

Experiencing Temporal Properties

Thomas Blackburn

UCL

MPhil Stud. Philosophy

Declaration

I, Thomas Blackburn, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed: Thomas Blackburn

Abstract

The topic of this essay is a class of properties that are naturally predicated both of the objects of experience, and experiences themselves: temporal properties. I begin by introducing the notion of a temporal property, before arguing that a basic ‘snapshot’ picture of temporal experience cannot accommodate our experience of these properties as properly conceived. Having completed this preliminary work, in chapter two I set out some of the key distinctions that will inform the argument of the essay; Dainton’s distinction between ‘extensionalism’ and ‘retentionalism’, and more importantly the notion of ‘temporal mirroring’. In chapter three, I focus on a major argument for mirroring from Phillips (2010), from the ‘transparency’ of experience. I present several ways by which this argument should be resisted, the main reason being that the notion of ‘transparency’ employed is not sufficiently clear to yield the conclusion Phillips requires. The discussion in chapter three also considers the notion of ‘seeming’ with regard to experience, and whether Phillips’ argument can be resisted by clarifying the notion of seeming at stake in the argument. In chapter four, I focus on some well-known empirical cases which I argue provide a further case against temporal mirroring; the Color-phi case and the case of the ‘cutaneous rabbit’. I consider a major response to these cases from Phillips (2014a), who argues that we should reconceive of experience as *anhomoeomeric*, and argue that this response faces numerous problems. The main conclusion of the essay is that we have good reason to reject temporal mirroring. In the final chapter, I briefly explore further issues that arise from the debate I have been considering, concluding that questions over the unity and continuity of experience over time still remain to be adequately addressed by all parties to the debate.

Experiencing Temporal Properties

Table of Contents

1. Temporal properties 9

- Introduction 9
- Temporal, spatial and colour properties 9
- Constraints on temporal experience 15
- The cinematic view 20

2. Mirroring and inheritance 26

- Dainton's debate 26
- Mirroring 32
- Mirroring, resemblance and inheritance 35

3. Mirroring and transparency 41

- Introduction 41
- Phillips' argument from transparency 42
- Analysis of Phillips' argument 44
- Different views on access 48
- Sensory deprivation and covert attention shifts 51
- Different kinds of seeming 55
- Conclusion 57

4. Postdiction cases and temporal experience 59

- Introduction 59
- Postdiction cases 60
- Interpreting postdiction cases: Orwellian vs Stalinesque 62
- A new conception of experience as anhomoeomeric 66
- The temporal correlation principle 71
- Russell worlds 75
- Conclusion 79

5. Conclusion and further issues: the diachronic unity of experience 80

- Conclusion 80
- Extensionalism and retentionalism revisited 82
- Unity and overlap 84
- Unity and continuity of experience over time 90

Bibliography 93

1. Temporal properties

Introduction

The topic of this essay is the question of how we perceive temporal properties, and more specifically the nature of the relation between the temporal properties (or, temporal structure) of the events we experience and the temporal properties of our experiences themselves. The literature relevant to this broad philosophical area is large and disparate, with numerous authors choosing to set out the central problems, and the available responses, in different ways. For example, Dainton (2000, 2010a) chooses to focus on what he terms ‘extensional’ and ‘retentional’ views of experience in his influential approach to temporal consciousness. Others, such as Lee (2014a) choose to focus on the relationship between experiences and their neural (physical) realizers instead, and implicitly rejects the taxonomy that Dainton has proposed. In this first chapter, I will set up what I take to be the central problem of temporal experience and in doing so focus on the *relation* between the two temporal structures mentioned above, rather than specifically the nature of experiences themselves (though that question will become relevant). Before, I set up the central problem, however, I first wish to discuss what I take to be some important features of our temporal experience that will inform the discussion throughout the essay.

Temporal, spatial and colour properties

In everyday scenarios, we are in constant experiential contact with a host of temporal phenomena. A simple walk down the street typically involves apprehending all sorts of temporal properties and relations; the sound of one bird’s call following another, the changing appearance of a tree as it blows in the wind, and the long wail of a siren as an ambulance drives past. In being acquainted with these features of the world, we are presented with temporal properties of objects and events; or, to use the language of representationalism, we represent temporal properties in our experience, properties that

are necessarily instantiated over time. The succession of one birdcall after another, the changing arrangement of the leaves on the tree, and the duration of the ambulance siren are all temporal properties that feature immediately in our experience, arguably as clearly as other properties of objects such as colour, or size (Dainton 2010a). I will take these kinds of experiences to be part of the data that a good theory of temporal experience should be able to account for, the implication being that any theory of experience (or consciousness more generally) that does not take seriously this basic phenomenology should be called into doubt. In Phillips' (2014a) terms, this way of conceiving of the data to be explained concerning temporal experience signifies a strong 'realist' commitment to temporal properties, and this is a commitment that most perceptual theorists (if not all) accept.

Our everyday experience of the world seems to involve the representation (or presentation) of a class of properties I will term 'temporal properties'. As previously mentioned, I take temporal properties to be the sort of properties that are necessarily instantiated over an interval of time: the key examples that will play a central role in the argument of this essay are succession (a kind of change), and duration, which are temporal properties of events. Whilst there may plausibly be more properties to add to this list, for the purposes of this essay I will consider just these two properties, taking them to be paradigm examples of the kinds of temporal properties we ordinarily experience. In setting up the central problem to which the argument of my essay will pertain, I first want to make more explicit the kind of claim on the part of the realist about temporal properties, with reference to the cases of succession, duration and motion.

Succession: consider hearing three tones, one after the other: A, B and C#, forming a seamless sequence. It seems very natural to say that there are two data perceived here; the tones themselves, but also their succession; we perceive B as succeeding A and C# as succeeding B. If this succession was merely inferred from the perception of A, B and C# (or perceived indirectly in any other relevant sense), the experience of the three

notes would not have the distinctive phenomenal character that it has; it would be no different phenomenally to the perception of A, B and C# each as 5 minutes apart, say. However, the experience of A, B and C# in succession is different; it has a distinctive phenomenal character which captures not only the notes and their respective temporal intervals, but also their succession, a relational property between events. (Phillips 2010).

Duration: consider a case where one hears a single note sustained for an extended period of time; perhaps at the opera. At first, one apprehends the note, but after an interval of time has passed, another datum enters the experience: the perception of the note as lasting for a sustained interval. The experience gives rise to the feeling that the note has lasted some time: this feeling forms an additional part of the phenomenal character of the experience, over and above the impression of hearing the note. If this duration was merely inferred from the perception of the note for more than an instant of time, the experience of the note would not have the distinctive phenomenal character that it has; we would get no extra phenomenal notion of the duration at all given in immediate experience. However, it certainly seems as if we do; we seem to hear something about the note's temporal extent, over and above the note itself. The phenomenal character captures the duration of the note (Kelly 2005).

Motion: consider a car moving down the street. The car is located at various spatial locations over the course of the short part of its journey that is perceptible to you; let's say that the car is located at position X, then Y, and then Z. It is certainly true that we perceive the car as being located at each of these three intervals, one after the other; however, it isn't true that this is all that we perceive, for we also experience the car as being in motion; we perceive the motion of the car itself. If the motion of the car was merely inferred from the perception of the car at different spatial locations, then it is doubtful that the experience would have the phenomenology it in fact does have. We seem to perceive the motion of the car as being something over-and-above its changing of location over time. The phenomenal character of the experience captures the motion

of the car itself. (Broad's example of seeing the second hand of a clock moving vs seeing the hour hand of a clock moving essentially makes the same point (Broad 1923)).

If 'temporal properties' are characterized by being necessarily instantiated over time, then succession, duration and motion are all alike in this respect. However, motion can be distinguished in the sense that as a property it takes as its object a physical object. In contrast, succession and duration can only be predicated of events in time. We might say in light of this that motion is only a temporal property in a derivative sense; Phillips (2009) places change (and therefore motion) in the category of properties that whilst not 'strictly' temporal in the sense of pertaining to events bear a logical connection to time, but nevertheless uses the phrase 'temporal property' in a loose sense as to include these properties (Phillips 2009, p.3).

