Zamir et al. (2003), Haggard et al. (2002)). In some cases, such adjustments can cause temporal order illusions (*e.g.* Stetson et al. (2006)). In general, even if a signal from A is detected and processed before a signal from B, B may be experienced as happening before A. This at least casts some doubt on Topological Mirroring, since it is unclear why the system would have to *reverse* the order of A and B at any stage, projecting B then A onto the movie screen of consciousness, to consciously represent this information<sup>17</sup>. In the next section, I explain in general terms why this is definitely *not* required.

To sum up, there are numerous possible examples of experiences that can represent certain temporal facts about the objects and events in the environment (either metrical or topological), without having a mirroring temporal structure. Not only do these examples provide

17. A similar point can be made about the “color phi” and “cutaneous rabbit” examples discussed in Dennett and Kinsbourne (1992). Although the Mirroring theorist can interpret the color phi experience in a way that is consistent with her view, it is unclear whether this interpretation is correct. I should note that Grush (2005) presents such cases, including also the “flash lag” effect, as counterexamples to the Process View, although the reason they are supposed to be counterexamples for him is different from the reasons I give here: his focus is on an alleged inability of the Process View to allow that the temporal content of experience can be “redrafted” from one moment to another, as the brain changes its mind about what happened at a given time.

evidence against Mirroring constraints (and therefore, indirectly, against the Process View) they suggest a deeper lesson, which is that temporal representation is a cognitive achievement over and above other kinds of representation like spatial representation, so that successfully producing a series of veridical experiences of static features of the environment does not guarantee veridical experience of any temporal relations, (even if the static experiences are unified), contrary to the “surveillance screen” model offered on a Mirroring View. The brain always has to do more work to extract temporal information from a temporal series of informational states in order for the information to figure in experience. In the next section, I elaborate this point, which I think provides us with the deepest objection to any kind of Mirroring/Resemblance or Process View, in the sense that it explains *why* these views are mistaken, and why counterexamples of the kind offered here are bound to arise.

### 3.2 *The Trace Integration Argument*

The most fundamental reason for rejecting any kind of Process View and favoring Extended Atomism is that one’s experiences are realized by neural states that do not represent temporal relations through resemblance, but rather through a kind of coding that is “simultaneous” in a sense that I will explain. We can appeal to this simultaneous coding to argue that temporal experiences are not experiential processes, but rather are atomic (that is, we can reject not just the Resemblance View, but also the weaker Mirroring and Process views).

Now, although claims about how the brain represents time are empirical claims, the claim that the neural representations of temporal features underwriting temporal experience aren’t processes with sub-representations as proper temporal parts (and, a fortiori, that time isn’t represented by resemblance by these realizer states) has less the status of an empirical discovery and more the status of a background assumption that is just taken for granted by those modeling the mental representation of temporal features. The background picture is something like this: at the periphery, for example on the retina,

### *Temporal Experience and the Temporal Structure of Experience*

we have temporal information spread out in time, represented by resemblance. The job of temporal computation is to compare or integrate the different temporal stages of this input representation to produce an explicit representation of the temporal information that can influence downstream processing and behavior in an appropriate way. The point is expressed nicely in the following passage from Mauk and Buanomano (2004):

Analyses of the neural basis of timing have generally focused on three general computational strategies: mechanisms based on neural clocks, mechanisms based on arrays of elements that differ in terms of some temporal parameter, or mechanisms that emerge from the dynamics of neural networks. In general, these models must accomplish some variant of the same computational task. *They must recode the temporal information present in an input into a spatial code.* That is, in some way different cells must respond selectively to temporal features of the stimulus. For example, to discriminate differences in the duration of two stimuli, there must be differential neuronal responses to each duration. (Mauk and Buanomano (2004)), italics added by me).

What Mauk and Buanomano mean by “spatial coding” has to be understood carefully. For example, neural firing rates can count as a form of spatial coding. What matters is that the system starts with information that is spread out in time, with different components of the information represented at different times, and then processes it in such a way that it is represented *all at once*, in an explicit form that allows the representation to play a certain causal role appropriate to its content, such as causing a report of the information represented, or reasoning with the information. This is what “spatial coding” is, and clearly a firing rate, although not strictly simultaneous, is in this category. For example, suppose I have to press a button just if A

happens before B rather than B before A; I might be wired up to press the button only if a neuron tuned to detect temporal order fires above a certain rate. I claim that these “simultaneous” representations will realize atomic experiences rather than process-like experiences. Let me now explain in more detail what this means and why it is true.

