8

Extensionalism, Atomism, and Continuity

Geoffrey Lee

Introduction: Extensionalism and atomism

A *temporal experience* is an experience that presents to its subject states of affairs that manifestly involve duration and change over time, such as the temporal order of sounds, the velocity of moving objects, or the duration of a brief flash of light in the visual field. There is a disagreement about temporal experiences between extensionalists, atomists, and snapshot theorists, which will be my subject here.

According to extensionalists like Barry Dainton, temporal experiences are themselves temporally extended processes that play out over time, having experiential stages that mirror the stages of the presented events (see Figure 8.1). For example, an extensionalist will typically think that experiencing the temporal order of two sounds S1 and S2 involves first experiencing S1, then experiencing S2, and a suitable relation holding between these experiences.

By contrast, atomists think that temporal experiences do not themselves have temporal structure that mirrors the temporal structure represented in their content. On this view, an experience of a temporallly structured state of affairs like the temporal order of two sounds is itself temporally unstructured, in that it does not have distinct experiences as proper temporal parts; for example, a temporal order experience does not involve first experiencing one event and *then* experiencing a second event (see Figure 8.2); it involves experiencing both events *at the same time*, even though they are experienced *as* happening at different times. For this reason, atomic temporal experiences might even be instantaneous, although atomists needn't think this, as I will explain below.

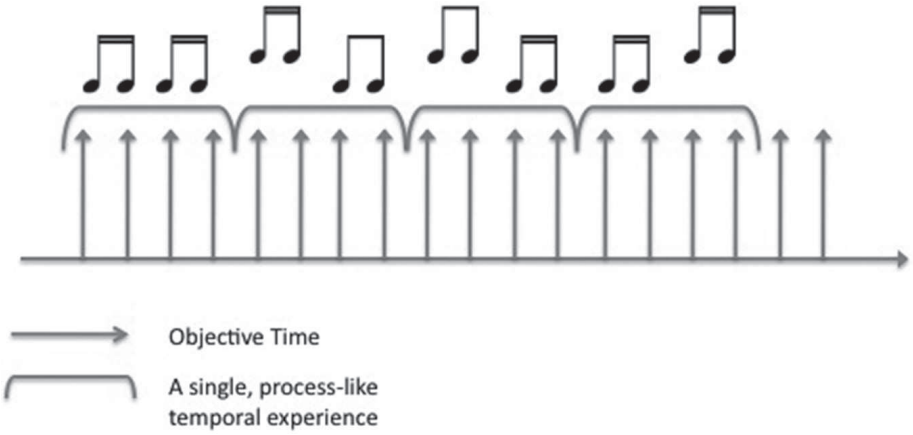

FIGURE 8.1 *The Extensionalist View*

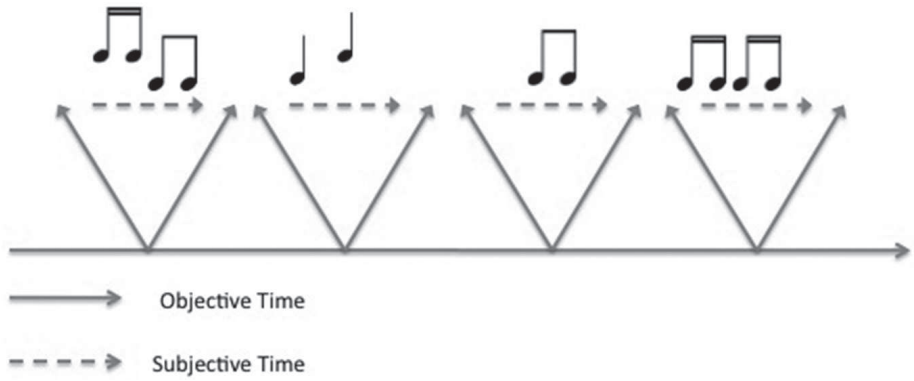

FIGURE 8.2 *The Atomist View*

Extensionalists and atomists both disagree with snapshot theorists, who deny that we have temporal experiences. On their view, the stream of consciousness is a series of photo-like presentations of properties of objects like their shape, spatial arrangement, surface texture, color or illumination, presentations that do not concern how the world is changing over time. In so far as we are able to make accurate judgments about temporal features, this involves a post-experiential process of comparing the snapshots, a process that doesn't require actually experiencing the temporal features. I will set aside the snapshot view here (for a defense of it, see Chuard 2011), as it is incompatible with obvious phenomenological data. Take auditory experience, for example—most of the information in it pertains to how events are organized over time; it is very hard to make sense of auditory experience without temporal content.

My aim in this chapter is to defend atomism, by responding to a number of arguments against it that appear in the literature, particularly those given by Dainton (2000, 2010) and Phillips (2010). Phillips (2010) argues that

atomism is incompatible with the kind of introspective knowledge we have of how experience is changing over time. Dainton argues that atomists can't adequately capture the sense in which the stream of consciousness is continuous and connected over time. I will respond to these arguments, among others. Elsewhere Lee (2014a) I give a detailed positive argument *for* atomism and *against* extensionalism, which I will also briefly recap here.

Lee (2014a) also contains a detailed discussion of what exactly is at issue between extensionalists and atomists. Here I will state more briefly what I see as the main issues.

It is helpful to distinguish three closely related but distinct ideas that extensionalists may or may not subscribe to. In increasing order of strength, these are:

The process view: Temporal experiences are experiential processes that unfold over time by having shorter experiences as proper temporal parts.[1]

The mirroring view: The process view, with the additional claim that the experiential temporal parts of a temporal experience are arranged in time in a way that mirrors or matches the temporal structure represented in their content.

The representation by resemblance view: The mirroring view, with the additional claim that the correspondence between the objective temporal structure of experience and the temporal structure in the experience's content is *explained* by the fact that temporal properties are represented by resemblance—time is used to represent time.

I treat experiences as instantiations of experiential properties by subjects (or parts of subjects, like brain areas). So an extensional temporal experience is a process involving distinct property instantiations at different times. In many cases, these will be *qualitatively* distinct, involving different experiential properties, for example when you experience an object moving or changing illumination, or hear the temporal relations between different sounds. However, we can also have distinct instantiations of the *same* experiential property as different temporal stages; for example, this is presumably what the extensionalist will say is happening when you have a qualitatively constant experience as of an unchanging scene.

Not any old series of experiences can make an extensional temporal experience. The process theorist thinks that the stages of a temporal experience have to be unified, or "co-conscious," in order to form an experience. So, for example, if I experience one event, then undergo a complete brain reset, and then experience another event, I will not have a unified experience as of

the temporal order of the events. What the relevant kind of "unity" is here is a difficult question. Dainton assumes that this diachronic unity relation is (at least) the same unity relation that holds synchronically between different parts of a subject's field of awareness: they are "experienced together" in a way that two random experiences enjoyed by different subjects are not. It is controversial even in the synchronic case what the unity relation is (or even whether there is a single relation here to focus on). I will grant Dainton that there is such a relation. The fact that in his view it holds diachronically is supposed to be a significant advantage over other views—I will discuss this in detail below.