However, in this essay, as in the literature, the temporal properties I will be most concerned with are succession and duration; the motion case above is merely used as it is a vivid way of characterizing the claim that temporal properties feature in our immediate experience. That the focus will be chiefly on duration and succession here is justified by the fact that these are two temporal properties which naturally seem to be had by experiences themselves; or at least, experiences seem to be the sorts of things on a naive conception that could possess duration, and succeed each other in time. It is worth considering which properties are plausibly predicated of experience itself, and which are not, in order to explain my restriction of 'temporal properties' to 'duration and succession' in this essay. Without committing to a view about the precise nature of experience, I assume here that we can sensibly distinguish between the sorts of properties that might be truly predicated of entities like experiences, and those that should not be. For example, it is natural to talk of 'an unpleasant experience', or 'a joyful experience'; however, in making such predications, we seek to highlight something of how our experiences *seem* to us, but we would not literally accept that unpleasantness and joyfulness are the sort of properties that could sensibly be applied to experiences *qua* (partly) physical events in time. In the case of temporal properties,

however, when we predicate properties of duration and succession of experiences, we mean to make the predications more literally; stating that ‘my experiences are successive’ is not merely a reflection of how our experiences seem to us, but is to predicate experiences *qua* events in time with temporal aspects. The key contrast I wish to make in claiming that experiences themselves naturally seem to have temporal properties is not between temporal properties and ‘seeming’ properties such as unpleasantness, but temporal properties and other properties instantiated by worldly entities, such as spatial and colour properties. Many events and objects can be truly predicated with spatial and colour properties, as well a temporal properties: for example, ‘a large explosion’, or ‘a red flash’. Explosions and flashes can be both large and red as well as long-lasting (i.e. have duration). However, the starting-point in my inquiry into temporal properties in this essay is the observation that it does not seem correct to predicate experiences with spatial or colour properties in this fashion, whilst it *does* seem natural to predicate experiences of temporal properties. In the ordinary way we conceive of our own experience, we much more naturally assent to propositions such as ‘I had a long-lasting taste sensation’, or ‘I heard the crash before I heard the bang’, than predications such as ‘I had a large experience’, or ‘I had a red experience’; the latter predications just seem misguided. It is intuitions of this form that have caused a large group of philosophers to take seriously the contrast in metaphysical terms, and formulate explicitly theories of temporal experience that posit a systematic relationship between the temporal properties of objects and the temporal properties of experiences (for example, Foster (1979), Dainton (2000), and Phillips (2009)). Whilst I will ultimately seek to reject a central principle of these theories in this essay, I take as my starting point the intuition that temporal properties are significantly different from spatial and colour properties when it comes to experience: experience itself appears to have a temporal aspect in a way that it does not appear to have a spatial or a coloured aspect.

Temporal properties, then, might plausibly be had of experiences themselves as well as their objects (these objects being, in the cases of succession and duration, events). We

might coherently talk of, for example, a succession of experiences, or an experience that has a duration of x seconds in objective time. This fact is not true of many other properties we would like to predicate of objects and events. Spatial or colour properties, as just seen, are arguably never predicable of experiences; experiences are just not the sort of things that could instantiate these properties, or at least this seems to be the case. Hence as we have seen, an experience of a large red car is in no cases itself red, or large. In the case of many temporal properties, however, the question as to the relationship between the temporal properties predicated of the objects we experience and the temporal properties predicated of those same experiences is one of the central questions in the philosophy of temporal perception, and it is in giving an account of this relation that the main approaches to the problem can be distinguished. As I argue, the problem of temporal experience is introduced explicitly upon consideration of one key similarity between temporal and spatial/colour perception, and one key difference. The intuition behind the above three cases is the intuition that we perceive temporal properties as immediately as we perceive spatial/temporal properties; as Phillips and others observe: *'we seem to be no less directly acquainted with the temporal structure of the world around us than with its spatial structure'* (Phillips 2010 p.177). Alongside this key similarity is the essential difference between temporal and spatial perception mentioned above: that experience *of* time necessarily takes place *in* time in a way that has no spatial analogue. The essential similarity plays a key motivation behind what Dainton has called the 'phenomenological constraint' on temporal experience: the assertion commonly made (including by myself above) that temporal properties are presented or represented in immediate experience; I formulate this constraint more precisely in a moment. However, the essential difference between temporal and spatial experience invites consideration of the temporal properties *of* experiences as entities existing in time, and the relationship between the temporal properties of the objects encountered in experience.

Constraints on temporal experience

Having introduced the starting point of my inquiry into the nature of temporal experience, I will now consider some plausible constraints on temporal experience, arguing that when such constraints are agreed upon, certain views in the literature are revealed to be implausible. The main purpose of this section is to set the common ground in the debate between two key views concerning temporal experience that I want to discuss; the common ground is that so-called ‘cinematic’ or ‘snapshot’ views of experience should be rejected because they cannot accommodate these plausible constraints. One essential consideration to be advanced here is perhaps one of the most prominent statements about temporal experience in the literature, sometimes referred to (and largely undisputed) as ‘James’ dictum.’ In his canonical statement of the notion of the specious present, James partly motivates the notion with the following claim: ‘*a succession of feelings is not, in and of itself, a feeling of succession*’ (James 1890, p. 628). Reading ‘feelings’ as equivalent to ‘experiences’, James’ point is that there is no entailment from succession as applied to experiences to succession as represented in experience. It is unclear whether James has in mind a lack of entailment from a succession of experiences to succession being represented *in those very same experiences*, or in experience generally (viz. the stream of consciousness). The point, widely affirmed including by Dainton (2010a), Phillips (2009) and others, is that a mere succession of experiences as events in objective time is not enough by itself to produce a representation or presentation of the property of succession in experience; we need an additional fact about experience to ground such a presentation. Consider three experiences; hearing a tone A at t1 (at breakfast), B at t2 (at lunch), and C# at t3 (at dinner). These experiences stand in temporal succession; A precedes B in time, and B precedes C#; and furthermore, on reflection we might even say that at t3 we perceive *that* it is the case that A, B and C# stand in succession. However, it is not true to say that we perceive A, B and C# *as* succeeding each other; their succession itself is not (re)presented in experience at t3, since A, B and C# are experienced as too far apart in time. This fairly straightforward failure of entailment between properties in 1 and

properties in 3 in the case of succession does not, however, rule out that there *could* be succession experienced as a result of a succession of experiences, however far apart they happen to be in time, hence ‘in and of itself’. However, an extra ingredient is needed in order to ground the fact that succession itself is experienced, as immediately as the tones A, B and C#.

James’ dictum is a claim that most parties to the various debates about temporal experience are happy to accept. Even so-called ‘anti-realists’ about temporal perception such as Dennett (1991) affirm this principle; that it shows that in order to have the temporal phenomenology we generally think we have requires something other than a series of successive experiences shows, for the anti-realist, that we simply don’t have the temporal phenomenology we generally think we have (Dennett just denies that we have the temporal phenomenology we think we have). For these reasons and others, most have preferred to take James’ dictum to demonstrate that we need to add something to the ontology of experiences in order to explain our temporal phenomenology, rather than subtract something from the phenomenology itself. Dennett’s view, which advocates subtracting from the phenomenology, I consider later on in this essay.

A further interesting question that follows from this is whether the converse entailment holds; that is, whether an experience of succession (i.e. succession (re)presented in experience) entails a corresponding succession of experiences. Certainly in the case of hearing A, B and C# in a short temporal interval, it seems very hard initially to conceive of experiencing a succession of tones A, B and C# without also having an experience of A, then an experience of B, then an experience of C#: a succession of experiences. However, the answer is far from straightforward, and depends on more fundamental questions concerning the nature of experiences and how they are individuated, and how we conceive of an experience in the first place. If experiences are individuated simply in terms of the objects (re)presented, then the entailment under consideration would surely follow; for any tone A experienced, there would be a corresponding experience E(A),

the experience of that object. However, there are reasons to think that such a simple criterion for individuating experiences is false; we may have an experience where the represented object does not in fact exist, for example, or an ‘undirected’ experience, such as a feeling of general anxiety. Alternatively, we might, for example, take a view similar to Tye (2003), whose ‘one experience’ view doesn’t allow talk of numerous individual experiences, strictly speaking; however, one could on this view still hold that succession is represented in experience nonetheless. Here is a view on which an experience of succession does not entail a succession of experiences, for talk of experiences plural is simply incorrect.