A very simple example to illustrate the point is the Reichardt motion detector – one of the simplest possible motion detectors, originally modeled on the way in which motion is detected by neurons in a fly, illustrated in the diagram below. The kind of Reichardt detectors actually used in motion detection, even by a fly, in fact have to be more complex than this one, and human motion detection involves a complicated pooling of information from different detectors. This is therefore not a realistic or complete model of how we detect motion, but it illustrates the point I want to make clearly.

The detector on the left is sensitive to the presence of an edge at a slightly different location from the detector on the right. An edge moving at the right speed and direction will set off the left detector just before the right detector, but because there is a delay filter on the connection out of the left detector, the signals from the detectors will reach the AND gate at the same time, and it will fire, indicating the presence of an edge with a certain velocity and location.




Fig. 3: A simplified Reichardt motion detector

If you consider the temporally extended event of the left detector firing followed by the right detector firing (ignoring what happens later), you can see that in some sense this composite event is a representation of the presence of a moving boundary – it happens just if a certain kind of motion occurs in the environment<sup>18</sup>. However, it is also obvious that we shouldn't conclude from this that the representation of motion after the AND gate is superfluous. The representation at the periphery can in fact be considered to involve a kind of representation-by-resemblance, the firing of each detector representing the spatial position of the edge, and the extended event of one firing after the other representing the motion – the firings “mirror” the motion of the object. But clearly there is an important sense in which the peripheral composite event represents the motion only implicitly, whereas the later representation can be regarded as a more explicit representation of motion.

What is the relevant understanding of the implicit/explicit distinction here? Again, what matters is that the temporal stages of the peripheral representation will have to be compared or integrated to get the information in a format that is useful to the system for such post-perceptual purposes as control of action and reasoning. We can spell this out as what I call the “trace integration” argument. Suppose we have temporal information at the periphery that is contained in a temporally extended pattern of receptor stimulation, say on the retina. In order for the representation of this information to causally impact post-perceptual processing, each relevant temporal part of the initial extended stimulation has to *leave a trace* in the brain. If you consider the process leading up to, say, a verbal report of the information, you can see that each of these traces will have to be *simultaneously present*

<sup>18</sup> This statement is in fact an oversimplification, since it's easy to imagine other stimuli causing this input. For example, two edges a certain distance apart moving in the direction opposite to that to which the detector is supposed to be sensitive could set it off. This is in fact one reason why a real Reichardt detector is more complicated, involving a separate detector for motion in the opposite direction that is combined in an opponent fashion with the original detector. This more complicated detector can distinguish a moving edge from this bogus stimulus.

before the report is made: the alternative is that traces from certain parts of the stimulus no longer exist, and therefore can have no causal impact on the report. Furthermore, the traces will have to be *integrated* in the right way for simultaneous representation of each temporal stage of the input to be able to control later processes in an appropriate way: if each relevant stage of the input leaves a trace but the traces are in completely different neural populations that aren't functionally integrated, the information is not present in a useful form; it is not "explicit" in the relevant sense. For example, you can see in the motion detection example that if the representation of motion is going to have an appropriate later effect, the triggering of the first detector has to cause a trace that is integrated with the trace from the triggering of the second detector. *Models of temporal computation implicitly assume that setting up such traces and then simultaneously integrating them is the task that the brain has to perform*<sup>19</sup>.

Supposing it is correct that temporal processing always involves trace integration, we can ask: at what stage of the process do we find the core realizers of the temporal content of *experience*? The defeasible criterion we usually employ in figuring out the content of experience is that the information in experience is "access conscious", to put it in Block's (1995) terminology (indeed, on many views it *couldn't* be conscious unless it was available in this way), that is, it is accessible to central processes such as verbal report. Given the above discussion, this suggests that experience is realized by the *output* of trace integration, since this is the process that codes temporal information in a form that makes it available for such uses.

