The Fragmentary Model of Temporal Experience and the Mirroring Constraint

(FORTHCOMING IN PHIL. STUDIES – PLEASE CITE THE OFFICIAL VERSION)

A central debate in the philosophical literature on temporal experience focuses on the following question: do temporal experiences themselves have a temporal structure that mirrors their temporal contents? *Extensionalists* argue that experiences do indeed have a temporal structure that mirrors their temporal contents. *Atomists*, on the other hand, deny this claim, and insist that experiences do not have a temporal structure that mirrors their contents.[1]

In this paper I argue that something has gone wrong with this debate. Atomism and extensionalism are proposed as general theories that describe the relationship between the temporal contents of experience and the temporal structure of experience. However, I will argue that both theories fail in this regard. Temporal experience is not a single undifferentiated phenomenon, but instead, breaks down into a number of distinct capacities to represent the temporal structure of our world. When we look closely at the different mechanisms that underpin these distinct aspects of our temporal experience we find that some of them preserve a mirroring between their temporal content and their temporal structure while others do not. The result is that atomism and extensionalism are false as general theories of temporal experience.

In section 1, I lay out the standard puzzle of temporal experience and situate the debate between extensionalists and atomists within that context. In section 2, I argue that introspection alone cannot settle this debate. Turning away from introspective arguments, in section 3, I consider two types of empirical arguments given in favor of atomism, *the argument from temporal illusions* and *the trace integration argument*, and argue that both fail to establish the truth of atomism. Having shown that the arguments for atomism and extensionalism fail, in section 4, I argue for *the fragmentary model of temporal experience*. Temporal experience is not a single unitary phenomenon, but is instead fragmented into distinct capacities to perceive the temporal structure of our world. Then in section 5, I turn to two specific aspects of our temporal experience and argue on the basis of these cases that we have counterexamples to both atomism and extensionalism. As a result, neither theory adequately characterizes how the temporal contents of experience relate to the temporal structure of experience.

1. The puzzle and the debate

Many philosophical discussions of temporal experience begin with the statement of a puzzle that calls into question *the very possibility* of experiences with temporal contents.[2] While different

[1] The literature on temporal experience lacks an accepted taxonomy of positions (both in how theories are divided and their labels). My use of 'extensionalism' and 'atomism' comes from Lee (2014a) and maps onto Hoerl's distinction between *molecularism* and *atomism* (Hoerl, 2009), and Lee's earlier distinction between *experiential process views* and *atomism* (Lee, 2014b).

[2] Other salient discussions that begin by posing a similar puzzle include (Boroditsky, 2011; Dainton, 2014a). Despite minor differences, the main puzzlement remains the same – *how is it that experience can represent time if experience is due to the causal relation between the senses and the current state of the world*.

authors put the puzzle in slightly different terms, we can follow Sean Kelly's formulation of the puzzle. Kelly says,

> How is it possible for us to have experiences as of continuous, dynamic, temporally structured, unified events given that we start with (what at least seems to be) a sequence of independent and static snapshots of the world at a time?
>
> (Kelly, 2005, p. 2)

According to Kelly, there seems to be a tension between two "facts" about experience. On the one hand, experience seems to provide us with access to the temporal structure of the world around us. We experience events *as having durations, as being earlier than* or *later than* other events, etc. Experience presents us with how the world is over an extended temporal interval.

Even though experience seems to present us with how the world is over an extended interval of time, Kelly claims that experience also seems to be ultimately composed of a sequence of *static snapshots* of the world as it is at specific moments in time. This fact, that experience is composed of static snapshots, is supposedly something that we know on the basis of introspection alone – our phenomenology tells us this is the case.[3] If experiences are ultimately composed of static snapshots that do not present any duration, then, Kelly argues, it is difficult (if not impossible) to explain how experiences can come to present us with temporally extended events.[4]

As most philosophers are reluctant to deny that we have temporal experiences, most reject the snapshot model.[5] However, while rejecting the snapshot model resolves the tension, we are still left wondering how it is possible that experiences have temporal contents. The debate over *the mirroring constraint* arises in trying to answer this question.

> *The Mirroring Constraint*: if experience represents some temporally structured event, then the experience itself will have temporal parts that correspond to the temporal stages of the

[3] While Kelly himself doesn't give an argument for the snapshot model, an argument for the view can be found in the early modern empiricists and their distinction between simple and complex experiences (Hume, 2000; Locke, 1689; Reid, 1855). Consider an experience that presents a non-zero duration. With the experience of any duration we can, at least conceivably, imagine what it would be like to experience just the beginning (or later) half of the duration. Therefore, the experience of any non-zero duration is a complex experience. The simple experiences that constitute all experience, must therefore be snapshot-like in that they cannot present any duration.

[4] Someone might try to salvage the snapshot model by arguing that *some process stitches these snapshots together into complex temporal experiences*. Therefore, the temporal contents of the complex temporal experience would be more than a sum of the temporal contents of the snapshots. This only shifts the same question to how that stitching process acquires its content.

[5] Notably Chuard (2011), Gallistel (1996) and Reid (1855), have chosen to bite the bullet and deny that we have temporal experience by relying on restricted understandings of *experience*. For instance, according to Gallistel, experiences are only brought about by domain specific sensory transducers, of which there are none for time, therefore there cannot be temporal experience. For Reid, experiences cannot involve cognition or memory, both of which are involved in temporal representation, so therefore there cannot be temporal experience. Despite their claims, however, we can still speak of *what it is like to represent the temporal structure of the world around us in a way that is different than pure recollection or thinking*. If we take experience in this broad sense of the term, then we are not beholden to any architectural restrictions on what counts as genuinely experiential, and it is in this broad sense of experience that our discussion will continue.

represented event, and the temporal properties that experience represents will be mirrored by the temporal properties of the experience itself.

Extensionalists[6] argue that experiences obviously satisfy the mirroring constraint, with Ian Phillips (2010) going so far as calling extensionalism the *naïve theory of temporal perception*. Part of the appeal of extensionalism comes from *the mirroring intuition*. Consider how someone might respond if asked to explain how a visual experience of a healthy field of grass comes to represent the grass as *being green*. Without appealing to vision science, it's hard to say anything more informative than the grass *simply looks green*. Now, consider how someone might respond if the question were to explain how an experience of a crash of lightning being later than a flash of lightning comes to represent this temporal relation. A perfectly respectable sounding reply would be that first there was the experience of the flash of lightning and then there was the experience of the crash of thunder. In the temporal case, it seems intuitive to explain how experience has its content by appealing to the temporal structure of experience itself and not just some brute phenomenal appearance.

Extensionalists take the mirroring intuition to reveal something very deep about experience. When an experience presents us with the crash of thunder (T) being later than the flash of lightning (L), extensionalists claim that this overall experience can be divided into at least two components, *the experience of T* (E_T) and *the experience of L* (E_L), and E_t will be later than E_L (see Figure #1).[7]