Following on from James’ dictum, the first constraint on temporal experience that I shall consider is the view that (strictly) temporal properties are immediately experienced, suggested by the three cases above. Following Dainton (2010a), I term it ‘the phenomenological constraint’ on temporal experience:

The phenomenological constraint (PC): *temporal properties such as succession and duration are immediately experienced.*

This formulation of the constraint, however, makes use of an additional important notion made both by Dainton and others (though rarely elucidated in the temporal experience literature); the notion of ‘immediate’ or ‘direct’ perception, a loaded concept. Why is this addition required in setting up the phenomenological constraint as illustrated by the three examples of succession, duration and motion above? The answer can be made clear by a distinction between two different kinds of perception: perceiving *that* x, and simply perceiving x. To take an example from Dretske (1995), a person might see *that* her car’s petrol tank is empty without actually perceiving the emptiness itself, as a property of the tank. This is the very distinction that the force of the above three cases is grounded in, and the examples make explicit this distinction. Nobody, presumably, would argue that we perceive *that* various events and objects stand in the temporal relation of succession; even the most hard-line anti-realist about temporal

perception would agree to this. It is because of this that it is necessary to invoke ‘direct’ perception to distinguish the second case from the first; direct perception in terms of ‘seeing/perceiving x’ as contrasted with ‘seeing/perceiving that x (is the case)’. The distinction is a straightforward one, but the task of specifying what exactly constitutes direct perception is far from straightforward; in what sense is the perception of temporal properties direct? The manner in which Dainton (2010a) sets up the phenomenological constraint takes as crucial the claim that the perception of temporal properties is of a similar nature as the perception of spatial and colour properties; that these properties are represented in experience in a similarly immediate manner. Phillips (2010) agrees, affirming the key similarity between temporal and spatial perception mentioned above: *‘we seem to be no less directly acquainted with the temporal structure of the world around us than with its spatial structure’* (p.177). If this is indeed the case (though there may be other ways in which one could set up the phenomenological constraint without relying on this analogy), the question of providing an account of ‘direct’ perception of temporal properties is significantly related to the question of providing an account of ‘direct’ perception of spatial and colour properties. The idea is that whatever grounds the fact that when I perceive a large red object in my visual field I directly perceive its magnitude and redness will be the same sort of thing that (at least partially) grounds my direct perception of temporal properties such as succession and duration.

Here is not the place to go into a detailed investigation of the various differing accounts of direct perception, but something must be said about the sense in which ‘direct’ should *not* be used here in relation to temporal properties. We can be neutral here on the question as to whether naive realism or ‘strong’ representationalism is true (Chalmers 2005); in the representationalist case, ‘veridical’ perception is enough to be consistent with the thesis that we perceive temporal properties directly (Crane 2011). The key point to note, I take it, is that temporal properties are partially perceived in the same basic manner as spatial/colour properties, and represented in experience as such, whether the nature of that perception be disjunctivist, strong representationalist, or something else. Chuard (2012) explicitly makes a contrast between two interpretations

of direct perception starting with a familiar, ‘metaphysical’ conception; that of Jackson (1977), on which to directly perceive x is to perceive x ‘in virtue of nothing’ (Chuard 2012 p.5). On this view, a perception of x does not depend on any other further perceptual experience, or indeed any other kind of psychological state. What reason might we have to reject this reading of ‘direct’ perception? The answer is suggestive of a second constraint on temporal experience, which Chuard (2012) terms the ‘relational constraint’:

The relational constraint (RC): for any relation R , a subject S can perceive R between x and y only if S perceives both its relata x and y [in a single experience] (Chuard 2012 p.3).

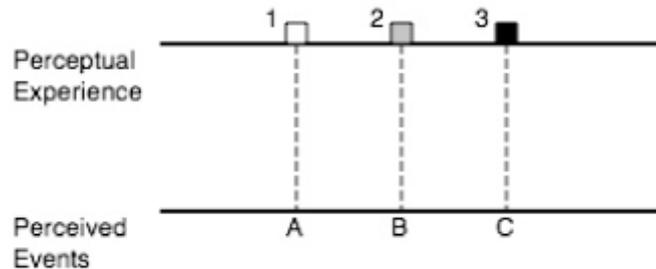
The relations at issue in the topic at hand are relations between non-simultaneous events, of which succession is the most prominent example. The relational constraint lies at the very heart of the problem of temporal experience; that it is true explains, along with **PC**, exactly why a ‘cinematic’ or ‘snapshot’ view of perception, which has it that both experiences and their contents as entirely momentary, is false. In his formulation of **RC**, Chuard curiously omits ‘in a single experience’, which results in a principle that is too weak. With this omission, perceiving succession merely requires both events in succession to be perceived at some point; however, it is doubtful whether we ‘perceive’ the succession at all between A and B if A takes place in 2013 and B in 2014, for example; certainly not if **PC** is correct. To make the principle consistent with Chuard’s own examples concerning colour/spatial perception (*‘it’s plausible that I wouldn’t be able to see the difference between the shade of red on the left and the darker one on the right if I couldn’t see both shades’* p.3), this addition is necessary. With reference to the above example of perceiving three tones A , B and $C\#$, the relational constraint more generally dictates that we cannot perceive the succession of tones A , B and $C\#$ without experiencing these tones themselves; that we perceive the succession is (partly) in virtue of experiencing the tones themselves. This is entirely plausible, and indeed it is very hard to see how ‘I perceive B as succeeding A ’ could be true if I do not

also perceive events A and B. However, if this is the case, then Jackson's 'metaphysical' reading of direct perception cannot be the correct reading for the above formulation of the phenomenal constraint on temporal experience; for we want to hold that the succession is directly perceived despite the fact that, according to the relational constraint, it is perceived in virtue of individual perceptions of the events in succession. For this reason, we should reject Jackson's reading of 'direct' as 'perceived in virtue of nothing' with regard to the phenomenological constraint on temporal experience. However, Chuard's proposed alternative, direct perception as characterized as being '*experienced by means of sensory perception exclusively, devoid of any meddling from any other kind of conscious state.*' (Chuard 2012 p.5) is also extremely problematic. It is highly questionable whether we possess a 'time-sense' in the Jamesian sense (James 1890 p.611); that is, something that James calls a 'myopic organ', which functions in a manner similar to the eye, for example. This simply seems like an unrealistic proposal; if the relational constraint is correct, and even if perception of temporal properties depends on perceiving events in time by means of the senses, it does not follow by the relational constraint that the perception of the succession of the events itself is sensory; such a proposal seems inherently unlikely, is not required for the sort of 'immediate' notion of direct perception we should have in mind, brought out by the contrast between perceiving that x, and perceiving x, or *as* x. Whilst the manner in which we enjoy direct perception of temporal properties is similar with how we perceive spatial and colour properties in terms of its *immediacy* in experience, this does not entail that both are perceived through exclusively sensory means. The crucial claim is that both the events in succession *and* their succession itself are represented in experience in a similar manner, though not necessarily through identical means. It is a further question as to the nature of these means.

The cinematic view

If the phenomenological and relational constraints as formulated above are correct, matters of direct perception set aside, we have a clear reason why a 'snapshot' view of

temporal experience cannot succeed in providing an adequate account of temporal experience. To see how **RC** and **PC** form a strong case against the snapshot view, it is worth setting it out as follows (from Grush 2007):



There are a number of ways this view of experience might be articulated, but all such views take as essential that both the contents of experience and the experiences themselves are momentary (or very short), and are correlated by the formation of static perceptual snapshots at different times. A dense succession of such momentary experiences is taken to ground the continuity of experience over time. On the snapshot view, our stream of consciousness exists merely in virtue of having different momentary experiences at different times; a series of static, momentary experiences with momentary contents all ‘run together’ in experience.