19. It is clearer what this means for the perception of temporal order than for the perception of *duration*, but I intend the principle to apply to this case as well. In the case of duration perception, each temporal part of a temporally extended input whose duration is computed "leaves a trace" in the sense that its duration contributes causally to the process of tracking the duration of the whole stimulus from beginning to end. For example, if duration is measured by an accumulator, each temporal stage of the stimulus contributes its duration to the accumulated representation of duration that is in place as the stimulus ends.

If this is right, then temporal experience is realized by states that do not code time by time itself, but rather use "simultaneous" or "spatial" coding for this information. For example, a trace-integrated representation of A as happening before B will involve *simultaneously* representing A and B, albeit *as* happening non-simultaneously. Given the principle that experiences that are realized at the same time happen at the same time (the Correlation Principle), we can infer that you have *experiences* of A and B that have the same timing, despite A and B appearing *as* happening at different times<sup>20</sup>. That is, your experience of "A then B" is atomic, rather than process-like. The same reasoning applies for *any* temporal experience: for the information to be accessible it must be trace-integrated, and so simultaneously represented. Therefore, Atomism is generally true as a theory of temporal experience. In summary:

- (1) In order to be accessible to post-perceptual processes like verbal report, high-level motor control and domain-general reasoning, temporal information has to be trace-integrated.
  - (2) Unless we have strong contrary evidence, we should assume that the contents of experience are accessible to these post-perceptual processes.
  - (3) Therefore, we should assume that temporal information in consciousness is realized by the *output* of trace-integration.
  - (4) The components of the content of a trace-integrated representation are represented at the same time (or over the same interval).
20. We could also argue as follows: if the neural states underlying experience do not represent time by resemblance, then neither does experience. If experience doesn't represent time by resemblance, there is no reason to think that it satisfies mirroring or is process-like. I prefer to argue directly against the Process View by appealing to the Temporal Correlation Principle, as this does not require assuming that experience is only process-like if it represents time by resemblance, an assumption not all process theorists would accept.

- (5) Experiences that have the same realizer timing have the same timing (the Temporal Correlation Principle).
- (6) Therefore, the different experiential parts of a temporal experience have the same timing.
- (7) Therefore, temporal experiences are atomic, not process-like.

This, briefly stated, is the reason why Atomism is true. In the next section, I respond to some objections to the argument, at the same time hopefully clarifying how exactly the argument works (in particular, I will say more about the role accessibility plays in the argument).

### 3.3 Replies to the Trace Integration Argument

I can envisage the following replies to the argument. Space does not permit me to give full counter-replies to all of them: I will try to at least cover some key points. I will give (A) and (B) most consideration, because I think they are the most likely objections, and it is particularly illuminating to see why they fail.

- (A) Trace Integration *is* required for temporal information to be accessed by central processes, but happens *downstream* of experience. (Rejects premise 3).
- (B) Like (B), except that trace integration is part of what *unifies* the stages of a process experience, and so *is* part of the realization of experience, although it is downstream of the activity that determines the temporal content of experience. (Rejects premise 3).
- (C) Trace Integration is not necessary for temporal information to be available to post-perceptual processes. (Rejects premise 1).
- (D) Experience happens in a separate processing stream from the one that contains the integrated traces, which therefore

### *Temporal Experience and the Temporal Structure of Experience*

provide no evidence for Atomism. (Probably rejects premise 2).

- (E) Experience happens downstream of simultaneously integrated temporal representations, which produce a further stage of representations in which integrated temporal representations are unpacked into a cinematic form where “time is represented by time”, and which realize temporal experience. (Rejects premise 3).
- (F) We should reject the Temporal Correlation Principle: the fact that two experiences are realized over the same interval does not give us reason to think they happen at the same time. (Rejects premise 5).