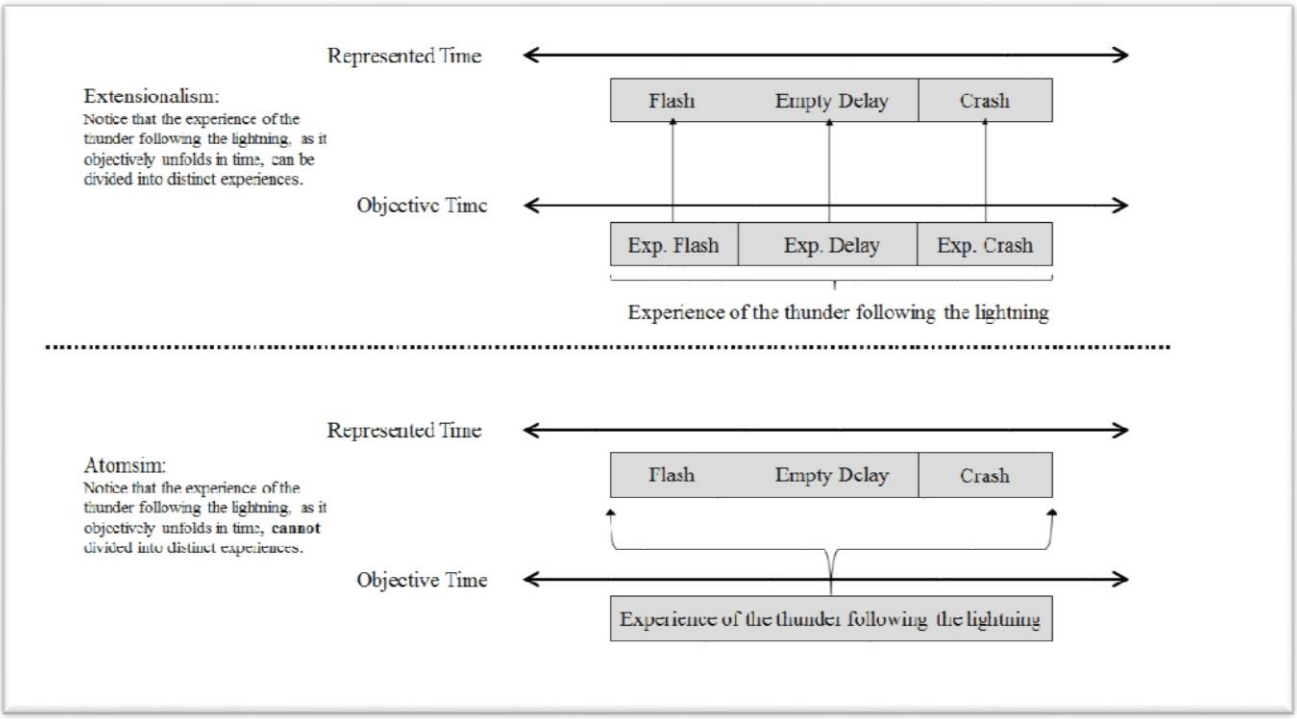

Figure 1: Extensionalism (top). Atomism (bottom)

[6] Extensionalists include (Dainton, 2008, 2011; Hoerl, 2009; Martin, 2002; Phillips, 2010, 2014b; Rashbrook, 2013).
[7] Extensionalists needn't be committed to experiences being divided into arbitrarily short temporal parts. Rather, they are committed to there *being some* such temporal parts.

Atomists[8], on the other hand, deny that we should make much of the mirroring intuition. Our experiences rarely, or even never, satisfy the mirroring constraint since our experiences represent the entire temporal interval *over the same span of time*, as a result, experience cannot be broken down into temporal parts that represent the distinct temporal stages of the represented interval (Figure #1).

As Lee (2014b) has emphasized, atomists are not claiming that experiences are instantaneous (as Dainton (2014) seems to characterize the position). If atomists were saying that experiences were durationless, then their view would be false, but for reasons having nothing specifically to do with *temporal* experiences. Plausibly, the temporal properties (or structure) of experience (e.g. how long experiences last and their temporal relations to other events / experiences) are inherited from the temporal properties of whatever physical processes *realize* those experiences. Since our best evidence suggests that experiences are realized in the brain by temporally extended processes (e.g. patterns of activity over brief time intervals, neural oscillations, etc.) it follows that experiences are not durationless. Therefore, atomism would be false but for reasons having nothing to do with temporal experience.

Instead, the debate must be understood in terms of mirroring. If temporal experiences have temporal parts that individually represent stages of the experienced temporal interval, and these temporal parts of the experience have temporal properties that mirror those temporal properties represented in experience, then the mirroring constraint will be satisfied and extensionalism will be true. Otherwise, if there is no mirroring, then extensionalism is false and atomism is true.

2. The introspective argument for extensionalism

As mentioned already, extensionalists take introspection to vindicate their position. While the specific arguments vary between authors, they are generally of the following form: The introspective evidence from temporal experience would not be available if the mirroring constraint were violated. Therefore, extensionalism must be true.

Both atomists and extensionalists agree that through introspection we form judgments about the temporal structure of experience. If you were to shift your fixation back and forth between the top and bottom lines of text on this page, it would be natural to judge that first you have an experience of the top line, then an experience of the bottom line, and back and forth. The temporal order of your experiences seems to be something that introspection allows us to form true beliefs about. However, to evaluate the introspective arguments for extensionalism, it's important to understand *how* these judgments are formed.

We can consider this introspective evidence as possibly coming in two forms. On the one hand, introspection may provide us with direct access to the temporal structure of experience. Alternatively, introspection may provide us with information about the world from which we *infer* the temporal structure of experience. That is, we may infer the temporal structure of experience from the contents of experience (which often include representations of how events are temporally structured).

[8] Representative atomists include (Grush, 2005, 2006; Lee, 2014b; Watzl, 2012).

If introspection were to give us *direct access* to the temporal structure of experience, then the argument for extensionalism would seem straightforward. Simply appeal to the widespread existence of the mirroring intuition and conclude that the temporal structure that experience appears to have is in fact the temporal structure that it does have. However, it's unclear whether experience gives us this sort of direct access. Consider again the case of shifting your fixation between the top and bottom lines of text. In this case, we feel comfortable making judgements about the temporal structure of our experience, and there's a good chance that our judgments are correct. Yet, the temporal structure of these experiences could be inferred from the changes in the content of our visual experience along with information about the top-down guidance of fixations. We may be inferring the temporal structure of our experience from how the world is presented to us and our actions in the world.[9]

If we did have direct access to the temporal structure of experience that was independent of our access to the temporal structure of the events in world, then we would expect there to be cases in which the apparent temporal structure of experience would come apart from the apparent temporal structure of the events in the world. There would be cases in which experience would seem to have a particular temporal structure, and the events represented by experience would seem to have a distinct temporal structure. However, to my knowledge, no cases of this sort exist, and in fact, the mirroring intuition speaks strongly against their existence since cases like this are just what the mirroring intuition denies. It also seems unlike that our access to the temporal structure of events in the world is inferred from the apparent temporal structure of experience, given the empirical evidence concerning the mechanisms of temporal perception to be described in section 5. Instead, as Phillips (2010) has argued in defending extensionalism, our introspective access to the temporal structure of experience is inferential.

If our judgments about the temporal structure of experience are inferential, then the extensionalist must provide some reason for why the contents of experience provide us with good evidence for how experiences are temporally structured. If no argument of this sort can be given, then we must simply conclude that introspective evidence cannot settle the mirroring constraint debate.

Phillips argued for the infallibility of this inference based on the assumption that experience does not allow for a *seems* / *is* distinction. According to Phillips, "experience is a matter of things *seeming* a certain way, and if, things seem to seem a certain way on rational reflection, then [experience is that way.]" (2010, footnote #15). If we combine this *seems→is principle* with the claim that experience seems to have a particular temporal structure, then it follows that experience has that temporal structure.

While Lee (2014a) surveys a number of different objections to the seems→is principle, we'll take a somewhat different approach. If, as the seems→is principle requires, the apparent temporal structure of experience *must be* the temporal structure of experience, and the apparent temporal structure is arrived at through an inference from the content of experience, then it must be the case that only experiences with a temporal structure satisfying the mirroring constraint could possess the contents from which we infer the apparent temporal structure of experience. That is, the contents of experience must place representational constraints on how our experiences themselves are temporally

[9] This is similar to the inferential approach to introspection proposed by some transparency theorists (e.g. (Dretske 1994)). Note that the argument here does not depend on experience in general being transparent, only that temporal experiences are transparent.

structured. In the remainder of this section, we'll look at possible theories of content for experience to show that extensionalists cannot point to any meta-semantic facts that would explain these constraints without begging the question. As a result, the introspective argument fails.

Since there are too many candidate theories of content to consider each one individually, we can make a rough division between two types of theory – *vehicle neutral theories* and *vehicle dependent theories*. Both types of theories attempt to explains for the content of mental representations are mapped onto and acquired by their underlying representational vehicles, however, the types of explanations these theories offer differ significantly.

According to vehicle neutral theories of content, a representation's content is determined by the representational vehicle's relations to the world beyond itself. According to vehicle dependent theories, on the other hand, a representation's content is determined by properties that lie within the boundaries of the representational vehicle itself. A prime example of a vehicle dependent theory is the *resemblance theory of content* according to which a representation has its content in virtue of the representational vehicle resembling its content. For example, a paint swatch represents the color of paint in a can by itself being the same color as, i.e. resembling, the paint. The vehicle internal properties of the representation determine what the representation applies to.