With this view thus described, we are now in a position to see how the two constraints on temporal experience that I have described **PC** and **RC** fit together, and in doing so provide a good argument against the snapshot view; in fact, the rejection of the snapshot view on these grounds directly leads to what has been described as ‘the problem’ of temporal experience. On the snapshot view, both experiences themselves and their contents are momentary. **PC**, however, requires that temporal properties, properties that are necessarily instantiated *over* time, are experienced directly, as immediately (re)presented in our experience as more basic sensory perceptions. However, in the case of succession, **RC** requires that if succession is to be perceived at all, its relata must also be perceived. If the relational constraint as described above is correct, no experience on the snapshot view will be able to provide perceptual contact with the various relata of the succession relation; there is no cross-temporal perceptual access which is required

(by **RC**) for such an experience. This fact directly leads to the above-mentioned James' dictum concerning succession; on the snapshot view, all we are provided with is a succession of experiences, which is not sufficient for an experience of succession, which follows from **PC** and **RC**. Since events in succession need necessarily to be spread across time (otherwise they would be simultaneous), and since succession is perceived partly in virtue of perceiving its relata (according to the relational constraint), then our perception of succession needs necessarily to be spread across time, at least in some sense: *temporal extension* of some aspect of our experience is required. That is, if both **PC** and **RC** are to be satisfied, some aspect of experience needs to be extended through time in order to ground the direct perception of relational temporal properties such as succession. If the argument above is correct, then the prospects for this view of temporal experience are dim. However, what are the prospects for a snapshot view that simply denies the truth of the phenomenological constraint (**PC**), and thus avoids the above argument against the experience of temporal properties such as succession? Unsurprisingly, the majority of theorists who might be thought of as subscribing to a snapshot view take this second option, since to combine such a view with **PC** appears unworkable, given the aforementioned considerations. I take the account of temporal experience set out by Dennett (1991) as the paradigm example of a snapshot view that denies **PC**. In discussion of this radical and revisionary view of experience, which holds that the ways we think and talk are poor guides to experiential phenomenal reality, I introduce a final constraint on temporal experience:

***The diachronic unity constraint (UC):** individual experiences stand in a certain relation/relations such that they are unified in a single, continuous stream of consciousness.*

As with the two previous constraints mentioned, I do not take this principle to be particularly controversial; in fact, it is a plausible and widely held claim about the character of our conscious experience. When we attend to experience, our consciousness possess a seamless, flowing quality that persists as long as we are

conscious, and characterizes our perception of the world; hence the widely used ‘stream’ metaphor. However, Dennett’s theory holds to a radically fragmented picture of consciousness, on which the dense series of momentary snapshot experiences is not supplemented with any further experiential (phenomenal) feature; a view which, I argue, leads to a rejection of UC above, despite UC being a basic desiderata of an adequate theory of consciousness. Dennett is keen to avoid what he terms a ‘content-vehicle’ confusion with respect to experience; we must make a distinction between the timings of the represented temporal content, and the timings of the neural realizers which give rise to the same content (Hurley 1998). This is an important suggestion, but in Dennett’s case it is motivated by a sharp disconnect between what we take to be the nature of our own experiences, and the true nature of those experiences. It is not the case, on Dennett’s anti-realist view, that we enjoy in immediate experience a unified, continuous stream of consciousness; rather that our brain simply receives disparate sensory inputs, and by means of ‘perceptual interpolation’ fills in the gaps in conscious experience, based on our beliefs and subconscious expectations about the apparent consistency of our visual field over time. Dennett first takes examples from ‘blind spot’ cases in the visual fields, cases where neuroscientific evidence is taken to suggest that there is in fact a fuzzy expanse in the field of vision of a subject despite the subject’s tendency to think and act as if there isn’t (Dennett 1991, p.331). Dennett then straightforwardly takes this case of spatial/visual perception and applies it directly to perception of temporal phenomena. Our consciousness over time is thus taken to be as fragmented and hole-ridden as our consciousness at a moment; both synchronic (at a time) and diachronic (over time) unity are sacrificed in a view that posits a fundamental disconnect between how our experience seems to us on introspection, how we form beliefs about it and report on its contents, and how our stream of consciousness *actually* is. The ‘multiple drafts’ model of experience that is developed in light of these considerations posits ‘*a parallel stream of conflicting and constantly revised contents*’ (Dennett and Kinsbourne 1992, p.1): there is no single stream of consciousness in which experiences are diachronically unified in this view; and the two parallel ‘streams’ of content posited instead hardly deserve the use of the metaphor,

since there is nothing continuous about them. Dennett's view is characteristic of a fully 'bottom-up' approach to consciousness that puts no evidential value in the deliverances of our own conscious introspection; but it is worth pointing out that in order to hold a 'bottom-up' approach to consciousness which gives weight to neuroscientific findings but also respects the phenomenology, we need not abolish the notion of evidence from phenomenology and introspection; for example, Lee (2014a, 2014c) gives such an approach that does not rely upon rejecting many plausible phenomenological considerations. Though it raises interesting questions concerning the relationship between phenomenology, introspection and belief, I argue that Dennett's account should be rejected partly on the basis that it denies **PC**, but chiefly because it leads to a rejection of **UC** as well, which I take to be an even more plausible constraint on our temporal experience.

I have argued that it is through considering three plausible constraints on temporal experience, **PC**, **RC** and **UC**, we must reject the basic 'snapshot' view of experience, and best articulate the basic problem of temporal experience. However, the methodological basis of this introductory chapter of this essay comes from the thought that the basic snapshot view of experience, whilst probably false, is to some extent intuitively plausible. At least, I find it to some extent intuitively plausible. The view plays an interesting role in the literature around the various problems of temporal experience; many theorists affirm it as the 'standard' view (e.g. Grush 2007 p.5) or as good basic, intuitive picture (Kelly 2005 p.1), after which it is usually quickly dismissed. In this sense the snapshot view plays a dialectical role rather than a role as a view to be seriously considered: only Chuard (2012) differs in this regard. I think there is much to be said for the basic thought behind the view; that perceptual experience is simple, unextended and simply takes as its input perceptual snapshots of the world at each moment. Many treatments of the various problems of temporal experience take it as given that this view is false; I have at least tried to give the view the more comprehensive treatment that it probably deserves. A proponent of the snapshot view may 'bite the bullet' by denying both the **PC** and **UC** (for example, Dennett 1991); or

they may deny **RC** (for example, Chuard 2012). I think on balance, however, such attempts involve denying other more deep-seated intuitions about the nature of experience. Following this discussion, I take the best exposition of the problem of temporal experience to be formulated in light of these three constraints on temporal experience. The question that naturally arises is the following: how can the basic snapshot ontology of experience be supplemented in order to satisfy the phenomenal constraint, given that the basic snapshot picture has been ruled out by considerations involving the relational constraint? Specifically, what sort of relationship must there be between the temporal structure of experience and the temporal structure of its objects in order to satisfy the above constraints? The point takes us back to James' dictum: '*A succession of feelings, in and of itself, is not a feeling of succession. And since, to our successive feelings, a feeling of their own succession is added, that must be treated as an additional fact requiring its own special elucidation*' (James 1890 p.628). In addressing the above concerns, the main conclusion of this thesis will be that we have good reason to doubt one prominent way of answering the above question; the view that experience 'mirrors' the temporal structure of its objects, or 'inherits' their temporal structure. In the next chapter, I elucidate upon this mirroring/inheritance view and the terminology involved, and then in later chapters give arguments to the conclusion that it should be rejected.