Let’s start with (A) and (B). They share the idea that trace integration *does* happen in the way described, but it is the *input* to trace integration, not the output, that determines the temporal content of experience. The difference is that (B) does at least postulate a role for trace integration in realizing temporal experience, in that it says that it is part of what *unifies* the stages of a process-experience. On these views then, there is an initial layer of processing which outputs a series of representations that underwrite experiences of “static” properties like shape, color, texture and illumination, and whose temporal relations realize the temporal content of experience. Temporal information is only explicitly integrated downstream of this initial layer, *post conscious perception*. Thus, although experience does have temporal content that is mirrored in its temporal structure, this content is only available post-perceptually through further processing. One could interpret this further processing as the subject introspecting how her experience is changing over time — a kind of “inner movie theater” view (Locke (1979) appears to have held such a view), but the post-perceptual integration could also be sub-personal.

The fact that on (B) this further integration is necessary for the initial stages to realize a unified temporal experience makes

it considerably more plausible than (A), as it is hard to believe that the processes involved in trace integration are not required *at all* for temporal experience. Consider motion experience; we know that some motion information in vision is computed in area MT, damage to which can cause akinetopsia – the inability to perceive motion (Zeki (1993), Zihl (1983)). Presumably MT is performing trace integration on its inputs (it contains neurons that respond selectively to motion in different directions (Albright (1984)), which are presumably fed by Reichardt detectors or functional analogues thereof). On view (A), MT will be interpreted as acting downstream of experience, merely making motion information which is already in experience available for later processes. This suggests that akinetopsic subjects may after all consciously perceive motion; it's just that they don't have any access to this aspect of experience. This (at the very least) violates our standard working maxim that subjects have access to the content of their experiences, a maxim that it seems very implausible to violate in this case (I do not say there are no cases where conscious contents are to some degree inaccessible – see Block (2007)). (B) avoids this implausible conclusion because its proponents can say that MT activity makes motion experience possible by unifying experiences of different positions along an object's trajectory.

Despite having this advantage, a similar point nonetheless holds against (B), and also (A). Both views identify the temporal content of experience with the temporal content of the *input* into trace integration. But if temporal content is realized *early* in this way, then it threatens to be peculiarly inaccessible to its subject. Consider, first, cases where integration goes wrong, so that *e.g.* a clock mechanism fails to accurately record the duration of an input, or a temporal order discriminator outputs a different order to the one inputted. In these cases the subject will *think* that her experience had a content corresponding to the output of integration, even though on views (A) and (B), they are wrong. Notice how this corresponds to a possible strategy for responding to the counterexamples in 3.1 – it could be that subjects are simply *mistaken* about their experiences in these

cases, which do after all satisfy Mirroring constraints. But also notice how desperate this seems – we need a very strong reason to ascribe this kind of error to subjects: to repeat, absent strong reasons, we should prefer interpretations on which subjects have access to what they are experiencing.

The problem actually runs deeper than this. Contrary to (A) and (B), often it is not the case that non-temporal features like color, shape, texture and illumination are computed first, and then temporal information about how the object is changing over time is computed on the basis of comparing this non-temporal information across time. Rather, sometimes the reverse is true: *information about how the world is changing over time is used to compute the static properties of objects*. One important example of this is the phenomenon of “structure from motion”, (*e.g.*, Ullman (1979)) whereby the brain uses motion information as a depth cue, to infer the 3-D structure of an object. This phenomenon is well illustrated by the “kinetic depth illusion” whereby a 2-D array of moving dots appears to form a 3-D solid because the different 2-D motions of the dots are interpreted as motions at different depths.

Go to this website for a nice illustration:

[www.georgemather.com](http://www.georgemather.com)

In this example, it is clear that your experience does not consist of a series of static experiences of 3-D arrangements of dots, from which your brain figures out how the dots are moving. One *could* hold that structure from motion processing happens wholly downstream from experience, so that the content of experience is realized much earlier than we thought (somewhere close to the retina??). But this would violate the accessibility maxim in an even more blatant way, both for temporal *and* non-temporal features.

There is also support for this view of the information processing hierarchy at the neural level. There are feedback connections between area MT (the area mentioned above that is involved primarily in temporal processing) and other areas involved in spatial processing (among other things), such as V1 (see *e.g.* Silvanto, Cowie, Lavie and Walsh (2005)). Furthermore, the processing of temporal information

in the brain happens in parallel with processing in other areas involved in spatial processing, not downstream of these areas; for example there are motion-sensitive neurons in early visual areas like V<sub>1</sub>, and although MT receives most of its input from V<sub>1</sub> (which also feeds forward to other areas involved in spatial processing), there is evidence that some information from the retina gets fed directly to MT without even passing through V<sub>1</sub>, let alone through the later visual areas involved in spatial processing (see *e.g.* Zeki (1993)).