If the seems → is principle is going to do the work that Phillips requires of it, then at least one of these types of theory must explain how the contents of experience appropriately constrain the temporal structure of experience. Let's take both types of theory in turn, starting with vehicle neutral theories, to see how the extensionalist cannot argue for this point without begging the question.

Even amongst vehicle neutral theories we find many alternative approaches. Causal theories (Fodor, 1994; Prinz, 2002; Stampe, 1977) explain how representations acquire their content by appealing to the (privileged) causes of the production of those representational states. Teleological theories (Dretske, 1995; Millikan, 1989, 1995; Ryder, 2004) explain how representations acquire their content by appealing to their selection history (either through evolution or learning). Finally, information-theoretic accounts (Dretske, 1981; Skyrms, 2010; Usher, 2001) explain how representations acquire their content by appealing to the way that the state of the representational system reduces the uncertainty regarding the possible states of the world.

To further simplify our discussion, let's focus on Dretske's information-theoretic account as what we say about his account will generalize to other vehicle neutral theories. According to Dretske, the content of a representation is determined by the following biconditional:

> A [system] *r* carries the information that *s is F* if and only if the conditional probability of *s's being F* given *r* (and [the background conditions] *k*), is 1 (but given *k* alone, less than 1).
>
> (Dretske, 1981, p. 51)

At the root of the information-theoretic account of mental representation is that given certain background conditions, i.e. physical states of the representational system and how it relates to the world around it, the fact that a representational system is in a particular state (e.g. instantiating some property) reduces the uncertainty regarding some other state out in the world. If the reduction in uncertainty is sufficiently complete, then knowing the state of a representational system puts you in a position to know about some state in the world. For instance, given how your car's gas gauge is

connected to the fuel tank (i.e. given the background conditions), the state of the gas gauge tells you about the state of the fuel tank.

Notice that the information-theoretic account says nothing about how representational systems must be structured. It only requires that the appropriate relation holds between the representational system and the world. With this in mind we can show that an information-theoretic account of content is capable of explaining how temporal information can be carried by systems that satisfy mirroring and systems that violate mirroring. To illustrate this point, let's consider a toy example.

Imagine that Susan must keep track of the temporal relation between two events, *a* and *b*. Since Susan cannot directly observe these events she relies on a pair of monitoring systems – a *mirroring device* and a *non-mirroring device*.

> *Mirroring Device*: this device produces two tones, t_a and t_b, that respectively stand for *a* and *b*. The temporal relation between *a* and *b* is mirrored by the temporal relation between t_a and t_b. Susan simply notes the temporal relation between t_a and t_b to know the temporal relation between *a* and *b*.
>
> *Non-Mirroring Device*: this device produces two lights. The first light can be blue or green. The second light can be any shade between yellow and red. The first light indicates whether *a* or *b* comes first. The second light indicates how long of a delay there is between *a* and *b*. By noting the color of the two lights Susan can know the temporal relation between *a* and *b*.

Simply by observing Susan's behaviors we have no idea which device she is using, since both devices do an equally good job of conveying the relevant information about *a* and *b*. Nothing about the temporal content the systems encode demands that the representational systems have a particular structure.

The analogy should be clear. If we adopt a vehicle neutral theory of content, then there is nothing particularly special about temporal information that constrains the temporal structure of the representational system. If a vehicle neutral theory were to explain how our experiences acquire their temporal contents, and our judgments about the temporal structure of experience are based on inferences from the contents of experience, then we cannot be certain that the mirroring constraint is satisfied. That is, the seems→is principle could not be justified, since the contents of experience would not provide us with evidence in favor of either theory.

Instead, to defend the seems→is principle and the introspective argument, extensionalists must appeal to a vehicle dependent theory, since only vehicle dependent theories of content provide the necessary connection between the content of experience and the structure of experiential vehicles. However, as I will argue, extensionalists cannot appeal to a vehicle dependent theory of content without begging the question.

Mental representations aren't generally thought to have their content in virtue of resemblance. For instance, while we may have visual experiences of gray, the fact that the brain is gray does not explain this fact.[10] The color of the brain simply isn't a good candidate for being the property of experiential vehicles that determines their content. One reason is that whatever properties of experiential vehicles determine their content must co-vary with changes in the content of experience,

[10] This example comes from Grush (2006).

and the color of the brain does not change in this way (i.e. your brain remains gray regardless of what you experience). The second reason is that the properties of representational vehicles that bear the content of the representation must be capable of causally influencing downstream processes, otherwise, the content would be inaccessible to any further processes, including verbal report, and the color of the brain doesn't have these causal powers.[11]

An extensionalist, however, could argue that in the case of temporal representation, these two concerns are avoided. If we accept the mirroring constraint, then the semantically relevant properties of the representational vehicle, i.e. its temporal properties, will co-vary with the content of the representation. This simply is what the mirroring constraint requires. The temporal properties of the experience itself are mirrored by the content of the experience. Second, differences in the temporal properties of experiences can have different causal influences on downstream processes. For instance, a neural system sustaining a particular level of activity for 200ms can have very different downstream effects than the same system maintaining that level of activity for 100ms. Accordingly, the temporal properties of experiential vehicles bring with them causal influences that the color of the brain does not.[12]

However, it's here where extensionalists are backed into a dialectical corner. If the argument for extensionalism requires the seems→ is principle, and justify this principle requires a resemblance theory of content, then the justification for adopting the resemblance theory must avoid appeals to facts about the temporal structure of experience. Yet, a resemblance theory for temporal representation is only viable if we assume mirroring. Therefore, the extensionalist cannot defend the seems→ is principle without begging the question in favor of extensionalism. Since we cannot accept this question begging, it follows that the introspective arguments for extensionalism fail. Other, non-introspective, evidence is needed to resolve this debate.

3. The case for atomism

In contrast to the introspective arguments given by extensionalists, atomists typically adopt empirical arguments. The most common strategy adopted is to argue from the existence of certain temporal illusions to the truth of atomism. Again, while specifics differ, the general strategy remains the same. In certain situations we misperceive the temporal structure of the world around us. If the misperceived temporal properties were mirrored by the temporal structure of experience, then we would be committed to implausible (or even contradictory) claims about experience. Therefore, mirroring must be false. [13]

For each type of temporal illusion, however, extensionalists seem to have a viable response. This literature on temporal illusions has received so much attention, that for our purposes, I will simply refer the reader to the relevant work. Phillips (2014a, 2014b) surveys a number of temporal illusions and shows how given plausible, and importantly not ad hoc, maneuvers extensionalists can accommodate these illusions. In some cases, extensionalists rely on alternative interpretations of the illusory

[11] Set aside the other problems that resemblance theories face.
[12] In this case, it's not the *mere* possession of temporal properties that is doing the causal work. Rather, it's some time-dependent property of the system that is doing the causal work.
[13] Examples of this strategy include (Grush, 2005; Lee, 2014b; Watzl, 2012).

experiences.[14] In other cases, extensionalists appeal to more general findings in the cognitive sciences that allow them to accommodate the illusory experiences.[15] In general, though, temporal illusions alone do not place sufficient constraints on the structure of experience to settle the debate between atomists and extensionalists. Some other source of evidence is needed to establish atomism.

It's at this point that Geoffrey Lee's *trace integration argument* comes into play. Lee argues that we should accept atomism given general considerations about information processing in the brain.

Lee's argument goes as follows (quoted from Lee (2014b)):

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 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).
5. Experiences that have the same realizer timing have the same timing [...]
6. Therefore, the different experiential parts of a temporal experience have the same timing.
7. Therefore, temporal experiences are atomic […].

Central to Lee's argument is that in typical cases the contents of experience are made available to downstream processes, and the only way for temporal contents to be made available is for them to undergo a process of *trace integration*. In order to understand what Lee means by trace integration, let's go through an example.

Consider the experience of lightning preceding thunder, and suppose that the lightning and thunder are occurring sufficiently far away that there is a noticeable gap between the two. At the point that the flash of lightning is being processed by the sensory system, it is too early for anyone to have the experience of the flash of lightning *as preceding the crash of thunder*, since the thunder has yet to impact the sensory system. However, once the sound of the thunder reaches the perceptual system, the lightning, and the experience of the lightning, will no longer be present. The experience of the lightning at this point will be in the past. Nevertheless, it is at this point, once the thunder has been perceived, that the person is in the position to have the experience *of the thunder following the lightning*.