The idea that temporal experiences are experiential processes is usually held in conjunction with a commitment to the mirroring view (at least a weak version of it)—indeed I would interpret extensionalists as thinking that the process view is true *because* mirroring obtains. Mirroring comes in various strengths. Extensionalists typically at least believe in topological mirroring—the view that experiencing A as happening before B requires first experiencing A and then B. One could also subscribe to the stronger Metrical mirroring thesis, that experiences of duration and rates of change are mirrored by the duration relations between the parts of the experience itself, so that, for example, an experience as of an event lasting 1 second itself lasts 1 second. Dainton himself rejects metrical mirroring; a process theorist who accepts mirroring for all experienced temporal features is Phillips (2010), although his position on metrical content is qualified in his more recent work (Phillips 2013).

The mirroring thesis is itself neutral on why mirroring obtains. If mirroring obtains, one possible explanation is that the mirrored temporal features are represented through *resemblance*—time in experience is represented by time itself. Recent defenders of extensionalism—including Dainton—do *not* in fact offer this as the explanation, so we probably shouldn't assume that representation by resemblance is part of the view, even though it is a very natural extension of it. This won't matter here, as I will be concerned only with the weaker claim that temporal experiences are process-like.

"Atomism" is not a term used by Dainton in his influential taxonomies of views in this area. I define atomism as the view that temporal experiences are *never* process-like, a view which implies that the process view is false, and therefore by implication, so are the mirroring and resemblance views. Atomism needs to be understood carefully. For one thing, atomic experiences needn't be instantaneous events. They might be realized by extended physical processes in the brain—indeed I would argue that all experiences are extended in this sense, because experiences require extended processes like neural firings in order to exist. Atomic experiences are atomic in the sense that they do not contain shorter *experiences* as temporal parts. If an atomic experience is realized by an

extended process (like 40hz neural firing), then this process may have shorter *physical events* as proper temporal parts, even if it doesn't have *experiences* as proper temporal parts; for example, these shorter physical events may not be sufficient for any experiences to exist. To be clear, if an experiential property instantiation occurs fundamentally over a short interval in this way, we can say that it has shorter experiential temporal parts in a *derivative* sense. For example, if I am feeling a certain painful sensation over interval I, there is a derivative sense in which I am feeling it at each time during I. The atomist can admit that experiences have such derivative temporal parts. The process theorist thinks that temporal experiences have temporal parts in a stronger sense than this. This is obvious when the temporal parts are qualitatively varied, but even if they are not, the idea of the process view is that experiential stages of temporal experiences are distinct property instantiations—for example, they are realized by different physical events happening at different times.

Atomism encompasses a number of different views that Dainton contrasts with extensionalism. Dainton himself focuses attention on "retensionalists," who think that atomic temporal experiences have a complex structure involving memory-like retentional experiences, direct perceptual experience of the present, and possibly also a "protentional" anticipatory awareness of what is about to occur (Husserl is the most famous proponent). Atomists needn't think that temporal experiences have this tripartite structure, however. Dainton rightly complains that we have awareness of temporal properties that is just as immediate as our awareness of "static" features like shape and color, and that the retentional view can't account for this fact. A better version of atomism would hold that there is a single kind of perceptual experience that presents both non-temporal features *and* facts about how these features are changing over time. Let us call this "non-retentional atomism." Dainton does acknowledge the possibility of such a view (versions of which have been defended by Broad 1925 and Grush 2005), but it doesn't fall very neatly in his taxonomy. To my mind, it is the most promising competitor to extensionalism.

Defending atomism

A full defense of atomism would involve a positive case against the process view and a defense of atomism against objections. Here I will focus more on objections against the view, although it will be worthwhile to first briefly describe the positive case for it (see Lee 2014 for more detail).

The reason why atomism is true is that temporal experiences are realized by physical events in the brain that do not code temporal information in a way that could realize an extensional experience; in particular, the neural

realization of a temporal experience does not have distinct temporal stages corresponding to the allegedly temporally separated stages of the experience itself. Rather, even if an experience involves being presented with events *as* happening at different times and standing in certain temporal relations, the experiences of these different events are each realized by neural events that happen at the *same* time (or over the same short interval). We can infer from this that the different parts of the temporal experience themselves happen at the same time, contrary to the process view.

Why think that the different parts of a temporal experience are realized at the same time? Information at the periphery of the perceptual system—for example, on the retina—is typically represented in a form that *is* process-like. For example, if A happens before B, then typically light from A will stimulate the retina, and *then* light from B will stimulate the retina (although due to different latencies in transmission from different sources, there is not always such a neat correspondence between arrival time and transmission time, something that the system is capable of correcting for, at least to a limited extent). Theories of how temporal information is extracted from such inputs typically assume that the point is to compare or integrate the different stages of the input to get an explicit representation of temporal information that is available "all at once." For example, the retinal stimulation from event A might leave a trace which is simultaneously compared with a trace from the retinal stimulation from event B, to get a simultaneous representation of "A before B." This might in turn enable the subject to act on the basis of this temporal information, for example by correctly pressing the "A before B" button, rather than the "B before A" button.

The reason such simultaneous integration seems necessary is that without it temporal information is not in a form that could cause appropriate effects downstream of experience, such as verbally reporting the experience, putting the information in memory, or pressing the appropriate button. The fact that the content of experience is typically accessible in this way is the main reason for thinking that temporal experiences are realized by the output of such a process of simultaneous "trace integration," rather than by earlier events in the processing stream, such as those on retina, which do in some sense "code time by time." If this is right, then we get the result that these experiences are not structured in time as the process theorist alleges, but rather represent temporal information all at once, as the atomist thinks.

A fuller explanation and defense of this argument is given in Lee (2014a). I will not pursue it here, instead focusing on the negative arguments that might be given against atomism. Whereas the Trace Integration argument is an empirical argument, these negative arguments tend to be based on phenomenological observation and philosophical considerations. Although such considerations *are* relevant to the debate at certain points, ultimately

it seems to me that a debate about the timing of experience has to take into account empirical evidence pertaining to the timing of the neural events that underpin experience, and that these considerations have far stronger weight than any phenomenological or philosophical considerations that point in a different direction. That said, even if we have independent reason for skepticism about the weight of such considerations, a satisfactory defense of atomism should have something to say about them, and saying it is my aim here.

The simultaneity argument and the multiple presentations argument

One argument against atomism is that if temporally separated events appear in your experience all at the same time, then they will appear to be happening at the same time—for example, tones occurring in experience at a single time will form a chord, or the different positions of a moving object would appear smeared through space. This is contrary to the atomist's claim that one can perceive, at a single time, consecutive events *as* happening at different times. This is obviously a question-begging objection, however—the whole point of atomism is that simultaneously presented events need *not* appear as simultaneous. This only seems problematic if one assumes a mirroring constraint, but the atomist rejects such constraints.