To be clear, I should stress that I do not want to deny that some temporal computation that is related to the contents of temporal experience happens downstream of experience. For example, Cavanagh (1992) argues for a form of high-level movement perception that involves recruiting attention to track an object over time based on continuity of high-level features like 3-D shape, rather than computations involving low-level motion detectors earlier in the visual system. Perhaps high-level motion awareness is in some sense “downstream” of the conscious perception of these high-level features (similarly, some temporal order perception may be fairly central (Pöppel (1988) p.20)). Still, my view would be that the experience associated with this high-level computation would be realized by the *output* of the computation, not the input. That is, the computation is not merely performing the function of making a content that is already in conscious awareness available for report and other central functions, but rather it is responsible for making the information conscious in the first place.

One modality where the correctness of my interpretation seems especially obvious is audition, because most of the relevant information in audition has to do with how things are changing over time (an important contrast with vision). Audition is so thoroughly temporal that we have no clear idea what “static” auditory experiences, to be compared post-perceptually on views (A) and (B), would even be.

Having dealt with (A) and (B), let me briefly say something about the other responses. I think it fairly safe to assume that (C) is mistaken for the reasons given in 3.2, although ideally we would look

in more detail at actual empirical models of temporal computation and the ways in which temporal information is used downstream of experience. An interesting challenge to the argument in this category comes from views in the dynamic systems tradition which eschew the whole idea that information is computed by perceptual systems and tend to favor a picture on which perceptual and actions systems are directly coupled together in a complex feedback loop that resists computational modeling. If temporal information is not explicitly computed at all, this would clearly be problematic for the argument I’m making. I am not at all inclined to completely reject the computational picture, although there are at least some cases involving timing where something like the dynamic systems picture seems quite plausible: consider, for example, the kind of complex feedback that enables a musician to play an instrument in time. A more complete discussion would need to consider cases like this.

The problem with (D) is that if temporal experience happens in a different processing stream from the one that involves trace integration — *i.e.* from the one that leads to explicit judgments, high-level motor planning, reasoning, and so on, then it would appear to be completely epiphenomenal with respect to these processes. One thing that makes this very implausible is that it would make it mysterious how we could have *knowledge* of the temporal content of experience — as presumably this requires temporal experience to feed into conceptualized judgment, *via* trace integration. In effect, we will have another violation of the presumed accessibility of experience.

(E) postulates that once temporal information is explicitly computed, it is unpacked again in order to be broadcast in consciousness in a form where time is represented by time (so it is downstream of the output of trace integration). By analogy, think of the way in which a musical score can be unpacked into a real-time musical performance. This model is similar to (D), in that this unpacking would, in effect, constitute a separate processing stream from the one leading to other downstream effects. It therefore suffers the same problem — it makes experience epiphenomenal with respect to judgment (although at

least a case could be made here that they are well correlated by having an immediate common cause). Relatedly, the view seems empirically implausible unless some reason can be given why unpacking temporal information in this way would have some function, or some direct evidence is given that this is happening. I do not say this can't be done — for example, performing actions that require mimicking the perceived temporal layout of events (*e.g.* reproducing a perceived duration or temporal pattern) requires just such unpacking — the actions themselves “represent time by time”, if you like. Still, as far as I know, there is no empirical reason to think all experienced temporal information is *immediately* unpacked. This is certainly an issue warranting further discussion.

A final response to the argument is to reject the Temporal Correlation Principle, and hold that the different parts of a temporal experience can happen at different times, even though each is physically realized (in either the core or total sense) over the *same* temporal interval (response (F)). I mentioned earlier that Phillips (2011, forthcoming) seems to commit to some such position, which is hard to understand, for the reasons I gave. Certainly more can be said to justify the Correlation Principle; I hope to address this issue elsewhere.