While we could interpret the temporally extended pattern of activity in early sensory areas as implicitly encoding information about the temporal relation between the thunder and lightning, Lee argues that this temporally extended patterns of activity *alone* (i.e. the successions of experiences), cannot make the temporal content *the thunder followed the lightning* accessible to downstream processes. At the time that these post-perceptual systems are attempting to access this temporal

[14] For instance, (Phillips, 2014b) argues that apparent cases of temporal dilations, in which we perceive events as having much longer durations than they in fact do, are actually cases in which *other* stimuli are contracted in duration.

[15] For example (Dainton, 2008, 2011) appeals to the distinction between perception for experience and perception for action, from the literature on the two visual streams, to argue that illusory experiences and the mirroring constraint needn't lead to sensori-motor problems.

content, the experience of the lightning will be in the past. But, merely past events cannot influence current or future neural processing without leaving behind a trace in the current state of the perceptual system. If the experience of the lighting really did leave no trace behind, then at the time the thunder is processed, it would be as though the experience of the lightning had never occurred. Assuming that in normal cases the content of experience is available to downstream processes our experience of temporal succession must be more than just the *mere* succession of perceptual states.

Instead, Lee argues, the temporally extended sensory processes that register temporally extended events in the world, must leave behind *traces* that are integrated into the current state of the perceptual system. It is because the earlier experience of the lightning leaves behind a trace, that the current experience of the thunder can be experienced *as following the lightning*. But, if the current state of the system represents the thunder and the lightning at the same time in the form of a trace-integrated representation, then the representation itself will not have a temporal structure that mirrors its content. Therefore, atomism must be true.

Saying that the temporally extended events are represented *at the same time* through a process of trace-integration does not mean that they are represented at *an instant*. Instead, following Mauk and Buonomano (2004), Lee describes how the brain must convert temporally extended patterns of activation into *spatial codes* whereby temporal information is encoded either by the distribution of activity within neural populations or by the firing rates of neurons. While the details of the spatial coding are to be determined by empirical models, some process of trace-integration is needed.

We can accept that Lee is correct about the need for trace integration in order to make temporal contents accessible. Nevertheless, the conclusion that mirroring is violated does not follow.

When Lee evaluates whether the mirroring constraint is satisfied he only considers the present state of the perceptual process as it is this state that makes temporal information available to downstream processes. Since this state presently represents a temporally extended interval, the mirroring constraint is supposedly violated. However, it's not clear that we should only be analyzing this present state of neural processes. Let's consider a case of duration perception.

Imagine you hear a tone that lasts 80ms. Due to the concerns about making this temporal content accessible to downstream systems, it must be the case that the entire temporal content of the experience that is accessible to downstream processes is contained in the current, causally efficacious, stage of the representational process. Yet, this process that represents the tone as having a duration of 80ms did not arise out of nowhere. This representational state, like all neural states, is brought about by a process that itself unfolds in time.

When we consider the temporal interval *represented* by this experience, we find that the interval can be divided into parts (or sub-experiences) *contained* by the overall experience. These *contained experiences* include *the tone lasting 70ms*, *60ms*, *50ms*, *etc.* Despite the fact that the causal efficacy of these contained experiences will be found entirely in the current state of the system that represents the tone as lasting 80ms, it is nevertheless possible that the process that gives rise to the *containing experience* of the tone lasting 80ms will itself have a temporal structure that satisfies the mirroring constraint. That is, if the process that gives rise to the experience of the tone lasting 80ms can itself be broken down into temporal parts that correspond to, and mirror the temporal relations

between, the contained experiences, then despite there being a process of trace-integration, the mirroring constraint would nevertheless be satisfied.[16]

Nothing in Lee's argument prohibits this from being the case. Lee's argument tells us something about what must be the case for temporal contents to be made accessible to downstream processes, but by itself this says nothing about how experiences must be temporally structured. In the next section, we will see that the empirical work on duration perception shows us that the processes that give rise to the experience of short durations do in fact satisfy the mirroring constraint.

4. The fragmentary model of temporal perception

We've seen that the arguments in favor of extensionalism and atomism fail to establish their intended conclusions. What we need is another source of evidence that can constrain this debate. It's here where emerging neuroscientific research becomes relevant. In this section we'll begin to address the debate over the mirroring constraint by arguing for *the fragmentary model of temporal perception*. 'Temporal experience' and 'temporal perception' do not pick out a single undifferentiated psychological phenomenon.[17] Instead, empirical evidence suggests that temporal experience is fragmented and composed of many dissociable capacities to perceive the temporal structure of our world. By looking at these individual perceptual capacities, and the mechanisms that underpin them, we will see that aspects of our temporal experiences satisfy the mirroring constraint while others do not. Ultimately, no general story about how the temporal contents of experience relate to the temporal structure of experience itself can be told.

While we perceive durations, temporal orderings, rhythms, etc, and while we perceive these properties across a range of timescales and through multiple sensory modalities, many scientists and philosophers nevertheless assume that a single mechanism or type of explanation will allow us to account for the human perception of time. Paraphrasing a recent paper by Hartcher-O'Brien and colleagues (2016), the search for a single perceptual timekeeping mechanism is motivated by the fact that in the world and experience the various temporal properties we perceived are intimately tied up with one another. The various temporal properties we perceive are all integrated into a single temporal ordering in experience, and therefore, we should expect there to be some unity to the timekeeping mechanism in perception.

In keeping with the search for *the* neuroscientific explanation of temporal experience, researchers often describe their favored explanations as being in direct conflict with those proposed by others (even when these models are developed in response to very different experimental methodologies).[18] Over the years, the most influential single mechanism account of temporal perception has been the *scalar expectancy theory* (SET) (Gibbon, 1977; Gibbon et al, 1984). Building off of

[16] Another way of putting the point is that if the process that gives rise to the experience of the tone lasting 80ms can be divided into temporal parts, and these temporal parts are themselves semantically interpretable such that they individually represent the contained experiences of shorter durations and stand in the same temporal relations to one another as the durations represented by those parts, then the mirroring constraint is satisfied.
[17] This point is similar to one made in (Craver et al., 2014) that 'temporal experience' is prone to *semantic drift*.
[18] For instance, *state-dependent models* (Ivry & Schlerf, 2008), are developed entirely in response to sub-second stimuli, while *default mode network models* (Lloyd, 2012) are developed with supra-second studies.

Treisman's (1963) *pacemaker-accumulator model*, SET consists of a three-part model. First, there is the pacemaker-accumulator system. A pacemaker produces *ticks* at a regular rate (e.g. neuron spikes) that are tallied by an accumulator system. Since the pacemaker "ticks" at a regular rate, the accumulates "ticks" can be used to tell time (think of this like the accumulation of sand in an hourglass). Second, a memory system stores the typical durations of common events and also allows for the short-term storage of particular events for comparison with ongoing events. Finally, a decision process compares the outputs of the accumulator with the stored representations in the memory system to determine whether the current reading is accurate and whether the perceived event is longer, shorter, or of equal duration to ones stored in memory.

While SET has been able to explain a large swath of temporal perception data, the flexibility the theory has in explaining this data is both a blessing and a curse. As Wearden (1999) emphasizes, the individual components of SET are sufficiently underspecified that almost any data can be accommodated by modifying the components of the theory. Furthermore, there are reasons to believe that the core of the theory, a centralized pacemaker-accumulator system, lacks biologically plausibility. Eagleman and Pariyadath (2009) cite evidence that no clock of this sort has been found despite years of searching. Similarly, it has been argued that no neural mechanism is sufficiently precise to be able to account for the reliability of temporal perception at scales spanning approximately 30ms to multiple seconds in length (Butts et al., 2007; Eagleman & Pariyadath, 2009; Panzeri et al, 2014).[19]

While many have rejected SET, or parts of SET, many of the alternative proposals still assume that there will be a single mechanism or explanation for temporal perception. Examples of these single mechanism approaches include the developments of SET (Hinton & Meck, 1997; Matthews et al., 2011), *the striatal beat frequency models* (Matell & Meck, 2000, 2004), and *default mode network models* (Carvalho et al., 2016; Lloyd, 2012). Other theories go further and reject that a single token mechanism can underpin temporal perception, but nevertheless assume that a single *type* of mechanism or explanation will suffice. Examples of this sort include *efficiency coding models* (Eagleman & Pariyadath, 2009) and *state-dependent network models* (Buonomano, 2000, 2005; Finnerty et al., 2015; Karmarkar & Buonomano, 2007).