2. Mirroring and inheritance

Dainton's debate

We have seen the snapshot view to be false: in its anti-realist form, it must deny **PC** and **UC**, whilst no realist formulation of the snapshot view can plausibly be given without denying **RC**. The rejection of the snapshot view prompts the consideration and comparison of two temporal structures (sets of temporal properties and relations): the temporal structure of experience, and the temporal structure of its objects (which are, for the temporal properties we are interested in, events). When we reject the snapshot view, a view which seems to inevitably lead to affirming that neither experience nor its objects have any discernible temporal structure, the question arises as to how the two structures are related. James' dictum provides an important insight into this question: it is not sufficient for an experience to represent a temporal property that the experience or a set of experiences instantiate that temporal property. More must be said. I will now introduce the key views in the debate I wish to focus on, concerning the relation the two temporal structures in question bear to one another.

The debate concerning experience and its temporal properties, and their relation to the temporal properties represented in experience, has been dominated recently by terminology introduced by Barry Dainton (2000, 2010a). Dainton's primary concern in much of his work is how to give a plausible interpretation to what has been referred to (most notably in James 1890, citing 'E.R. Clay', p.609) as the 'specious present'. Motivated by the need to give a satisfactory account of this notion, Dainton formulates his taxonomy of views concerning experience as providing different interpretations of the basic specious present idea. Dainton's main distinction is between 'extensionalist' and 'retentionalist' families of views, but though I will reference these views in the course of my critique, the proposed distinction between extensionalism and retentionalism is not the distinction that I take to lie at the heart of the problem I am concerned with. As I detailed above, the question on the table is not simply the question

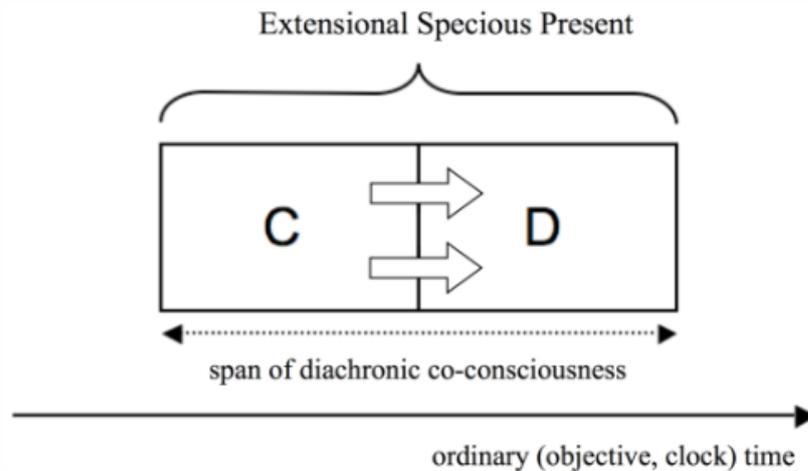
as to the particular ontology of experience (e.g. extended or momentary), nor the question of how temporal properties are (re)presented in experience, though both these further issues are clearly relevant. The key question I wish to discuss in this essay concerns the relationship between the temporal properties represented in experience, and the temporal properties of experiences themselves. Dainton's characterization of the debate about temporal consciousness does not make this concern prior; in his account, the question as to this special relationship follows from what he takes to be a more fundamental concern, which is whether experiences are extended in time or not. I present the issues in reverse: we should decide on the nature of the relationship between the two temporal structures before filling in the details as to the precise nature of each structure, for no reason other than the source of initial puzzlement was this very relationship. Therefore, I will only briefly outline Dainton's 'extensionalism' and 'retentionalism' here, before outlining what I take to be the more relevant debate, on the premise that the distinctions Dainton begins with do not cut to the heart of the issue that we should be most interested in here.

Dainton's key concern, as is ours, is how **PC** can be satisfied; specifically, what sort of nature experience must have in order that **PC** is satisfied. Having rejected as we have done the basic 'snapshot' picture of experience, though in less explicit terms, Dainton goes on to make a distinction between what he takes to be the two main competing views of temporal experience: 'extensionalism' and 'retentionalism' (Dainton 2010a). Dainton begins with James' notion of the specious present, and cashes out these two competing views as rival interpretations of the specious present. For James, the 'specious present' refers to the present moment as it is actually experienced, as contrasted to the 'strict' or 'mathematical' present, which has no duration at all (James 1890). For James, the present moment as experienced cannot be momentary: it is *'no knife-edge, but a saddle-back, with a certain breadth of its own on which we sit perched, and from which we look in two directions in time'* (James 1890 p.609). For James, the experienced present must apprehend a duration, or interval, of time owing to **PC**; we could not perceive temporal phenomena at all if all we were presented in

experience with was a moment, because temporal phenomena are themselves necessarily extended in time. Though the actual, ‘mathematical’ present moment may be strictly momentary, the present as given in experience is temporally extended. Thus, in motivating the idea of a specious present, James implicitly rejects the ‘snapshot’ family of views which reject that experience gives us apprehension of a duration; his rejection can be seen as closely related to the rejection of snapshot theories presented in the previous chapter. In his highly metaphorical description of the specious present, James explicitly states that the specious present sits either side of the strict, mathematical present; that is, in a single experience, we are able to apprehend a duration which includes objects and events that are located both before and after the actual present. Later in his development of the idea in the *Theory of Knowledge* manuscript, Russell modifies the notion of the specious present to incorporate only ‘backwards’ awareness: for Russell, we might say that the specious present sits entirely in front of the mathematical present in experienced time, thus only allowing access to a short duration of the recent past (Russell 1913). Whichever way we choose to understand these specifics, however, the key claim in the notion of the specious present is that in a single experience, we are presented with (or represent) an interval of time. For Russell, we can only experience temporal properties in a manner consistent with **PC** if we experience time within a specious present. More specifically, the specious present allows for multiple entities to be experienced together, or *all at once*, i.e. in a single experience. So for example, when I experience the succession of notes A and B, I am able to apprehend both A and B, and their succession, in a single experience; the two notes are contained within a single specious present.

The specious present is primarily a thesis about how the world is presented to us in a single experience; it is not a thesis concerning the nature of experience itself. Dainton introduces the debate between ‘extensionalists’ and ‘retentionalists’ as a way of interpreting the basic specious present idea in the case of experience itself; these two views outline very different conceptions of how experience itself must be if the notion of the specious present is upheld. ‘Extensionalists’ hold that individual experiences

themselves are extended in time, and commonly across the same interval as their specious presents. Thus, if a typical specious present is 2 seconds long, then the corresponding experience will also be 2 seconds long; the length of a single experience matches the length of the specious present associated with it. The general extensionalist picture can be represented in this manner:

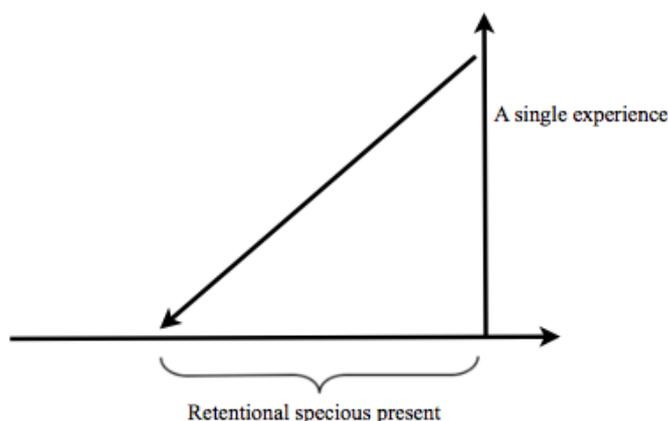


(From Dainton 2010a, section 5.1)

As can be seen from the diagram, the temporal extent of a single, temporally extended experience matches the temporal extent of its corresponding specious present. Dainton holds that experiences are unified over time (diachronically) in virtue of being related by the relation of co-consciousness: we needn't enter into an analysis of Dainton's chosen unity relation here, so can set aside this part of the view for now. Apart from the fact that the experience *itself* is extended on this picture as well as its content, the diagram illustrates another crucial aspect of the extensionalist picture: the claim that experiences decompose into further experiences. Thus, the experience of C (whatever event or object it is) and the experience of D are unified by the relation of co-consciousness to form a further experience, CD; and it is the length of CD that is the same as the length of the specious present. The relation that Dainton takes to hold between C, D and CD is a mereological relation; that is to say, the parts of an experience (which are also experiences) compose a distinct experience in a similar manner in which, for example, a wooden seat and four wooden legs might compose a

wooden chair. Again, since my explanation of the view is brief and chiefly methodological, we needn't at this point get into discussion about this mereological view of experience. The key claims of the extensionalist thesis are 1) that experiences are things that have a temporal extension which matches that of their specious presents and 2) experiences decompose into parts, which are also experiences. On the extensionalist picture, then, there is a real sense in which the temporal structure of an experience and the temporal structure of its objects are intimately related.