Another easily deflected objection is the "multiple presentations" objection. Since the content of atomic experiences covers a temporally extended portion of what is happening in the world, a single event may appear in the content of a series of consecutive experiences in the stream, not just a single experience. This is supposed to be counterintuitive because we do not seem to experience the event over and over again. However, all that "multiple presentations" really amounts to is the claim that you experience the event *for an extended period of time*, longer than the duration of a "single experience" (which presumably is the duration of the minimal amount of neural activity sufficient for an experience to exist). There is nothing counterintuitive about this. It is true that the *onset* of an event may be presented for more than a single moment, but this does not mean that the event will seem to you to be starting over and over again; all the experiences you have present it as starting only once (Tye (2003) makes the same point in response to this objection.)

Phillips (2010) thinks that this response to the multiple presentations objection—and atomism in general—requires we lack introspective access to how our experience is changing over time—in this case, because we are not aware of enjoying more than one experience of the onset of an event.

He thinks that we have introspective awareness of the temporal layout of experience in a stronger sense than this allows for, and that introspection reveals that temporal experience does not in fact have an atomic structure, but rather is process-like. Let's look at this argument.

Phillips' transparency argument

Phillips thinks that we have introspective access to the temporal relations between our experiences (I'll call this the "introspectibility premise"), but also that experience is "transparent": introspection of perceptual experience only reveals the properties that external events are experienced as having, not any psychological properties (the "transparency premise"). He thinks that we can reconcile these two data by holding that we introspect the temporal features of experience *by* introspecting the apparent temporal layout of external events (the "reflection premise"). Thus, an experience as of A happening before B *also* is an appearance that the *experience of* A happened before the *experience of* B. Furthermore, he takes it as a premise that it is part of the nature of experience that if experience appears a certain way to me, then it must be that way (the "infallibility premise"). So if it seems introspectively as if my experience of A happens before my experience of B, then these experiences really must be related in this way.

Clearly, if the apparent temporal layout of the world is an infallible guide to the temporal layout of experience itself, then a strong form of mirroring (and hence a process view) must be correct. Phillips' argument is valid; atomists must reject one of the premises.

The atomist could deny that *by* presenting a certain temporal layout, an experience automatically seems to have a corresponding temporal layout (the reflection premise); for example, they could deny that an experience *of* temporal order automatically seems to have a corresponding temporal order. An analogy with the spatial features of experience is relevant. If physicalism is true, then experiences are physical events happening in spacetime and therefore have spatio-temporal features like all other physical events—for example, at a given time, an experience will be occurring in a region of space with a certain *size* (for example a region of a brain). These spatial features of experiences are obviously not made available to you by introspecting the apparent spatial layout of the world; for example, an experience of a teacup does not in any sense appear to occupy a teacup-shaped region. The atomist could say something similar about temporal experience.

The spatial features of experience are not available to introspection *at all* (this is presumably why people are sometimes tempted to say that experiences do not happen in physical space). It seems too strong to make the

analogous claim that we have *no* introspective sense of our experiences as events happening in time. We talk about the "stream of consciousness," a description that seems apt given the way experience seems from the inside. Kant described time as the "form of inner sense," again giving voice to the intuition that we are introspectively aware of experiences as happening in time. Admittedly, if the atomist denies the reflection premise, they must give an alternative account of how we are aware of experiences as playing out in time, or at least explain away the strong appearance that we have such awareness.

Consider the option of denying *any* introspective awareness of experience changing in time. It might be that at any given moment you are simply enjoying the atomic experience you are having, and have no introspective access at all to the atomic experiences that were happening at other times. If you try to introspect the temporal relations between experiences, you will only be able to attend to the experienced temporal relations between events presented by your current atomic experience. Perhaps there is a tendency to confuse these temporal relations with those that hold between experiences themselves, giving us the mistaken impression that we are aware of experience itself changing, when really we are only aware of the external world changing.

Although I think Phillips has not done enough to disarm this view, and that the view is not obviously false, I also think it may be too strong. For one thing, instead of denying introspectibility, the atomist could have a view that rejects transparency, or a view that accepts reflection but denies infallibility; if either of these options can be made to work, they will allow for introspective awareness of experience changing in time.

Consider the latter option. The atomist could say that, after all, we *can* become aware of how experience is changing by attending to worldly changes (the reflection premise), but that the argument fails because apparent external changes are not an *infallible* guide to experiential changes. For example, it may *very often* be true that if you experience A as before B, then you had an experience of A before you had any experience of B, because A was detected before B, and the information that A occurred was available before the information that B occurred was. This is something that is perfectly consistent with atomism (although if atomism is true, your experience of B may be accompanied by a further experience of A). There may be exceptions to this rule of thumb,[2] but if it is usually true, then we can use temporal appearances as a *rough* guide to the temporal structure of experience, gaining knowledge of experiential structure that is consistent with atomism.

Phillips will reply that this alleged status as a rough rule of thumb is incompatible with the principle that there is no appearance–reality distinction for

experience. But in so far as there is a correct principle of this kind, it cannot apply to *all* the properties of experience. Take spatial properties again: it might be that experiences as of different regions of space are typically located in distinct neural regions. If this were right, we could, after all, use the apparent spatial properties of things as a reliable, but not infallible, guide to certain spatial properties of experiences themselves. This obviously wouldn't conflict with whatever the truth is behind a "no appearance–reality distinction" principle.

The other way the atomist could block the argument without rejecting introspectibility would be to reject transparency. If transparency fails, then even if we aren't aware of the temporal features of experience via reflection, the implausible conclusion that we have no awareness of experience changing over time does not follow. In fact, a pretty strong case can be made that transparency fails in a way that is relevant here: we can be aware of our experience changing in a way that does not go via awareness of external changes. Consider shifting your attention from one object to another. These could be covert attention shifts that do not require moving your eyes or other body parts. And the scene you are looking at might not be changing at all. Still, you experience a change as happening—a *psychological* change, not an external change. Note that it is not merely *that your experience changes* when you shift your attention: you have an experience *of* change happening. You can even be aware of certain features of this change. For example, if you shift attention back and forth between two objects at a certain *rate*, then you can experience the rate that this is happening, just as if you perceive a light turning on and off at a certain rate, you can be aware of this rate. The proponent of transparency makes the mistake of thinking that perception is a kind of bare confrontation with the world. However, perception involves actively shifting your attention around the environment, and an awareness of this shifting contributes to what it feels like to consciously perceive. We might compare such awareness of psychological changes with the awareness of bodily changes that can occur during touch. We can conceive of an analogue of touch that is more "transparent" than our actual sense of touch: a creature might get information about the surfaces of objects by touching them, but the resulting experiences might present the properties of the surfaces without ever making the subject consciously aware of the process of exploring the surface of the object with a part of the body; such bodily interactions with an object might be controlled entirely sub-personally. Similarly, visual awareness could have been designed to make you aware of the properties of external things without any experienced sense of visually exploring the world through controlled changes in attention, but this is not the experience we actually have.