One explanation for the apparent disagreement between these models is that *we have yet to uncover the correct model*. However, another possibility is that *there is no genuine conflict at all*. Rather, temporal experience is composed of many distinct capacities to perceive specific aspects of the temporal structure of the world, and different models apply to different component mechanisms. It is this second approach that I'll develop here.

Some researchers have recently started to adopt this approach. As Wittman and van Wassenhove put it:

> In recent years, new ideas have emerged regarding the neurobiological mechanisms underlying the experience of time. In particular, major progress has been made with the realization that

[19] The general idea is that there is a trade-off between resolution and noise. If a system that is capable of discriminating durations as short as 30ms, then the worry is that whatever representational process is employed at that scale cannot be scaled up without significant amounts of noise – more than what would be expected from the Weber fraction.

time perception may engage distinct brain mechanisms (and areas) depending on the time scale at which events occur…

(Wittmann & Wassenhove, 2009, p. 1809)[20]

One aspect of Wittman and van Wassenhove's argument for this divide has to do with the differential contribution of systems in the perception of time above and below one second. Suprasecond timing seems to be sensitive to attention, emotion, and bodily movement, while subsecond timing does not seem to be sensitive to these factors (Wittmann & Wassenhove, 2009). However, a problem with this argument is that It may simply be the case that in suprasecond timing these auxiliary systems have more time to be utilized by a single timekeeping mechanism, while subsecond timing tasks simply do not allow for these systems to contribute. As a result, the behavioral differences due to timescale changes may be accounted for by a single timing mechanism.

Instead, like most attempts to show that multiple mechanisms underpin some psychological phenomenon, the main evidence for the fragmentary model comes from specific dissociations between timing capacities and the lack of any unilateral timing deficits.[21]

Through both pharmacological and mechanical interventions, it is possible to selectively disturb aspects of temporal perception. Rammsayer (1999) showed that haloperidol, a dopamine receptor antagonist used in the treatment of schizophrenia, and midazolam, a benzodiazepine used in the treatment of seizures and insomnia, both impair the discrimination of temporal properties of approximately 1 second in length. However, only haloperidol impaired the discrimination of durations of approximately 50ms. Similarly, the hallucinogen psilocybin seems to leave the ability to perceive time at scales of less than 2 seconds unimpaired, but severely impairs the perception of longer timescales (Wittmann et al., 2007).

Similar effects are found with the selective application of *rapid transcranial magnetic stimulation* (rTMS). Koch et al (2007) showed that the application of rTMS to dorsolateral frontal cortex impaired timing in the seconds range, whereas rTMS to the cerebellum impaired millisecond timing (Jones et al., 2004).

Clinical populations also show interesting dissociations. Schizophrenia and autism are associated with distinct disturbances in the ability to perceive temporal order.[22] Parkinson's patients, on the other hand, show deficits in short timescale timing while longer timescales remain untouched (Riesen & Schnider, 2001).

Finally, modality specific temporal illusions suggest even more fine-grained dissociations. Burr et al (2007) showed that saccades compress the perceived duration of visual stimuli yet leave temporal perception in other modalities untouched. Importantly, the duration compression resulting from

[20] Eagleman (2008) and Grush (2016) both argue that different temporal properties are perceived via different types of mechanisms.

[21] While Wittman and van Wassenhove also appeal to dissociations in their argument, they primarily focus on pharmacological and mechanical interventions.

[22] People can typically tolerate small temporal asynchronies in stimuli that are nevertheless perceived as simultaneous – e.g. when the audio and visual tracks of a movie are out of step (Dixon & Spitz, 1980). The perception of simultaneity in people with schizophrenia allows for greater asynchrony (Martin et al, 2013), while people on the autism disorder spectrum allow for less asynchrony (Stevenson et al., 2014).

saccades does not influence the perception of rhythms or rates (Eagleman, 2008), suggesting that the perception of different temporal property-types relies on different mechanisms.

The flipside to these selective deficits is the absence of any global deficits in temporal perception. A thorough search of the literature has failed to bring up a single case of someone that completely lacks the capacity to perceive time. Even people with severe cases of anterograde amnesia still retain some ability to perceive time in that they are able to perform complicated sensorimotor tasks, and understand some speech, which some capacity for temporal perception.[23]

All this evidence points towards there being multiple distinct mechanisms that underpin our ability to perceive time. If we identify psychological capacities through a multi-level approach, then it follows from the behavioral and neurological evidence that temporal perception is not a single unified phenomenon, but is instead fragmented and composed of many dissociable capacities to perceive various temporal aspects of the world.

While this evidence establishes the existence of multiple token mechanisms, it does not establish that the mechanisms underpinning these distinct capacities are type-distinct. In the next section, we'll consider evidence suggesting that temporal experience relies on type-distinct mechanisms, some of which satisfy mirroring, while others do not.

5. Two cases: against the atomism / extensionalism debate

The two cases presented in this section respectively provide counterexamples to atomism and extensionalism. The difference, however, between the cases presented here and other purported counterexamples in the literature is that the behavioral / perceptual data will be considered alongside our best understanding of the relevant neural mechanisms. This behavioral and mechanistic evidence constrains our theorizing in such a way that we can make progress on the atomism / extensionalism debate.

Case #1: The unimodal perception of duration

Until relatively recently, the models, notably SET, proposed to account for the perception of duration within individual modalities at sub-second timescales have employed *extrinsic* timing mechanisms. In addition to the various sensory specific mechanisms that process non-temporal aspects of stimuli these models posit dedicated timekeeping mechanisms that bind temporal information to the non-temporal representations delivered by the specific senses. However, alternative models that do away with dedicated timekeeping mechanisms have begun to gain traction.

These *intrinsic* timing models explain temporal perception as the result of the *intrinsic* properties of neural systems that also encode for non-temporal aspects of stimuli. [24] Some intrinsic

[23] For a discussion of relevant amnesia cases see (Craver, Graham, & Rosenbaum, 2014; Craver, Kwan, Steindam, & Rosenbaum, 2014).
[24] See (Ivry & Schlerf, 2008) for general discussion of *intrinsic* models.

timing models see temporal perception as being an intrinsic property of *particular* neural systems.[25] Others see intrinsic timing mechanisms as a much more ubiquitous feature of the brain. It is in looking at these more ubiquitous models that we'll find our counterexample to atomism.

Of the intrinsic timing models, *the state-dependent network* (SDN) *models* have gained significant traction (Buonomano, 2000b; 2005; Buonomano & Maass, 2009; Finnerty et al., 2015; Mauk & Buonomano, 2004). The basic idea behind these models can be illustrated through an analogy (Karmarkar & Buonomano, 2007). Imagine a pebble falling into a still pool of water. The pebble's impact causes ripples to travel across the water's surface. Provided enough information about the pebble, and how it impacted the water's surface, you would be able to use the state of the ripples as a means of determining how much time had passed since the pebble was dropped. Now, imagine a second pebble being dropped into the water. This pebble would cause a change in the ripples in the surface such that if you knew enough about the state of the ripples due to the first pebble and the characteristics of the second pebble, then you would be able to use the ripples again to determine how much time had passed since the dropping of the second pebble. As Goel and Buonomano put it, "ripples thus establish a short-lasting and dynamic memory of the recent stimulus history of the liquid, and it is possible to estimate the amount of time elapsed based on the current state of the liquid." (Goel & Buonomano, 2014).

SDN models appeal to a similar *rippling* found in biologically inspired artificial neural networks and in vitro neural populations (Finnerty et al., 2015). The *ripples* in these systems consist in the subset of neurons within the network that are active at any given moment (call this the system's *active state*). As time passes, the active states change as a result of the time-dependent evolution of the mechanisms for short-term synaptic plasticity that govern how the modulation of the system's active states (Buonomano & Maass, 2009; Goel & Buonomano, 2014). In this way, when an incoming signal acts as the pebble and activates the system, the time-dependent evolution of the active states of the system (i.e. the ripples) carries information about how much time has passed.