The contrasting view to extensionalism Dainton terms 'retentionalism', which I admit I find a slightly perverse labelling. If 'extensionalism' is named due to its claim that experiences are extended and built up of other experiences as parts, then why not label the alternative view 'atomism', denoting that experiences need not be extended, and do not decompose into further experiences (e.g. Lee 2014a, 2014b)? Instead, Dainton uses 'retentionalism', in reference to the fact that this kind of view is often supplemented with cognitive faculties such as 'retention' or 'memory' in order to make it coherent. To understand how this is so, the retentionalist these can be represented as follows:



On the retentional picture, we can see that the experience (the vertical arrow) is not necessarily extended, but point-like, or momentary. Nevertheless, the content of the experience, or what is presented in the experience, has temporal extension; the single experience still represents an interval of time, and so can be said to honour the notion of the specious present. This is possible because of the diagonal arrow, a 'backwards'

representation or retention. The temporal interval necessary for the perception of temporal properties such as duration and succession is made possible by the contribution to the content of experience made from such retentions. Retentionalism differs from extensionalism in two key aspects: firstly, the temporal structure of the experience itself (if it has one) and the temporal structure of the objects apprehended (within the specious present) come apart in terms of duration; only the specious present is had over an interval. And secondly, experiences on this picture can be momentary, and so do not necessarily decompose into parts that are also experiences. The view is termed ‘retentionalism’ in Dainton’s terminology because in order for a momentary experience to apprehend a temporal interval, other cognitive faculties need to be employed: in this case, retention, or memory, a notion taken from Husserl. The retentionalist’s momentary experience is able to apprehend an interval of time because that interval is partly constituted by recent memories; the cognitive faculty of memory plays a vital role in the representation of a temporal interval. On retentionalism, the temporal structure of an experience and the temporal structure of its objects can significantly diverge.

Again, since this exposition of Dainton’s view of the debate is brief, there is no space to fully explore these issues; I introduce carving up the debate about temporal experience in terms of ‘extensionalism’ and ‘retentionalism’ in order to contrast it with an alternative, and in my view better, way of carving up the debate. I don’t believe that Dainton’s contrast between extended (extensionalist) and momentary (retentionalist) experiences makes the central distinction in the inquiry that I wish to pursue: namely, the question of the relationship between the temporal structure of experience and the temporal structure of its objects. For example, one could hold a view on which an experience is momentary with an extended content (Dainton’s ‘retentionalism’) without holding that retention is the cognitive faculty making this picture possible; furthermore, one could hold to a view on which an experience is *not* momentary, but just not extended to the same extent as its content (e.g. Lee 2014a). As we have seen, something about the relation between our two temporal structures can be drawn from Dainton’s

two positions; extensionalism would most naturally include the thesis that the temporal structures match, whilst retentionalism would most naturally come together with the view that the temporal structures can and do diverge. However, this is not the whole story: I would like to instead investigate the notion of ‘temporal matching’ or ‘mirroring’ further, what it involves, and what sort of conception of experience arises from it. To this end, I will now outline the debate in different terms: there is disagreement in the literature concerning the extent to which a ‘mirroring’ constraint on temporal experience is upheld. I will also consider the notion of ‘inheritance’, which is found in Phillips (2009) and Soteriou (2013), and which is taken to be a development of the mirroring thesis. After setting up the debate as framed around the mirroring thesis, I present some challenges to the mirroring these in chapters 3 and 4. In the concluding chapter 5, I return to a specific view of experience that a denial of mirroring might suggest (something akin to Dainton’s retentionalism, albeit without specific reliance on the notion of retention), and consider some of its weaknesses.

Mirroring

A temporal structure is a collection of entities ordered by temporal relations, and possessing temporal properties. We experience the world as having a temporal structure, in the sense that we are aware of successions of events which have duration, and objects that change over time. However, it is also intuitively plausible that our experiences themselves have their own temporal structure; experiences succeed one another in time, change, and have duration as well. In this essay, I am concerned with these two temporal structures and the relation between them: the temporal order and durations of objects experienced, and the temporal order and durations of the experiences themselves. We generally take both experiences and their objects to have temporal structures; and so one view open to us is to hold that these temporal structures mirror, or match, each other. The basic ‘mirroring’ principle can thus be articulated as follows:

Mirroring - For a temporal property T, if the objects of an experience [for our purposes, events] instantiate T, then the experience itself instantiates T.

Lee (2014a) describes mirroring in less formal terms, as the view that *'the way an experience of a melody [for example] unfolds over time mirrors the way the melody itself unfolds over time'* (p.1). What this means in practice is that if the object of an experience has a given temporal property, then the experience itself has the same temporal property. Of course, the mirroring principle as presented here is too imprecise to analyze: which temporal properties satisfy mirroring, and how fine-grained we are to take the principle is left open. As mentioned earlier, the mirroring principle cannot be true of all properties loosely characterized as 'temporal properties', since this would include motion, and experiences are not plausibly thought to move around like physical objects. The mirroring principle can only be considered properly if it is decided which temporal properties satisfy it. As stated previously, I will take the temporal properties in question to be duration and succession (ordering); I will assume that these are two viable candidates that could satisfy the principle. Typically, mirroring theorists (for example, Phillips (2010, 2014a)) hold that at least these two key temporal properties satisfy the mirroring principle.

With this in mind, we can consider three different readings of the principle. Lee articulates these three readings as follows:

Metrical Mirroring: the ordering *and* duration relations between the temporal parts of a process-experience match those of the apparent perceived scene.

Topological Mirroring: only the ordering relations between the temporal parts of a process-experience match those of the apparent perceived scene.

Structural Mirroring: only distinct temporal stages of the perceived scene are presented by distinct temporal stages of experience.

(Lee 2014a, p.9)

In his treatment of mirroring, Lee reads the mirroring constraint as only being an option for what he calls ‘process’ views of experience: views that hold that experiences are extended processes that unfold alongside their objects. This specific view needn’t concern us at this point, though we can note that Dainton’s ‘extensionalist’ would most naturally be drawn to such a view. We might read these three interpretations of the mirroring principle as varying from strong to weak respectively. *Metrical* mirroring states that both the ordering and duration relations of an experience match those of the perceived scene¹. For example, if I have an experience of three notes in quick succession, A, B and C#, then that experience will have three parts that correspond to A, B and C#, which occur in that order. Furthermore, if the duration of A is 0.5 seconds (say), then the duration of the experience of A will also last 0.5 seconds. As Lee puts it, ‘an experience of a 1 second gap between two sounds is mirrored by a 1 second gap between the experiences of the sounds themselves’ (Lee 2014a p.9). This is the strongest mirroring constraint one might endorse with regard to the properties of duration and succession: both temporal properties are mirrored by the experience itself. In the case of *topological* mirroring, only the temporal property of succession is mirrored: the ordering of the respective experiences is fixed by the ordering of A, B and C#, but the individual experiences of these notes need not have the same duration as the notes themselves. And in the case of *structural* mirroring, the weakest reading of the principle, it is only the general temporal structure of the perceived scene that is mirrored with regard to the experiences themselves; on this reading, the experience of A, B and C# might have three parts corresponding to the three notes, but might not necessarily occur in this order.

It is easy to see how the mirroring principle might be attractive to an ‘extensionalist’ in Dainton’s terminology. Recall that the extensionalist’s two main commitments were to experience being extended, and experiences decomposing into parts that are also experiences. The strongest interpretation of mirroring, *metrical* mirroring, seems to entail these two commitments: for an experience to mirror the successive structure of A,

¹ Phillips uses ‘*apparent* object’, in order to account for cases of hallucination and illusion where the experience may be of something other than a genuine physical object or event.