Even if we are aware of attentional changes in this transparency-violating way (the idea deserves further discussion, but I will not pursue it here), it is

not clear that they *exhaust* our awareness of changes in experience. Below I will consider another kind of experience we have—the experience of a constant flowing change in our temporal perspective—that is arguably also inconsistent with transparency, and that might constitute another way in which we are aware of experience changing over time (these experiences are also part of my response to Dainton, which is why I will delay discussing until later). Furthermore, as I already mentioned, it may be that there is something right about the reflection premise: perhaps we can be aware of experience changing, albeit in a fallible way, just by attending to external changes. I think it is plausible that postulating these various forms of awareness of experiential changes is enough to explain our introspective sense of a "stream of consciousness" and the fact that we seem to be aware of the temporal features of experience in a way that we are *not* aware of its spatial features; no appeal to an infallible awareness of the kind Phillips postulates is necessary.

To sum up, Phillips has not made a compelling case that external temporal appearances are an infallible guide to the temporal structure of experiences itself: there are a number of plausible moves an atomist can make in response. They can deny introspectibility, holding that we mistake awareness of external changes for awareness of internal changes. Or they can deny infallibility, holding that awareness of external changes is a fallible, not infallible, guide to internal changes. Finally, there is a good case to be made that transparency is false: we can be aware of experience changing in a way that does not depend on awareness of external changes (as I said, I will discuss another possible counter-example below). Even though I think that Phillips' argument can be resisted, I think it has the great merit of emphasizing that there is a real puzzle accounting for the sense in which we are aware of experience changing over time: it is really not obvious how, if at all, this happens (for example, I am suggesting in this chapter that we should believe in a kind of temporal introspection that violates transparency, but I will not be able to give a full account of how exactly it is that such inner awareness exists).

Having looked at Phillips' case against atomism, let us now move on to consider the reasons Dainton gives for rejecting it.

Dainton's continuity argument

Dainton argues that atomism cannot capture the sense in which experience is continuous through time, whereas extensionalism is tailor-made to account for this. The idea that experience is "continuous" is extremely ambiguous, and to understand his argument properly, we need to first think about the different meanings that we could attach to this.

Experience could be continuous in the sense that it does not have any gaps in time. If time itself is continuous, this means that a section of the stream fills a continuum of moments from its beginning to end. This is weaker than saying that experience is continuous in the sense that it *changes continuously* (i.e., it can be represented as a continuous function from real numbers representing times to sets of real numbers representing experience states). For example, experience might change discretely, even if it is non-gappy. It is also weaker than saying that experiences are instantaneous (not just that they exist derivatively at instants), and that these instantaneous experiences form a continuum in time. The stream could fill a continuum of time, without itself forming a precise continuum in this way. For example, imagine a pain whose intensity is realized by the firing rate of a single neuron (obviously in reality many neurons would be involved). Each intensity is realized over a period of time (because a firing rate is an average calculated over a period of time); nonetheless, the firing rate, and hence the intensity of the pain, can increase gradually over time. In this way, the temporal parts of an experience might be all temporally extended, but overlap in time in such a way as to form a non-gappy covering of all the instants in a certain interval of continuous time, without precisely occupying any of these instants.

It is an interesting question as to whether experience is continuous in any of these senses (for some relevant empirical literature, see Van Rullen and Koch, C. 2003; Van Rullen et al. 2011; Kline et al. 2004; Mathewson et al. 2009, 2011). Atomism is consistent with continuity in any of these senses, but also with strong discontinuity—for example, it is consistent with experiences coming in discrete bursts that are separated in time from each other. Some may think that it is obvious from introspection that consciousness is *not* gappy in this way. However, it is not clear that a gappy structure would be apparent introspectively; awareness of it would seem to require a higher-order quasi-perceptual monitoring of first-order experiences that is set up to be sensitive to such gaps, and we may lack any such capacity. Furthermore, the continuity intuition can be explained away as confusing the fact that we lack an awareness of gaps, with our having an awareness *of* a lack of gaps. The moral is that to discover whether experience is discretely gappy, we need to figure out empirically what the temporal structure of neural states underwriting experience is, not search for gaps introspectively.[3]

Dainton's objection to atomism is not that it can't accommodate continuity in any of the above senses. The kind of continuity that he is interested in (he calls it "strong continuity") involves a feature of experience that we allegedly experience from the inside. It is a little hard to pin down what it is exactly, but it involves a sense of the interconnectedness of experience over time, of one experience flowing into the next—he expresses this by saying that consecutive experiences seem to be *experientially connected*. He quotes William James as one source of the idea:

> My experiences and your experiences are "with" each other in various external ways, but mine pass into mine, and yours pass into yours in a way in which yours and mine never pass into one another. (James 1912: 47)
>
> … the parts of experience hold together from next to next by relations that are themselves parts of experience. The directly apprehended universe needs, in short, no extraneous trans-empirical connective support, but possesses in its own right a concatenated or continuous structure. (James 1912: xii)

One example Dainton gives to try to make this vivid is the case of hearing a long continuous tone; one's current experience seems to flow directly from previous experiences of earlier stages of the tone. Or when you apprehend a melody, your experience of one sound or a portion of melody seems to flow seamlessly into your apprehension of the next part of the melody.

Allegedly this strong continuity can easily be accounted for in Dainton's version of extensionalism. He postulates diachronic unity relations linking together one moment of experience with the next. He thinks the unity relation holds between all experiential stages that are sufficiently close together in time, so that if we consider maximal groups of mutually co-conscious experience-stages (i.e., "total experiences"), we will find that these form overlapping blocks in time. Hence he calls his version of extensionalism "the overlap model" (see Figure 8.3).