Since these timekeeping properties are due to network internal structure, there is theoretically no need for dedicated extrinsic timing mechanisms. Yet, showing that networks have these properties isn't enough to show that this is how animals perceive time. Downstream processing might simply not exploit this information. What we need is some reason to think that timing behaviors depend on this type of mechanism.

The most straightforward way of determining whether some property of a neural system is what bears the semantic contents of some mental representation is to see whether there are any peculiar behavioral / functional effects that arise from those properties being the bearers of content. In this particular case, the behavioral / functional effects arise from the context dependent nature of the activity in the neural networks. The accuracy with which we perceive time is a function of the variability of the temporal interval between target stimuli and pre-target stimuli for very short durations (up to 350ms) presented within individual modalities. Think of this variability as effecting the activation of the network in the same way that the temporal interval between the first and second pebbles that impact the water has on the rippling of the water's surface. The accuracy of temporal judgments of a tone's

[25] For example, Lebedev et al (2008) see timing as the result of ramping activity in pre-motor areas. However, the same conclusions would be reached by looking at their model.

duration was directly influenced by whether the target tone was presented within a fixed or variable context (i.e. whether the temporal delay between the target tone and lead-in tone was fixed or varied). The findings of Karmarkar & Buonomano (2007) were further confirmed by Ivry et al (unpublished work cited in (Ivry & Schlerf 2008)). However, when the duration of the target tone was increased beyond 300ms, the context effects were lost, indicating that at this timescale, duration judgments rely on mechanisms other than SDN mechanisms.

While extrinsic models, like SET, *could* account for this data, they would not predict this sort of context dependency. Furthermore, the types of changes to SET needed to account for this context dependency would be entirely ad hoc. Given that we have reasons for thinking that neural networks behave in the way that the SDN models require, and that the extrinsic models require ad hoc adjustments to explain the data, there is no need to posit an extrinsic timekeeping mechanism in addition to the machinery that an intrinsic model requires.

Let's turn back to the mirroring debate. Consider the experience of a tone as lasting 80ms. This experience is realized by a temporally extended neural process. Furthermore, as described in the discussion of the trace integration argument, the experience of the tone as lasting 80ms *contains* experiences as part of its content (i.e. the tone lasting 70ms, 60ms, 50ms, etc). If the neural process that realizes the experience of the tone as lasting 80ms can be divided into temporal parts that each correspond to the contained experiences, and the temporal relations between these parts of the experience are mirrored by the represented temporal relations between the contained experiences, then mirroring will be satisfied.

According to SDN models the extended neural process that realizes the experience of the tone lasting 80ms is divisible into temporal parts (i.e. distinct active states) that represent the tone as having specific durations. Given how changes in the system's temporal content and changes in the system's active states are linked together, the temporal relations that hold between the temporal stages of the experience will be mirrored by the temporal relations represented in the experience. The temporal relation that holds between the *experiences of the tone* as lasting 70ms and 80ms, will be mirrored by the temporal relations that are *represented as holding between the tone lasting 70ms and 80ms*. That is, mirroring holds in this case.

One could object, however, that at the moment that one is having the experience of the tone as lasting 80ms, the process *at that moment* serves as a form of memory for the entire experienced 80ms interval, thereby violating mirroring (since the represented extended interval would be represented *now*).[26] This is actually a common form of atomist argument.[27] If it can be shown that some aspect of our temporal experience relies on memory that presently encodes some temporal interval, then the mirroring constraint must be violated. In large part, this strategy is generally effective because there seems to be nothing in principle that requires memory systems to be temporally structured in any particular way to encode information. Therefore, once memory systems are on the table, it is easy to construct an account of those memory systems that violate mirroring.

However, this is exactly where mechanisms matter. Atomists cannot insist that we evaluate the mirroring constraint with respect to the very current state of the system, since as discussed in section 2,

[26] Thank you for an anonymous referee at this journal for pushing this objection.
[27] A notable example is found in Grush (2006).

atomists cannot claim that experiences are durationless (and constrained to the *mere present*). Experiences are temporally extended processes. So, when we consider whether the experience of the tone satisfies the mirroring constraint, we need to look at the neural process that constitutes that experience (not a mere timeslice of that process). Once we do this, however, the mechanism by which the SDN models account for the perception of duration shows us that in this case there is reason for thinking that the processes by which *this type of memory* is formed requires the memory system to possess a particular temporal structure that satisfies the mirroring constraint. Memory in general is not a threat to mirroring. It all depends how those memories are structured.

As a result, despite all atomists have said, here we have an empirically informed case of mirroring. Atomism, cannot provide a general characterization of temporal experience.

Case #2: Multimodal perception of order and simultaneity

We just saw that the perception of duration at short timescales within individual modalities satisfies mirroring. Now, we'll see that the multimodal perception of temporal order violates mirroring, by focusing on *temporal order judgments* (TOJ) at the timescale of less than one second.

For a pair of stimuli to be perceived as simultaneous, they needn't objectively simultaneous. In fact, the stimuli don't even have to simultaneously impact sensory receptors to be perceived as simultaneous. When you see, hear, and feel, your foot kicking a ball, the various sensory signals impact their transducers, and arrive at cortical sensory areas, at different times – it takes longer for signals from the foot to reach the brain that they do from the eye. Yet, we perceive these aspects of the event as occurring simultaneously. Furthermore, as our bodies change (e.g. a child's leg grows) these temporal asynchronies change, and therefore TOJs must be flexible.

One particular form of flexibility in TOJs arises in the form of adaptation effects (Cai et al., 2012; Stetson et al., 2006). Similar to how constant presentation of a visual stimulus alters visual responses, similar adaptations occur with the perception of temporal order. By consistently inserting a short delay between a motor action (e.g. pressing a button) and a visual event (e.g. a flash of light) the two events come to be perceived as occurring closer in time than they in fact do. Once this initial recalibration occurs, and the delay is removed (or reduced), and the stimuli are presented objectively closer in time, the effect is that the stimuli are perceived as occurring *much closer* together than they in fact are. The perceptual system seems to be compensating for the no-longer-present delay. In some cases, the effect is so dramatic that the supposed effect (i.e. the flash of light) is perceived as occurring prior to the supposed cause (i.e. the pressing of the button)! Our experience of temporal order is so sensitive to recalibration effects that the perceived order of events can drastically differ from the actual order of events in the world and the order in which these events impinge upon the sensory receptors.[28]

Whatever mechanism we propose to account for TOJs would seem to cause problems for the mirroring constraint as perceived temporal order becomes more a matter of the state of calibration than the timing of stimulus processing. Yet, extensionalists could try and accommodate these findings by arguing that the temporal structure of the sensory processing leading up to experience is not identical to the temporal structure of experience. They could insist, as they do with explanations of postdictive

[28] Similar effects have been found with other stimuli (Vroomen & Keetels, 2010; Vroomen et al., 2004).

effects (Dainton, 2014b), that the seemingly mirroring violating sensory processing is *pre-experiential* and that only later are they integrated to form a coherent experience of time that plays out in time in a mirroring friendly way. Therefore, recalibration effects say nothing about experience's temporal structure, since recalibration occurs pre-experientially. However, once again, this is where mechanisms matter and show how this extensionalist reply fails.

Cai et al (2012) propose that TOJs are underpinned by an opponency system. The individual sensory signals from the distinct sensory systems are fed into a collection of delay-tuned neurons responsive to particular temporal asynchronies (e.g. one may be tuned for *signal a trailing signal b by 20ms*, another for *signal a trailing signal b by 10ms*, another for *signal b trailing signal a by 20ms*, etc). These neurons project to a pair of opponent summation nodes that sum the signals emerging from the individual delay-tuned neurons. One summation node signals that *a precedes b* while the other signals that *b precedes a* and the relative activation of these opponent cells signals whether *a* or *b* precedes the other and by what interval. This opponency system can then explain the recalibration of TOJs due to the adaptation of either delay-tuned neurons or the summation nodes themselves.