This ends my discussion of possible replies to the Trace Integration Argument, and my case against the Process View, and for Atomism. In sum: the most likely objection to the argument is that temporal content of experience is realized by the *input* to trace integration not the output. My objection to this is that it violates the maxim that absent strong reasons to the contrary, we should assume subjects have access to the contents of their experiences. As my discussion suggested, there is certainly more to be said about these issues, so I do not regard what I have said as providing a watertight case against the Process View — although I do hope to have pinpointed its main weakness.

#### 4. Conclusion

I argued that the best way of framing the debate about temporal experience that I and other philosophers have been interested in, is as a debate about three closely related issues: (1) Are temporal experiences atomic or process-like? (2) Are there any “Mirroring Constraints” on temporal experience? (3) Are any aspects of the temporal content involved in conscious experience fixed by resemblance? I discussed a number of considerations that relate to these questions, ultimately arguing for an Atomic View of the stream that implies a negative answer to the latter two questions. A general lesson of the paper is methodological: we can't hope to get the right view of the constitution of the stream of consciousness without thinking carefully about the possible metaphysical forms it could have, and about the relationship it stands in to its neural realization in the brain.<sup>21</sup>

#### References

- Albright, T.D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology*, 52(6), 1106–1130.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30(5), 481–498.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences* (1995) 18, 227–287.
- Broad, C.D. (1925). *The mind and its place in nature*. T.K. Paul (ed.). London: Routledge & Kegan Paul.
- Cavanagh, P. (1992). Attention-based motion perception. *Science*, 257(5076), 1563–1565.

21. This paper evolved through multiple versions over a considerable period of time. Thanks to everyone I've talked about the topic with during that time. I owe special thanks to Ned Block, David Chalmers, Adam Pautz, Ian Phillips and two anonymous referees for helpful comments and discussion.

- Chuard, P. (2012). Temporal Experiences and their Parts. *Philosopher's Imprint* 11, no. 11, 1–28: [philosophersimprint.org/011011/](http://philosophersimprint.org/011011/)
- Dainton (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. Routledge.
- Dainton, B. (2010). Temporal Consciousness. Stanford Encyclopedia of Philosophy entry: [plato.stanford.edu/entries/consciousness-temporal/](http://plato.stanford.edu/entries/consciousness-temporal/)
- Dennet, D. and Kinsbourne, M. (1992). Time and the Observer — The Where and When of Consciousness in the Brain. *Behavioral and Brain Sciences*, 15, 183–247.
- Droit-Volet, S., Meck, W. H., & Penney, T. B. (2007). Sensory modality and time perception in children and adults. *Behavioural Processes*, 74(2), 244–250.
- Eagleman, D.M. (2008). Human time perception and its illusions. *Current Opinion in Neurobiology*, 18(2), 131.
- Foster, J. (1982). *The Case for Idealism*. London: Routledge & Kegan Paul.
- Grondin, S. (2010). Timing and time perception: a review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, & Psychophysics*, 72(3), 561–582.
- Grush, R. (2005). Internal models and the construction of time: generalizing from state estimation to trajectory estimation to address temporal features of perception, including temporal illusions. *Journal of Neural Engineering*, 2(3), S209–S218.
- Haggard P., Clark S., Kalogeras J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, 5, 382–385.
- Hirsh I.J., Sherrick C.E. (1961). Perceived order in different sense modalities. *Journal of Experimental Psychology*, 62, 423–432.
- Husserl, E. (1964). *The Phenomenology of Internal Time-Consciousness*. Indiana University Press.
- Johnston, A., Arnold, D.H., & Nishida, S. (2006). Spatially localized distortions of event time. *Current Biology*, 16(5), 472–479.
- Lee, G. (forthcoming). Extensionalism, Atomism and Continuity. In Oaklander, N. (ed.) *Debates in the Metaphysics of Time*. Continuum — Bloomsbury.
- Locke, J. (1979). *An Essay Concerning Human Understanding*, Book 2 Chapter 14. Oxford University Press.
- Mauk, M., Buonomano, D. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, 27, 307–330
- Meck, W. H. (1996). Neuropharmacology of timing and time perception. *Cognitive Brain Research*, 3(3), 227–242.
- Mellor, D. (1985). *Real Time*. Cambridge University Press.
- Miller, I. (1985). *Husserl, Perception and Temporal Awareness*. MIT Press.
- Mitrani, L., Shekerdjiiski S., and Yakimoff, N. (1986). Mechanisms and asymmetries in visual perception of simultaneity and temporal order. *Biological Cybernetics*, 54(3), 159–165.
- Morein-Zamir, S., Soto-Faraco, S., & Kingstone, A. (2003). Auditory capture of vision: examining temporal ventriloquism. *Cognitive Brain Research*, 17, 154–163.
- Morrone, M. C., Ross, J., & Burr, D. (2005). Saccadic eye movements cause compression of time as well as space. *Nature Neuroscience*, 8(7), 950–954.
- Nijhawan, R. (2008). Visual prediction: Psychophysics and neurophysiology of compensation for time delays. *Behavioral and Brain Sciences*, 31, 179–198.
- Pelczar, M. (2011). Must an Appearance of Succession Involve a Succession of Appearances? *Philosophy and Phenomenological Research*, 81(1), 49–63.
- Penney, T.B., Gibbon, J., & Meck, W.H. (2000). Differential effects of auditory and visual signals on clock speed and temporal memory. *Journal of Experimental Psychology: Human Perception and Performance*, 26(6), 1770.
- Phillips, I. (2008). Perceiving temporal properties. *European Journal of Philosophy*, 18(2), 176–202.
- Phillips, I. (2011). Indiscriminability and experience of change. *The Philosophical Quarterly*, 61(245), 808–827.
- Phillips, I. (forthcoming). The Temporal Structure of Experience. In D. Lloyd and V. Arstila (eds.). *Subjective Time: The Philosophy, Psychology, and Neuroscience of Temporality*. MIT Press (in press).