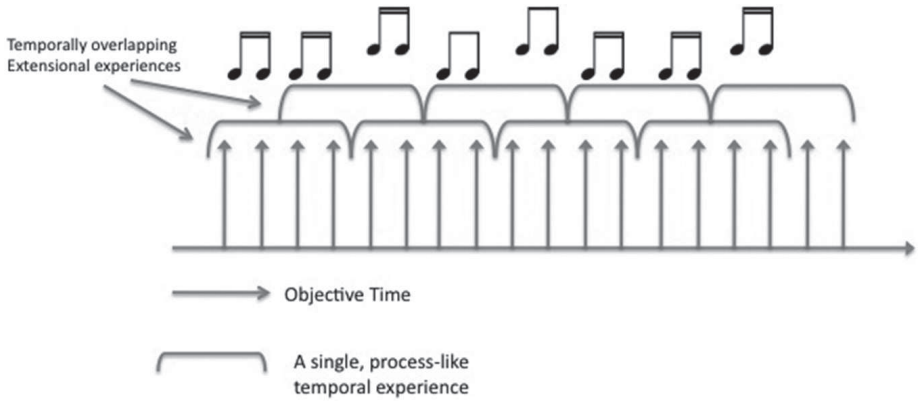

FIGURE 8.3 *Extensionalism with Overlap*

Supposedly the fact that total experiences overlap by sharing experiential parts, and the fact that the whole stream is interconnected by diachronic unity, gives us an explanation of the felt sense of continuity, connection, and flow within conscious experience. On the other hand, on the atomic view, consecutive experiences within the stream of consciousness are not connected together by unity—they do not form larger experiential wholes

(although see the discussion below). In this sense, there is a kind of independence between your experience at one moment and the very next moment for the atomist. Dainton thinks that this is implausible. One objection is phenomenological—it doesn't capture the felt flowing connectedness of experience. Another is more metaphysical: if successive experiences are independent in this way, then they are no more connected than experiences in completely different streams of consciousness:

> In the absence of experienced transitions between pulses, successive pulses might as well belong in entirely different streams of consciousness. (Dainton 2004: 17)

> … individual acts are totally isolated from one another. From a purely experiential perspective, the successive phases of our streams of consciousness might as well exist in different universes. (Dainton 2004: 21)

One way that the atomist could respond here is by saying that it is obscure what Jamesian experiential flow is supposed to be, and so it is unclear what exactly the atomist is allegedly failing to account for; pending clarification, there is no real objection here. I think they can do better than this, however. There *is* something intuitive about the idea of experiential flow, and there are fairly well-defined ways to elaborate an atomist view in order to accommodate it. Moreover, the atomist can challenge whether overlap extensionalism really does itself accommodate the relevant phenomena here. Finally, the atomist can also explain why their atoms are not "totally isolated" from one another in an implausible way.

Let us discuss each of these points, beginning by addressing the metaphysical worry. According to Dainton, consecutive atomic experiences "might as well be in completely different universes." I find it hard to see a compelling worry here. The atomist should reply by detailing the ways in which experiential stages belonging to a single stream may well be intimately connected with each other in the atomic view, connections that won't hold between experiences in different subjects, or in completely different universes. As we will see, the atomist can even acknowledge a limited kind of experiential unity across times.

First, what one experiences at a particular moment in time may well be *causally* relevant to what one experiences a moment later, so the stages of an atomist's stream are at least bound by causation. To be clear, I do not think that there is any *a priori* reason to think that there are direct causal relations between experience stages. For all we know *a priori*, there is no top-down feedback within perceptual processing, so that consecutive experiences can be appropriately compared with consecutive images projected on a screen

by a movie projector, which are not directly causally related to each other, despite having correlated properties. In the experience case, if the "images" are close enough in time, they may have overlapping causes at the input end (because, as I mentioned above, perceptual processing integrates information that arrives over an extended interval), but unless the output of one chain of perceptual processing feeds back into the next chain, consecutive experiential "outputs" will not be directly causally related.

As it happens though, there is empirical evidence that there *are* such feedback connections in perceptual processing. One reason for thinking this is that there are neural connections feeding back within many different stages of perceptual processing (see, for example, Lamme et al. 1998). Another is that computational models that assume such feedback give good predictions. Grush (2005) describes one such model in giving his "emulation theory," in which perceptual and motor control systems estimate the next state of the world by comparing current input with a prediction of how the world will evolve (based on previous estimations of the state of the world and reference copies of current motor instructions), and creates a new model by computing a weighted average between the prediction and the input-based estimate. It is as if visual experience has a natural path that it will follow on its own (like a ball rolling down a hill), and the role of external input is as an external force correcting the direction it travels. If a model like this is correct, then it is plausible that there are feedback relations causally linking adjacent stages of an atomic stream.

Another kind of intimate connection that might exist between temporally adjacent atomic experiences is *overlapping realization*. If an experience is neurally realized by temporally extended neural activity, then another experience close enough in time might be partly realized by some of the same neural activity. Consider again the example of the intensity of a pain being realized by the firing rate of a neuron. Because temporally adjacent firing rates might be realized by some of the same neural firings, temporally adjacent experiences of different intensity might be realized by some of the same neural activity. Thus two experiences with different intensities may "overlap" in the sense that they have overlapping realizers, even if they don't overlap by sharing experiential parts (the overlapping neural activity may be too brief to realize any experiences). In this way, stages of a stream of atomic experiences that are close enough in time may not be completely metaphysically independent events.

Dainton's response is likely to be that the existence of both causal and realization connections between stages is not the same as the existence of the kind of *experiential* connections that he thinks exist, and which the atomist allegedly can't account for. In particular, he thinks we experience the transition from one experience into another, in virtue of diachronic unity

relations holding between them. I think the points I have made so far at least show that different atoms are intimately connected in ways that make the "different universes" complaint at least somewhat misleading. I now want to argue that (1) there *is* a limited kind of diachronic unity between experience that the atomist can acknowledge, and (2) that the atomist can explain the sense of continuity that Dainton is gesturing at.

With regard to (1), recall that above I said that atomists probably ought not to hold that their atomic experiences are instantaneous—they are best off saying that they are extended in time, because they are realized by extended processes like neural firings. If all experiences are extended in this way, this makes it tricky to draw a sharp line between synchronic unity and diachronic unity, suggesting a form of atomism in which synchronic unity inevitably gives us a form of diachronic unity also. For example, suppose we believe, as theorists like Bayne (2010) do, that synchronic experiences belonging to a single subject at a single time are typically "unified" into a single field of awareness (a weak version of what Bayne calls the "unity thesis"). If experiences are extended in time, the "at a single time" constraint is rather fuzzy. How are we to interpret it? As meaning *overlapping* in time, or something else?

Here is one way in which the unity thesis might work out in an atomic view, with synchrony interpreted as temporal overlap (see Figure 8.4). Consider modality-specific sub-streams, such as visual experience and auditory experience (I will assume that experience can be divided up this way). Suppose it is usually the case that (1) visual experiences which overlap in time with auditory experiences are unified with them, but (2) that consecutive visual experiences (which may overlap in time, in virtue of being realized over overlapping temporal intervals, as in the pain intensity/firing rate case discussed above) are not unified. (1) is a kind of unity thesis, and (2) is plausible, because the continuous updating of visual experience is likely to mean that adjacent atomic visual experiences have incompatible contents, even if they overlap in time.

Given this set-up, it is possible that a particular auditory experience, such as A2 in Figure 8.4, will be unified with consecutive stages of visual experience (V1,V2, and V3) that are not themselves unified and which do not overlap by sharing experiential parts (even though they overlap in time).