The temporal content needed to explain TOJs is accounted for by these opponency mechanisms. There is no need for an additional *experiential phase* of processing to get temporal order content into perception. Furthermore, the opponency mechanism shows us how identically presented stimuli may be perceived as standing in very different temporal orders because of calibration effects and not because of differences in the temporal structure of sensory processing. As a result, we have a case in which mirroring is violated.

The initial sensory processing cannot be divided into temporal stages that satisfy mirroring because the initial, stimulus dependent, sensory processing can remain identical despite changes in perceived temporal order. Similarly, we cannot divide the operation of the opponency mechanism into temporal parts that satisfy mirroring, since the ramping of activity in the delay-tuned neurons and the opponency nodes are not themselves representational, and therefore do not correspond to, let alone mirror, the content of experience. Finally, even though the current state of calibration is a result of prior experiences, the dependence of experience on these prior experiences does not save mirroring. Mirroring requires more than a mere dependence between the contents of current experience and past experience, and the prior calibrating experiences do not correspond to any component of the currently experienced temporal sequence.

On no interpretation of how this system carries information about time do we have an interpretation that satisfies the mirroring constraint. As a result, we have a counterexample to extensionalism, and therefore, neither atomism nor extensionalism can provide us with a general theory for how the temporal contents of experience relate to the temporal structure of experience.

6. Conclusion

The original puzzle about temporal experience was motivated by a particular type of puzzlement – how could experience come to have its temporal contents. In trying to resolve this puzzle, we saw that the philosophical literature turned to a discussion of the mirroring constraint. Yet the arguments put forward in favor of atomism and extensionalism fail in establishing their intended conclusions.

This failure, however, shouldn't come as a surprise. Temporal experience is not a single phenomenon, but is instead composed of distinct capacities to represent time underpinned by type-distinct representational mechanisms. It is only by considering the behavioral and perceptual evidence, in light of what we know about the neural mechanisms that underpin these capacities, that we can make progress regarding the mirroring constraint. No single story can be given for how the temporal contents of experience relate to the temporal structure of experience. Both atomism and extensionalism are false as general theories of how the temporal content of experience relates to the temporal structure of experience.

What then of the puzzle? How is it possible to have experiences with temporal contents? Answering that question gets us to central questions concerning mental representation. A theory of content is needed that can explain how the various timekeeping mechanisms and processes of the brain come to carry information about time. In this paper, I've suggested that perhaps an information-theoretic analysis could explain how this occurs, as our discussions have been loosely framed in information-theoretic terms. However, this is where the work on temporal experience has to make progress. What we need is a theory of content that meshes with our neuroscientific understanding of the mechanisms that underpin temporal perception.

References

Boroditsky, L. (2011). How Languages Construct Time. In S. Dehaene & E. Brannon (Eds.), *Space, Time, and Number in the Brain* (pp. 333–341). Elsevier.

Buonomano, D. V. (2000a). Decoding temporal information: A model based on short-term synaptic plasticity. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *20*(3), 1129–1141.

Buonomano, D. V. (2000b). Decoding temporal information: A model based on short-term synaptic plasticity. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *20*(3), 1129–1141.

Buonomano, D. V. (2005). A Learning Rule for the Emergence of Stable Dynamics and Timing in Recurrent Networks. *ResearchGate*, *94*(4), 2275–83. https://doi.org/10.1152/jn.01250.2004

Buonomano, D. V., & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, *10*(2), 113–125. https://doi.org/10.1038/nrn2558

Burr, D., Tozzi, A., & Morrone, M. C. (2007). Neural mechanisms for timing visual events are spatially selective in real-world coordinates. *Nature Neuroscience*, *10*(4), 423–425. https://doi.org/10.1038/nn1874

Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M., & Stanley, G. B. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature*, *449*(7158), 92–95. https://doi.org/10.1038/nature06105

Cai, M., Stetson, C., & Eagleman, D. M. (2012). A Neural Model for Temporal Order Judgments and Their Active Recalibration: A Common Mechanism for Space and Time? *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00470

Carvalho, F. M., Chaim, K. T., Sanchez, T. A., & Araujo, D. B. de. (2016). Time-Perception Network and Default Mode Network Are Associated with Temporal Prediction in a Periodic Motion Task. *Frontiers in Human Neuroscience*, *10*.

Chuard, P. (2011). Temporal Experiences and Their Parts. *Philosophers' Imprint*, *11*(11).

Craver, C. F., Graham, B., & Rosenbaum, R. S. (2014). Remembering Mr B. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, *59*, 153–184. https://doi.org/10.1016/j.cortex.2013.11.001

Craver, C. F., Kwan, D., Steindam, C., & Rosenbaum, R. S. (2014). Individuals with episodic amnesia are not stuck in time. *Neuropsychologia*, *57*, 191–195. https://doi.org/10.1016/j.neuropsychologia.2014.03.004

Dainton, B. (2008). The Experience of Time and Change. *Philosophy Compass*, *3*(4), 619–638. https://doi.org/10.1111/j.1747-9991.2008.00153.x

Dainton, B. (2011). Time, Passage and Immediate Experience. In C. Callender (Ed.), *Oxford Handbook of Philosophy of Time* (p. 382). Oxford University Press.

Dainton, B. (2014a). Temporal Consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014). Retrieved from http://plato.stanford.edu/archives/spr2014/entries/consciousness-temporal/

Dainton, B. (2014b). Temporal Consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014). Retrieved from http://plato.stanford.edu/archives/spr2014/entries/consciousness-temporal/

Dixon, N. F., & Spitz, L. (1980). The detection of auditory visual desynchrony. *Perception*, *9*(6), 719–721.

Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press.

Dretske, F. (1994). Introspection. *Proceedings of the Aristotelian Society*, *94*(n/a), 263–278.

Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.

Eagleman, D. M. (2008). Human time perception and its illusions. *Current Opinion in Neurobiology*, *18*(2), 131–136. https://doi.org/10.1016/j.conb.2008.06.002

Eagleman, D. M., & Pariyadath, V. (2009). Is subjective duration a signature of coding efficiency? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1525), 1841–1851. https://doi.org/10.1098/rstb.2009.0026

Finnerty, G. T., Shadlen, M. N., Jazayeri, M., Nobre, A. C., & Buonomano, D. V. (2015). Time in Cortical Circuits. *The Journal of Neuroscience*, *35*(41), 13912–13916. https://doi.org/10.1523/JNEUROSCI.2654-15.2015

Fodor, J. A. (1994). *The Elm and the Expert*. MIT Press.

Gallistel, C. R. (1996). The Perception of Time. In K. Akins (Ed.), *Perception: Vancouver Studies in Cognitive Science* (Vol. 5). Oxford University Press.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, *84*(3), 279–325. https://doi.org/10.1037/0033-295X.84.3.279

Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar Timing in Memory. *Annals of the New York Academy of Sciences*, *423*(1 Timing and Ti), 52–77. https://doi.org/10.1111/j.1749-6632.1984.tb23417.x

Goel, A., & Buonomano, D. V. (2014). Timing as an intrinsic property of neural networks: evidence from in vivo and in vitro experiments. *Phil. Trans. R. Soc. B*, *369*(1637), 20120460. https://doi.org/10.1098/rstb.2012.0460

Grush, R. (2005). Brain Time and Phenomenological Time. In A. Brooks & K. Akins (Eds.), *Philosophy and the Neurosciences*. Cambridge.

Grush, R. (2006). How to, and How Not to, Bridge Computational Cognitive Neuroscience and Husserlian Phenomenology of Time Consciousness. *Synthese*, *153*(3), 417–450.

Grush, R. (2016). On the Temporal Character of Temporal Experience, its Scale Non-Invariance, and its Small Scale Structure.

Hartcher-O'Brien, J., Brighouse, C., & Levitan, C. A. (2016). A single mechanism account of duration and rate processing via the pacemaker–accumulator and beat frequency models. *Current Opinion in Behavioral Sciences*, *8*, 268–275. https://doi.org/10.1016/j.cobeha.2016.02.026

Hinton, S. C., & Meck, W. H. (1997). The "internal clocks" of circadian and interval timing. *Endeavour*, *21*(2), 82–87.

Hoerl, C. (2009). Time and Tense in Perceptual Experience. *Philosophers' Imprint*, *9*(12), 1–18.

Hume, D. (2000). *A Treatise of Human Nature*. Oxford University Press.