B and C#, it must have parts that correspond to these notes that are also experiences, and for an experience to mirror the duration of a note, it must itself have duration, i.e. be extended in time. It is less clear that *topological* mirroring or *structural* mirroring entail the extensionalist thesis explicitly, but Lee holds that each of the three mirroring constraints imply a ‘process’ view of experience; and since the most natural way to hold experience to be extended is to hold that experiences are processes, extended in time, we might agree with Lee on this point. If one subscribes to a mirroring constraint on temporal experience, whatever its strength, then the most natural view to take regarding experience itself is, in Dainton’s terminology, ‘extensionalism.’ Dainton himself subscribes to this position in Dainton (2000, 2010a). In the course of this essay, I will advocate the denial of at least the metrical and topological mirroring principles previously outlined, in the case of temporal experience.

Mirroring, resemblance and inheritance

If it is accepted that any of the mirroring constraints above, or interpretations of the mirroring principle, are genuine constraints on temporal experience, one question that naturally arises is the question of why should such a mirroring constraint hold? It would surely not be a mere coincidence that the two temporal structures that we are considering should be systematically related in this manner with regard to duration and succession: what explains the relation? Does one temporal structure depend for its nature on the other, and if so, which way does the dependence run, and what is the nature of the dependence relation? Lee considers one possible explanation of mirroring (though does not endorse it): what he calls ‘representation by resemblance.’ If such a resemblance theory of experience is true, then ‘*experiences present certain kinds of temporal features partly by themselves having the very same features*’ (Lee 2014a p.10). On this view, experiences are able to present / represent certain temporal properties partly by having those very same temporal properties: the specific temporal properties that satisfy this theory (duration, succession, or both) depending on the particular mirroring constraint endorsed. In this proposed explanation of mirroring, the content of experiences depends at least in part on the nature of the experience itself: the temporal

structure of the experience is metaphysically prior to the temporal structure of its objects. An explanation is thus given of why a given experience has a given content: it is the temporal structure of the experience itself that explains the temporal structure of its content. As Lee states (p.11), the representation by resemblance view does not violate ‘James’ dictum’, mentioned earlier, which is that experience having a certain temporal property (i.e. succession) is not sufficient for that property to be represented in experience. Rather, succession of experience is necessary but non-sufficient for experience of succession; other constraints will bear on whether experience is able to (re)present succession, including (perhaps) the proximity of experiences in time, and whether such experiences are unified in the right way over time. On the resemblance picture, the temporal structure of experience is explanatorily prior to the temporal structure of the objects, and so an explanation is given for their mirroring: the temporal structure of the objects of experience matches the temporal structure of experience itself owing in part to this representation by resemblance.

An alternative picture of why mirroring might be the case is suggested by Soteriou in *The Mind’s Construction* (Soteriou 2013). Motivated by a relational ‘naive realist’ view of experience, on which the character of perception is (partly) constituted by the objects of the perceived scene, Soteriou sees the dependence relation hold in the opposite direction. On the relational view, which is often contrasted with ‘representationalist’ views, a special perceptual relation is present between a subject and his/her objects of perception in cases of true perceptual contact with the world. When we genuinely perceive an object or event, the object or event partly constitutes the experience itself. This is not the case in cases of hallucination, where we appear to be in genuine perceptual contact with an object or event, but in fact the object or event does not exist. In these cases, the perceptual relation does not hold, despite the fact that the scene might appear to the subject to be qualitatively identical. Cases of genuine perception are therefore metaphysically distinct from cases of mere hallucination in terms of the sort of experience enjoyed: genuine perceptual experiences and hallucinations are experiences of fundamentally different kinds. The representationalist disagrees, positing a difference between ‘veridical’ and ‘non-veridical’ perception; cases of genuine perception and

hallucination are fundamentally the same kind of experience, except that in cases of hallucination, the perception is non-veridical, or false. Soteriou subscribes to a ‘process’ view of experience which, as we have seen, is extensionalist according to Dainton; on this view, experiences are extended processes which unfold alongside their objects. Given the relational picture Soteriou also subscribes to, his explanation for the mirroring constraint is that the temporal structure of the experience depends on the temporal structure of its objects; the object (in this case, an event) perceived constitutes perception in such a way that the temporal structure of experience mirrors the temporal structure of the event perceived. Whilst Lee’s ‘representation by resemblance’ had the temporal structure of the events (in the veridical case) dependent on the temporal structure of the experience, Soteriou sees the dependence going the other way: the temporal structure of experience depends on the temporal structure of the events perceived, in the case of genuine perception. The *explanation* for why mirroring holds on Soteriou’s account is that, on the relational view, the objects of perception partly constitute experience itself, and this is the basis for the latter to gain its temporal structure from the former.

We might refer to the pictures Lee and Soteriou present with regard to the explanation of temporal mirroring as presenting ways in which one entity ‘inherits’ the properties of the other. In Lee’s case, the objects of experience inherit their temporal structure from experiences themselves, whilst in Soteriou’s case, experience inherits its temporal structure from its objects. The language of ‘inheritance’ when it comes to temporal properties is used most prominently by Phillips (2014a) in describing the relation between our two temporal structures. However, and as Lee (2014a) notes, Phillips uses the language of inheritance simply to advocate a mirroring constraint on temporal experience, and doesn’t intend his account to provide an *explanation* for why mirroring is true. Therefore, whilst we might interpret an ‘inheritance’ claim as suggestive of a metaphysical dependency relation, as in the case of Soteriou, in Phillips the term is simply a metaphor for the existence of a mirroring constraint, plus a certain thesis concerning order of explanation. Phillips’ presentation of inheritance is dominated by the claim that

‘...for any temporal property apparently presented in perceptual experience, experience itself has that same temporal property’ (Phillips 2014a, abstract);

in other words, a mirroring thesis. Rather than explicitly adding to the mirroring thesis a claim about the metaphysical dependence of one set of properties on the other, Phillips simply discusses the mirroring thesis in the language of inheritance, suggesting that inheritance in his account is simply a metaphor that denotes a general matching of temporal structures: *‘our stream of consciousness inherits the temporal structure of the events which are its contents’* (Phillips 2014a p.6). Since Phillips’ presentation of mirroring is the account which I will be paying most critical attention in this essay, it is worth getting clear on a couple of preliminary questions concerning Phillips’ conception of mirroring and inheritance, with reference to the quote from Phillips 2014a, above.

Firstly, it is a crucial question as to what kind of mirroring constraint Phillips has in mind when advocating mirroring by way of the metaphor of inheritance. Something stronger than merely *structural mirroring* is intended: if *‘any property apparently presented in perceptual experience’* is mirrored for experience itself, then this is a claim not simply about experience-phases, but about duration *and* ordering/succession, since as we have seen, these are both properties that can be said to feature in immediate experience. This observation alone is suggestive of the strongest mirroring constraint, *metrical mirroring*, on which both ordering and duration relations are mirrored in experience. However, when Phillips makes an argument for the mirroring/inheritance claim explicitly (Phillips 2010), his focus is solely on the temporal property of *succession*: a key premise in his argument (which will be the main focus of the next chapter) is that *‘we will always rationally judge an experience of succession to be itself successive in temporal structure...’* (Phillips 2010 p.183). It is unclear whether Phillips intends his argument to apply only to the temporal property of succession, or whether this is simply an example property used to demonstrate the conclusion of the argument, which concerns any temporal property found in experience. For the purposes of this essay, I take Phillips to be advocating a *metrical* mirroring constraint, and this is the mirroring constraint I will seek to cast doubt on later on. I take it as read that when

Phillips states that *'for any temporal property apparently presented in experience...'* the class of temporal properties is restricted; *motion*, for example, is a temporal property on one reading, but is never plausibly had by experiences themselves. Quibbles about the class of 'temporal properties' aside, I assume here that Phillips' claim should be restricted to duration and succession only, since these are properties which at least are candidates for being had by experience themselves, when we reflect on our own experience.