If a "simultaneous total experience" at time t is a group of mutually unified experiences that all happen during a period that *includes* t, then the situation depicted would involve A2 being contained in three different total experiences that occur over different intervals, total experiences that contain different experiential elements (A2+V1, A2+V2, A2+V3). This shows how, even in an atomic view, total experiences at slightly different times can overlap by sharing experiential parts (the overlap is quite different from

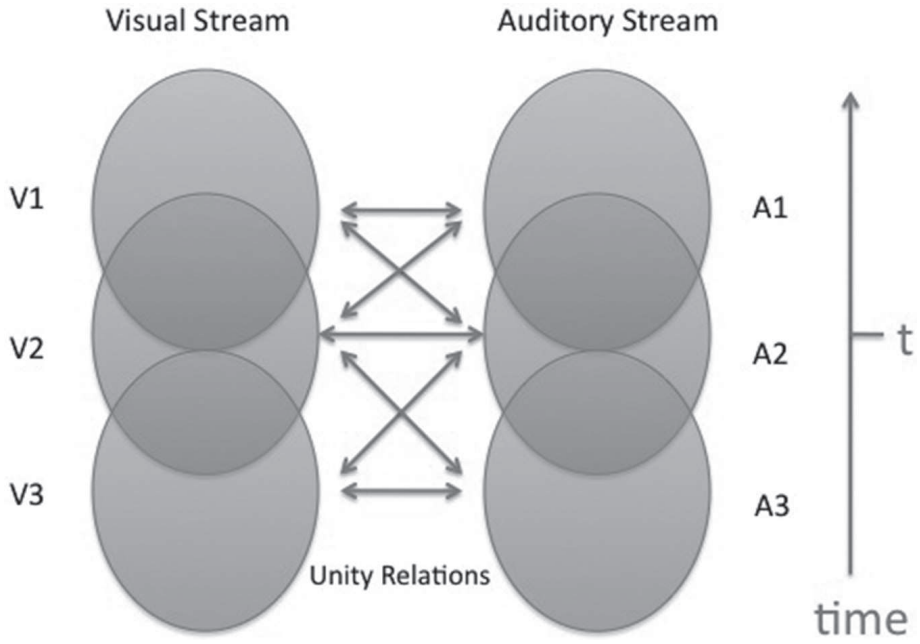

FIGURE 8.4

the kind postulated in Dainton's overlap model, however). By implication, the momentary stages of an atomic stream of consciousness may not be experientially isolated islands, but may be connected to each other by a criss-crossing web of unity relations. That said, I should stress that the atomic view is *also* compatible with a lack of any such diachronic connections, and further empirical evidence may reveal that they do not exist. For example, if your total experience is a series of discrete pulses, there may be no overlapping realization or unity connections between stages.

I don't want to suggest that I think that these diachronic unity relations, if they exist, would explain our sense of one experience flowing into the next. I don't think they would. But then again I don't think Dainton's diachronic unity relations would explain the sense of flow either. I bring them up simply to point out that even on an atomic view, there can be experiential connections between stages that are similar to those that appear in the extensionalist view.

So, let us now turn to the question as to whether the atomist can explain the sense of "flow" in experience. Before getting to that, one point that an atomist should make here is that there are reasons in advance to doubt whether the extensionalist can do any better than them in capturing this extra sense of "flow." Consider a total experience that involves apprehending, for example, a section of music, including the temporal relations between various sounds. The atomist and the extensionalist can agree about what the overall content of this total experience is, and agree that it has as parts experiences

of individual sounds and their relations, and agree that these parts are *unified* as part of the complex whole. The main difference between the theories involves the *temporal arrangement* of these parts (this is especially clear in Dainton's version, because he thinks that the unity relation that holds diachronically is the same as the one that holds synchronically). Since, aside from this temporal difference, atomic total experiences can be isomorphic to extensional ones, it is not at all clear that there is any phenomenal fact that can be explained by extensionalists but not by atomists. So if a "feeling of flow" is not accounted for in the atomic view, this suggests that it is *also* left out in the overlap model[4].

Dainton's view must be that the mere fact that, in the overlap model, the relevant parts of temporal experience happen sequentially in time rather than at a single time, gives a feeling of flow from one part of the experience to the next, which would otherwise be absent. But this is problematic. First, it is not at all clear why having a sequential rather than synchronic arrangement in time would make any phenomenal difference at all, let alone why it would produce a sense of "flow" (when we consider the kinds of features that *would* explain "flow" we will see that this is especially clear). Second, whatever phenomenal role is played by these temporal relations in Dainton's model has an analogue in the atomic view in the temporal *content* of an atomic experience. For example, the temporal order of the parts of an extensional experience has an analogue in the presentation *of* temporal order in an atomic experience. So if the mere temporal arrangement of an extensional experience can explain a sense of "flow," it is unclear why the corresponding temporal content of an atomic experience cannot do the same work.[5]

So, it is unclear that extensionalists do any better than atomists in explaining the sense of "flow." What would explain it? It may be that it can be fully explained in terms of our awareness of changes in external events, in ways that I will describe immediately below. But if not, I would suggest that this is because there is a kind of awareness of our experiential perspective changing over time that is inconsistent with temporal transparency: I will try to describe what I mean by this in more detail below.

As just mentioned, I think our experience of the continuity of external events and processes is probably at least one source of the sense of "flow." External events may be experienced *as* continuing on from before we apprehended them, or at least we may *lack* an experience of them as starting when *we* start to experience them. Take the simple example of experiencing a long continuous tone. Your auditory system—indeed your perceptual system in general—is designed to detect discontinuities in stimuli, and make them perceptually salient to you. There is therefore a big perceptual difference between experiencing a sound *as* starting, and experiencing a sound without experiencing it as starting. There may also be such a thing as positively

experiencing the sound *as* continuing on from before. All three of these can easily be accommodated by an atomist, because they involve different conditions under which experience is veridical, and there is no reason why an atomic experience cannot have the relevant veridicality condition.

The tone example is not an isolated curiosity—we constantly perceive processes that are not bounded in time at the point of apprehension. Much of ordinary experience is of activities and processes that are already in motion; more generally, even the unchanging state of a boring material object like a teacup or muffin is an example of something continuing on from before that we either do not perceive as bounded, or positively sense as continuing on.

Interestingly, it is not actually clear how an extensionalist will capture the distinctions we are talking about here. Suppose, for example, that I have a "total experience"—a maximally unified experience—that presents a sound as continuing on from before. What is the difference between this and a case where the sound is heard as bounded in time? Dainton's idea seems to be that it is the fact that the beginning of the total experience is unified with a prior experience of sound, rather than an experience of silence, that explains the appearance of continuity. But this connection with earlier parts of the stream is actually an *extrinsic* property of the total experience, whereas the appearance of continuity is surely an *intrinsic* property of the experience, having to do with how it presents a current sound as related to the past. This suggests that the atomist may actually do *better* than the extensionalist in explaining our experience of continuity.