Ivry, R. B., & Schlerf, J. E. (2008). Dedicated and intrinsic models of time perception. *Trends in Cognitive Sciences*, *12*(7), 273–280. https://doi.org/10.1016/j.tics.2008.04.002

Jones, C. R. G., Rosenkranz, K., Rothwell, J. C., & Jahanshahi, M. (2004). The right dorsolateral prefrontal cortex is essential in time reproduction: an investigation with repetitive transcranial magnetic

stimulation. *Experimental Brain Research*, *158*(3), 366–372. https://doi.org/10.1007/s00221-004-1912-3

Karmarkar, U. R., & Buonomano, D. V. (2007). Timing in the Absence of Clocks: Encoding Time in Neural Network States. *Neuron*, *53*(3), 427–438. https://doi.org/10.1016/j.neuron.2007.01.006

Kelly, S. D. (2005). The Puzzle of Temporal Experience. In A. Brook (Ed.), *Cognition and the Brain: The Philosophy and Neuroscience Movement* (pp. 208–238). Cambridge: Cambridge University Press.

Koch, G., Oliveri, M., Torriero, S., Salerno, S., Lo Gerfo, E., & Caltagirone, C. (2007). Repetitive TMS of cerebellum interferes with millisecond time processing. *Experimental Brain Research*, *179*(2), 291–299. https://doi.org/10.1007/s00221-006-0791-1

Lebedev, M. A., O'Doherty, J. E., & Nicolelis, M. A. L. (2008). Decoding of temporal intervals from cortical ensemble activity. *Journal of Neurophysiology*, *99*(1), 166–186. https://doi.org/10.1152/jn.00734.2007

Lee, G. (2014a). Extensionalism, Atomism, and Continuity. In N. Oaklander (Ed.), *Debates in the metaphysics of time*. Bloomsbury.

Lee, G. (2014b). Temporal Experience and the Temporal Structure of Experience. *Philosophers' Imprint*, *14*(3).

Lloyd, D. (2012). Neural Correlates of Temporality: Default Mode Variability and Temporal Awareness. *Consciousness and Cognition*, *21*(2), 695–703.

Locke, J. (1689). *An Essay Concerning Human Understanding*. Hackett Publishing.

Martin, B., Giersch, A., Huron, C., & van Wassenhove, V. (2013). Temporal event structure and timing in schizophrenia: preserved binding in a longer "now." *Neuropsychologia*, *51*(2), 358–371. https://doi.org/10.1016/j.neuropsychologia.2012.07.002

Martin, M. G. F. (2002). The Transparency of Experience. *Mind and Language*, *4*(4), 376–425.

Matell, M. S., & Meck, W. H. (2000). Neuropsychological mechanisms of interval timing behavior. *BioEssays: News and Reviews in Molecular, Cellular and Developmental Biology*, *22*(1), 94–103. https://doi.org/10.1002/(SICI)1521-1878(200001)22:1<94::AID-BIES14>3.0.CO;2-E

Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Brain Research. Cognitive Brain Research*, *21*(2), 139–170. https://doi.org/10.1016/j.cogbrainres.2004.06.012

Matthews, W. J., Stewart, N., & Wearden, J. H. (2011). Stimulus intensity and the perception of duration. *Journal of Experimental Psychology. Human Perception and Performance*, *37*(1), 303–313. https://doi.org/10.1037/a0019961

Mauk, M. D., & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, *27*, 307–340. https://doi.org/10.1146/annurev.neuro.27.070203.144247

Millikan, R. G. (1989). Biosemantics. *Journal of Philosophy*, *86*(July), 281–97.

Millikan, R. G. (1995). Pushmi-Pullyu Representations. *Philosophical Perspectives*, *9*, 185–200.

Panzeri, S., Ince, R. A. A., Diamond, M. E., & Kayser, C. (2014). Reading spike timing without a clock: intrinsic decoding of spike trains. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *369*(1637), 20120467. https://doi.org/10.1098/rstb.2012.0467

Phillips, I. (2010). Perceiving Temporal Properties. *European Journal of Philosophy*, *18*(2), 176–202.

Phillips, I. (2014a). Breaking the Silence: Motion Silencing and Experience of Change. *Philosophical Studies*, *168*(3), 693–707.

Phillips, I. (2014b). Experience of and in Time. *Philosophy Compass*, *9*(2), 131–144.

Prinz, J. J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. MIT Press.

Rammsayer, T. H. (1999). Neuropharmacological evidence for different timing mechanisms in humans. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, *52*(3), 273–286. https://doi.org/10.1080/713932708

Rashbrook, O. (2013). An Appearance of Succession Requires a Succession of Appearances. *Philosophy and Phenomenological Research*, *87*(3), 584–610.

Reid, T. (1855). *Essays on the Intellectual Powers of Man*. Early Modern Texts Online.

Reutimann, J., Yakovlev, V., Fusi, S., & Senn, W. (2004). Climbing Neuronal Activity as an Event-Based Cortical Representation of Time. *The Journal of Neuroscience*, *24*(13), 3295–3303. https://doi.org/10.1523/JNEUROSCI.4098-03.2004

Riesen, J. M., & Schnider, A. (n.d.). Time estimation in Parkinson's disease: normal long duration estimation despite impaired short duration discrimination. *Journal of Neurology*, *248*(1), 27–35. https://doi.org/10.1007/s004150170266

Ryder, D. (2004). SINBAD Neurosemantics: A Theory of Mental Representation. *Mind & Language*, *19*(2), 211–240. https://doi.org/10.1111/j.1468-0017.2004.00255.x

Skyrms, B. (2010). *Signals*. Oxford University Press. Retrieved from https://global.oup.com/academic/product/signals-9780199582945

Stampe, D. W. (1977). Towards a Causal Theory of Linguistic Representation. *Midwest Studies in Philosophy*, *2*(1), 42–63.

Stetson, C., Cui, X., Montague, P. R., & Eagleman, D. M. (2006). Motor-Sensory Recalibration Leads to an Illusory Reversal of Action and Sensation. *Neuron*, *51*(5), 651–659. https://doi.org/10.1016/j.neuron.2006.08.006

Stevenson, R. A., Siemann, J. K., Schneider, B. C., Eberly, H. E., Woynaroski, T. G., Camarata, S. M., & Wallace, M. T. (2014). Multisensory Temporal Integration in Autism Spectrum Disorders. *The Journal of Neuroscience*, *34*(3), 691–697. https://doi.org/10.1523/JNEUROSCI.3615-13.2014

Treisman, M. (1963). Temporal discrimination and the indifference interval. Implications for a model of the "internal clock." *Psychological Monographs*, *77*(13), 1–31.

Usher, M. (2001). A Statistical Referential Theory of Content: Using Information Theory to Account for Misrepresentation. *Mind & Language*, *16*(3), 311–334. https://doi.org/10.1111/1468-0017.00172

Vroomen, J., & Keetels, M. (2010). Perception of intersensory synchrony: A tutorial review. *Attention, Perception, & Psychophysics*, *72*(4), 871–884. https://doi.org/10.3758/APP.72.4.871

Vroomen, J., Keetels, M., de Gelder, B., & Bertelson, P. (2004). Recalibration of temporal order perception by exposure to audio-visual asynchrony. *Brain Research. Cognitive Brain Research*, *22*(1), 32–35. https://doi.org/10.1016/j.cogbrainres.2004.07.003

Watzl, S. (2012). Silencing the experience of change. *Philosophical Studies*, *165*(3), 1009–1032. https://doi.org/10.1007/s11098-012-0005-6

Wearden, J. H. (1999). "Beyond the fields we know...": exploring and developing scalar timing theory. *Behavioural Processes*, *45*(1–3), 3–21.

Wittmann, M., Carter, O., Hasler, F., Cahn, B. R., Grimberg, U., Spring, P., … Vollenweider, F. X. (2007). Effects of psilocybin on time perception and temporal control of behaviour in humans. *Journal of Psychopharmacology (Oxford, England)*, *21*(1), 50–64. https://doi.org/10.1177/0269881106065859

Wittmann, M., & Wassenhove, V. van. (2009). The experience of time: neural mechanisms and the interplay of emotion, cognition and embodiment. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *364*(1525), 1809–1813. https://doi.org/10.1098/rstb.2009.0025