- Pöppel, E. (1988). *Mindworks: Time and Conscious Experience*. Harcourt Brace Jovanovich: Boston.
- Pöppel, E. (2004). Lost in time: a historical frame, elementary processing units and the 3-second window. *Acta Neurobiologiae Experimentalis (Wars)*, 2004, 64, 295–301.
- Prinz, J. (2012). *The Conscious Brain*. OUP USA.
- Rammsayer, T. H. (1999). Neuropharmacological evidence for different timing mechanisms in humans. *The Quarterly Journal of Experimental Psychology: Section B*, 52(3), 273–286.
- Sacks, Oliver (2004). Speed: Aberrations in Time and Movement. *The New Yorker*, 23 August 2004.
- Shoemaker, S. (1981). Some varieties of functionalism. *Philosophical Topics*, 12(1), 93–120.
- Silvanto, J., Cowey, A., Lavie, N. & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience*, 8(2), 143–44.
- Stetson C, Cui X, Montague PR, Eagleman DM (2006). Motor-sensory recalibration leads to an illusory reversal of action and sensation. *Neuron*, 51, 651–659.
- Tse, P. U., Rivest, J., Intriligator, J. & Cavanagh, P. (2004). Attention and the subjective expansion of time. *Perception & Psychophysics*, 66(7), 1171–1189.
- Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.
- Ullman, S. (1979). The interpretation of structure from motion. Proceedings of the Royal Society of London. Series B. Biological Sciences, 203(1153), 405–426.
- Vroomen, J., Keetels, M., de Gelder, B., & Bertelson, P. (2004). Recalibration of temporal order perception by exposure to audio-visual asynchrony. *Cognitive Brain Research*, 22(1), 32–35.
- Watzl, S. (2012). Silencing the experience of change. *Philosophical Studies*, 1–24.

- Wearden, J. H., Todd, N. P. M., & Jones, L. A. (2006). When do auditory/visual differences in duration judgments occur? *The Quarterly Journal of Experimental Psychology*, 59(10), 1709–1724.
- Yarrow, K., Haggard, P., & Rothwell, J. C. (2004). Action, arousal, and subjective time. *Consciousness and Cognition*, 13(2), 373–390.
- Zeki S (1993). *A Vision of the Brain*. Blackwell.
- Zihl, J., von Cramon, D., & Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain*, 106, 313–340.