Whether or not they have a response to this last point, the extensionalist might object that Jamesian continuity is a quite different phenomenon from an experience of temporal continuity in the world. Jamesian continuity is a matter of feeling the flow from one *experience* to another, not from one external state of the world to another. It is this that the atomist cannot capture.

I suspect that theorists like Michael Tye, who often press the idea that experience is transparent to introspection (for example, Tye 2002), will be inclined to deny the phenomenological intuition here. They will say that introspection only reveals how the world appears to be arranged, and so any sensed flow or continuity must be an apparent feature of external events. We have already discussed transparency, and noted that it may be subject to counter-examples. I want to end by considering an additional kind of awareness we have that would fit Dainton's job description of an experience of flow very nicely, and is an additional counter-example to temporal transparency. As we will see, "flow" in this sense is perfectly compatible with atomism.

The relevant "flow" is best described as a sense of a constant change in your temporal perspective, that is, of which events are presented by

experience as being in the present. It is related to the idea that we are aware of time itself passing—not just in the sense that we are aware of temporal relations between external events, such as temporal order and distance relations, but in the sense that we are aware of a constant flowing change in what is in *the present*. If there is an experience of this kind, it means that there is a sense in which we "experience time" that does not just involve perceiving relations like temporal order and duration between external events. Now, I prefer a metaphysical view of time in which it doesn't literally involve a "moving present"—*a fortiori* we don't perceive such a thing. Nonetheless, it seems to me that we do genuinely have a sense of a constant change in what is present, a sense which arguably helps ground the intuitions about time that believers in a moving present have[6] (I will not say any more about the fascinating topic of the metaphysics of the moving now here). Supposing this is right, what would account for it?

Let us begin by noting that our perceptual system does, in some sense, present events as happening in the present. We automatically assume that what we see is happening in the present—for example, we make present tense judgments based on experience, and act on the assumption that perceived events are current. (This can be true even if we are also capable of accommodating cases where we know that there has been a significant delay between information transmission and reception at our sensory periphery.)

Slightly more contentiously, we have a conscious sense of events having happened in the immediate past. For example, a briefly perceived event like a bright flash of light in some way lingers in consciousness longer than the brief moment when it is perceived as present. Or the presently experienced passage in a piece of music is somehow experienced as having a particular musical context (the phenomenology I am talking about is also at its most obvious in cases where we are anticipating that an event will happen, and deliberately attend to it when it does in fact happen). These phenomenological claims are partially substantiated by evidence for the existence of short-term memory mechanisms linking perception with other processes downstream. 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 seconds, but their performance rapidly drops off beyond this range. Perhaps these working memory representations form a "long specious present," which contrasts with a shorter temporal window of events presented in vivid phenomenal awareness, and which gives us a sense of the immediate past of events that are experienced as present.

If this is right, then we can think of perception as a conveyor belt of information, being first processed unconsciously, then appearing in phenomenal

consciousness as a presentation of a short temporal array of current events, and then briefly lingering in short-term memory, still contributing to what it is like for the subject, as a conscious sense of the immediate past. In order to keep up with present events, notice that there must be a constant *updating* of the information that is represented at each stage—the conveyor belt has to keep moving, otherwise we won't keep up with the world. The existence of these constantly updated conscious representations of what is present and immediately past is perfectly consistent with the basic version of atomism, and probably with extensionalism too (although maybe it is easier for an atomist to accommodate a sense of what is immediately past *as* being immediately past).

Now take note of the following important point: the mere *existence* of these constant changes is not enough to give us a sense *of* the flowing present—for that, it seems that we need to be in some sense aware *of* the process of updating occurring, not merely be such that it is *in fact* occurring. Consider an experience of a brief event happening and then seeming to quickly fade into the past. To get such a sense of fading, it will not be enough that the representation of the event is *in fact* moved along the conveyor belt from perception-as-present into conscious short-term memory, and then out of consciousness entirely. This process will have to be one that we have a higher-order awareness of. Simply having a series of first-order experiences of external world events (whether atomic or extensional), even if they are accompanied by short-term memory experiences, will not give us this awareness. That would only amount to a "frozen" sense of the layout of events from one temporal perspective, not a sense of a constant flowing *change* of temporal perspective. Arguably, to explain the sense of experience flowing from one moment to the next, we need to postulate a higher-order awareness that can be directed to the changing array of first-order conscious experiences.

Notice that, since first-order experiences, even if tensed, could not on their own give us a sense of changing temporal perspective, it must be that this higher-order awareness is inconsistent with "transparency" as articulated by Phillips. If such higher-order awareness exists, then we are not stuck inside our present perspective, aware only of how it presents the world, but have a sense of this perspective *itself* evolving through time. (Note that there is a difference between this awareness of temporal flow, and a mere awareness *that* our temporal perspective has changed, of a kind that we might get from episodic memory.)

I am not the first person to suggest that such a thing exists—Husserl held a similar position; furthermore, he may have intended it to explain the sense of temporal flow (see Miller (1985) for an interpretation of Husserl along these lines). I also do not say that the idea is without problems, or that it is completely obvious that there is any need for such a thing. We might

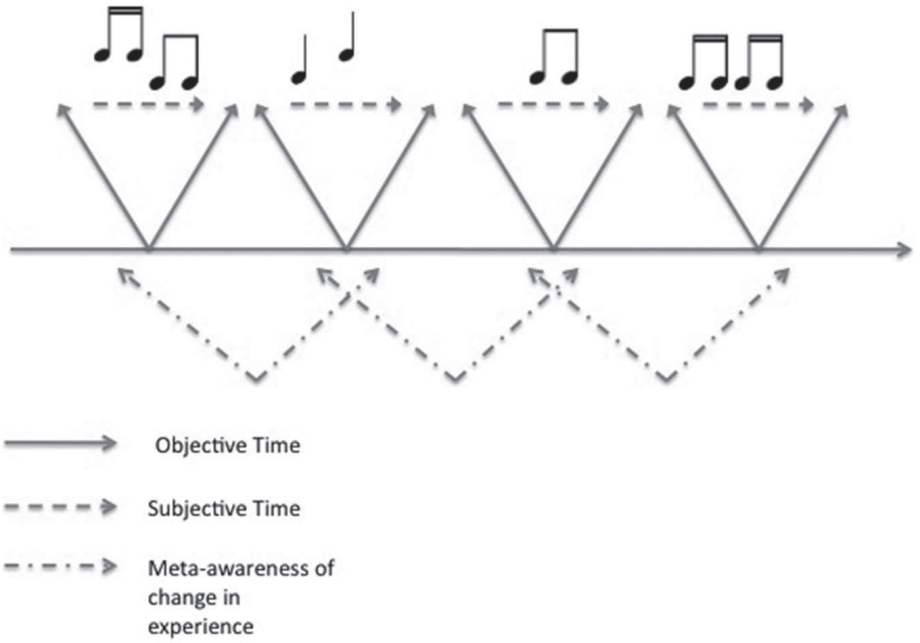

FIGURE 8.5

wonder if these higher-order states are quasi-perceptual experiences that can misrepresent first-order states, or stand in some other kind of relation to their targets. We might also wonder if there is any independent reason for thinking that they exist. These are important questions, but I will not address them here. Perhaps pursuing them will reveal that I have misdescribed experiential flow, or that the idea that it exists is an illusion. What I am saying here can perhaps best be put as a dilemma. If there is a flow in experience that needs explaining, it *either* has to do with a sense of external events continuing on from before, *or* a sense of a constant change in temporal perspective of the kind discussed here. The former can easily be accounted for by the atomist in terms of the contents of first-order experiences, as explained above. One might doubt that there really is an experience of the latter kind, but in so far as there is, a story of the kind that I have just told, involving a higher-order awareness of changes in your temporal perspective, might be true. Furthermore, this is a story that is perfectly compatible with atomism, and extensionalists do not have an alternative story to offer that requires their view to be true. This shifts the burden to the extensionalist to explain how there is a kind of experienced continuity or flow between experiences that is incompatible with atomism.

Conclusion

The main aim of this discussion has been to defend atomism against the criticisms of it that appear in the current literature, especially those in Phillips (2010) and Dainton (2000, 2010). I looked at Phillips' (2010) argument, suggesting that one or more of the main premises can plausibly be rejected by the atomists. I then considered Dainton's argument from Jamesian continuity, suggesting that, if anything, atomists are better placed to explain our sense of flow and continuity in experience than extensionalists are.

Notes

1 These temporal parts need not be completely independent events capable of existing on their own. What matters is that they are distinct events capable of happening at different times; this may (or may not) be compatible with various forms of mutual interdependence (for example, they might have independent core realizations, but identical or overlapping total realizations). For more discussion of various forms of interdependence between experiences, see Lee (2014b).

2 Perceptual systems are able to recalibrate information in time to account for discrepancies in transmission time from different sources. For example, simultaneous taps on the nose and feet feel simultaneous even though it takes longer for the signal from the feet to be transmitted to the brain. Consider a case where your nose is tapped slightly in advance of your feet. You might correctly experience this temporal order, even though the signal from the feet does not arrive before the signal from the nose. Given atomism, this makes it at least doubtful whether you experience the foot tapping before you experience the nose tapping.

3 Interestingly, it is not clear whether extensionalism is compatible with such a discrete gappy structure: the extensionalist's commitment to mirroring might suggest that if experience had this structure, events in the external world would seem to also have a discrete gappy structure, which they do not. I suspect that there are versions of extensionalism that avoid this problem by adopting a sufficiently weak version of mirroring, so I won't pursue the objection here.

4 One way to see this is, is to note that Overlap Extensionalism is (apparently) completely consistent with strong transparency. So even if Overlap Extensionalism is true, we may only ever introspectively apprehend the worldly appearances provided by total experiences, not diachronic connections between experiences themselves. In other words, it is not clear why the fact that, on this view, total experiences overlap with each other, or the fact they are spread out in time, is something that would be revealed by introspection, or contribute to a sense of continuity or flow that is absent on the Atomic picture.

5 It is true that the Atomist will deny that having an experience with a certain temporal content is the same thing as having awareness of how experience is changing over time. But as we saw above in discussing Phillips' argument, it is contentious to assume that ordinary temporal phenomenology, involving experiences of events in the external world playing out in time, *also* involves direct awareness of experience itself changing. If the idea of Jamesian flow depends on such an assumption, then that atomist can reasonably say that it is suspect.

6 We can distinguish two aspects to these intuitions : (1) an intuition that the present moment is metaphysically special; and (2) an intuition that there is a constant change in which events are highlighted as present. The psychological phenomenon I am interested in is more relevant to (2). For attempts to explain (1), see Butterfield (1984) and Callendar (2008). A recent attempt to explain (2) (or something close to it), in terms different from those I discuss here, is Paul (2012).

References

Bayne, T. (2010). *The Unity of Consciousness*. Oxford: Oxford University Press.

Broad, C. D. (1925), *The Mind and its Place in Nature*. London: Routledge & Kegan Paul.

Butterfield, J. (1984), "Seeing the present," *Mind,* 93, 161–76.

Callendar, C. (2008), "The common now," *Philosophical Issues*, 18, 339–61.

Chuard, P. (2011), "Temporal experiences and their parts," *Philosopher's Imprint*, 11.

Dainton, B. (2000), *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

—(2004), "Precis of 'stream of consciousness'," *Psyche*, 10, 1.

—(2010), "Temporal Consciousness," in Edward N. Zalta (ed.), *The Stanford Online Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/fall2010/entries/consciousness-temporal/

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, S209–18.

James, W. (1912), *Essays in Radical Empiricism*. New York: Longmans.

Kline, K., Holcombe, A. O., and Eagleman, D. M. (2004), "Illusory motion reversal is caused by rivalry, not by perceptual snapshots of the visual field," *Vision Research*, 44, 23, 2653–8.

Lamme, V. A., Super, H. and Spekreijse, H. (1998), "Feedforward, horizontal, and feedback processing in the visual cortex," *Current Opinion in Neurobiology*, 8, 4, 529–35.

Lee, G. (2014a), "Temporal experience and the temporal structure of experience," *Philosopher's Imprint*, 14, 3, 1–21.

—(2014b), "Experiences and Their Parts," in D. Bennet and C. Hill (eds), *Sensory Integration and the Unity of Consciousness*. Cambridge, MA: MIT Press.

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., and Ro, T. (2009), "To

see or not to see: prestimulus α phase predicts visual awareness," *Journal of Neuroscience*, 29, 9, 2725–32.
Mathewson, K. E., Lleras, A., Beck, D. M., Fabiani, M., Ro, T., and Gratton, G. (2011), "Pulsed out of awareness: EEG alpha oscillations represent a pulsed-inhibition of ongoing cortical processing," *Frontiers in Psychology*, 2.
Miller, I. (1985), *Husserl, Perception and Temporal Awareness*. Cambridge, MA: MIT Press.
Paul, L. (2012), "Temporal experience," *Journal of Philosophy*, CVII, 7, 333–59.
Phillips, I. (2010), "Perceiving temporal properties," *European Journal of Philosophy,* 18, 2, 176–202.
—(2013), "Perceiving the passage of time," *Proceedings of the Aristotelian Society*, CXIII.
Pöppel, E. (2004), "Lost in time: a historical frame, elementary processing units and the 3-second window," *Acta Neurobiologiae Experimentalis (Wars)*, 64, 295–301.
Tye, M. (2002), "Representationalism and the transparency of experience," *Noûs* 36, 1, 137–51.
—(2003), *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.
Van Rullen, R. and Koch, C. (2003), "Is perception discrete or continuous?," *Trends in Cognitive Sciences*, 7, 5, 207–13.