Secondly, we can question Phillips' formulation of the mirroring thesis above by noting the distinction between properties of objects/events and *apparent* properties of objects/events. Like Lee (2014a), Phillips formulates the notion of mirroring/inheritance as holding between apparent objects events and experiences, rather than the objects/events themselves. This is intended to allow for cases of hallucination, where what is presented in experience is (seemingly) qualitatively identical to a real object or event, but no such object or event exists. On a relationalist account such as Soteriou's where the inheritance relation is read in a 'metaphysical dependence' sense, and where the contrast between hallucination and genuine perception is crucial for the nature of the experience itself, the contrast is not necessary: since in genuine perception the real objects partly constitute the perceptual experience itself, mirroring will only hold in cases of genuine perception. For Phillips, introducing the 'apparent' clause allows neutrality on the perceptual question: representationalism vs a naive realist/relationalist account. There is certainly a question here of how the properties of an experience are supposed to depend on properties of apparent objects; since apparent properties depend on experiences themselves (to be 'apparent' is just to feature in experience), there seems to be something circular in claiming that the properties of experiences themselves depend on apparent properties or objects, which themselves depend on experiences. The point would be more acute if Phillips advocated inheritance as a thesis of metaphysical dependence: merely apparent properties do not seem to be the sort of entities on which the nature of experience could depend in this manner. However, as we have seen, we should think of inheritance in Phillips as used merely as a metaphor by which to

advocate a mirroring constraint. For the purposes of this essay, we may bracket worries about apparent vs real properties, and cases of hallucination; restricting the mirroring/inheritance thesis to veridical cases (cases of genuine perception, where properties of real objects are (re)presented) will be sufficient for my critique of mirroring.

Having set out these preliminaries, the main purpose of this essay is to cast doubt on the view that experience either *metrically* or *topologically* mirrors its objects in terms of temporal structure; to this end, I will in the next chapter discuss a major argument from Phillips for mirroring, which I will give reasons to reject. The key form of my argument in the next two chapters is as follows. Metrical mirroring entails that if I experience a note A then a note B, then my experience of A must precede my experience of B in time (owing to the inheritance of ordering relations between the objects of experience); and furthermore, if A is experienced as being of a certain duration, the experience of A itself must have the same duration. Topological mirroring entails only the first of these two claims, the claim concerning succession. However, there may be cases whose most natural interpretation suggests that A and then B can be experienced without the corresponding experiences possessing that order, or cases where A is experienced as having some duration, but the corresponding experience lacks this duration. If such cases can be substantiated, then this is a reason to doubt that either metrical or topological mirroring is the case with regard to temporal experience.

3. Mirroring and transparency

Introduction

In the previous chapter, I set out the closely related notions of ‘mirroring’ and ‘inheritance’ proposed by some theorists as constraints on temporal experience, as a means of accounting for some features of our experience of temporal properties. I observed that whilst what sometimes characterizes the ‘inheritance’ thesis is a claim about dependency (namely, that the temporal properties of an experience depend for their explanation in some way on the temporal properties (re)presented in that experience), all readings of the inheritance thesis entail that there is a ‘mirroring’ constraint on temporal experience. That is, if inheritance is true, the temporal structure of an experience must mirror or match the temporal structure of that which is (re)presented in experience, at least with regard to the key temporal properties of duration and succession. As stated previously, though inheritance is a dependency relation, and therefore asymmetric, it entails a matching of two temporal structures, and it is this matching I shall be focussing on in this chapter. That the two temporal structures under investigation match in this manner is the conclusion of an argument made by Phillips (2010) from observations concerning the so-called transparency of experience. In this chapter, I present a number of challenges to this argument, which if reasonable provide a reason to reject that such mirroring or matching is a constraint on our temporal experience; a prominent argument for the mirroring view can be rejected. If we reject the claim that experience temporally mirrors its objects in the temporal case, then this entails that the inheritance thesis is false, since inheritance entails a mirroring constraint. Phillips restricts his argument to the temporal property of succession; I take my rejection of this argument to provide good reason to doubt that experience mirrors its (apparent) objects with regard to succession. Having made this negative case, I go on in chapter 4 to present a further argument against the mirroring view, this time concerning the duration of experience.

Phillips' argument from transparency

I will begin by presenting the argument as Phillips presents it in Phillips (2010); I then present what I take to be a clearer reconstruction of the argument presented in Frischhut (2014). Phillips presents the argument as an argument against the view that we can experience temporal intervals at a moment, and furthermore that the argument demonstrates that this view is '*revealed to be impossible when we reflect on the nature of our experience*' (Phillips 2010 p.183). However, the argument is also an argument for the matching of two temporal structures, and so presents a challenge to any view that denies this claim, including views (such as the view I advocate later with Lee 2014a) that deny that experiences are momentary or durationless, but that nevertheless hold that the two temporal structures in question can diverge. The argument is intended by Phillips not to show that such divergence (i.e. a rejection of temporal mirroring) is not theoretically impossible or contradictory, but that it is not consistent with manifest facts revealed when we reflect on the nature of our experience.

- *'When one attends to the temporal structure of experience (that is: reflects upon its nature...it is rational to judge that one's experience is temporally determined....only by taking its temporal structure to mirror the apparent temporal structure of the world experienced, i.e by making a judgement concerning the apparent temporal structure of the world experienced, and then taking the experience to have that same temporal structure.'*

- *'Thus, we will always rationally judge an experience of succession [for example] to be itself successive in temporal structure as opposed to instantaneous.'*

- *'Experience cannot systematically seem some way to rational introspective reflection and yet be some other way.'*

- *Therefore, 'we cannot be systematically in error when we judge our experiences of succession to themselves be successive in temporal structure [for example] as opposed to instantaneous.'*

(Phillips 2010 p.183)

Here Phillips focusses solely on the temporal property of succession in this formulation of the argument, though I take it that it is at least a possibility that the reasoning applies to the property of duration as well. Whilst Phillips' formulation of the reasoning from the so-called transparency of experience to the conclusion that experience temporally mirrors its objects is clear enough, there are many notions appealed to in the argument that call out for further explanation and elucidation. To fully assess the argument, we require better notions of 'transparency', 'introspection', 'rational judgment' and 'seeming'; and as I will later argue, the argument can be challenged once we get clearer of how these concepts are being used and the role they play in the argument. Furthermore, Phillips' own presentation of the reasoning is casual in nature; his presentation does not fit the form of a tight, clear argument. I will therefore follow Frischhut (2014) in reconstructing the argument in a tighter form, in order to make clear the various claims on which each stage of the argument relies.

- 1) **Transparency:** when we attend to our experiences, we are attending to, at least in part, the objects of those experiences; we have indirect introspective access to the temporal structure of our experiences through perceiving the temporal structure of their objects.
- 2) **Rational judgment:** we will always rationally judge an experience of e.g. succession, to be successive in temporal structure as opposed to instantaneous. It is always rational to judge as such.
- 3) **Judging -> seeming:** If S judges that experience is a certain way based on rational introspection, then it seems to S that experience is that way.
- 4) Therefore, it always seems as if our own experiences mirror the temporal properties of their objects.

5) **Seems -> Is:** if it seems to S that their experience is a certain way based on rational introspection, then S's experience has that nature; seeming is a good guide to the nature of experience.

6) **Conclusion:** Therefore, experience has the same temporal structure as its objects.

(from Frischhut 2014)

The conclusion of the reconstructed argument as I have construed it is that some form of mirroring is the case, to say nothing further of the relationship between the two sets of temporal properties. In Phillips' original statement of the argument above, reference is made to the *determination* of the temporal structure of experience by the temporal structure of the world; however, the inheritance claim implicit here is not strictly necessary to establish the conclusion that Phillips wants, which is that an experience cannot be momentary whilst apprehending a temporal spread of content. In arguing for the conclusion that the two temporal structures must match, Phillips wishes to rule out the possibility of the divergence of the two temporal structures. Therefore, I will treat the above argument as an argument for mirroring, rather than inheritance; whether the claim concerning explanatory dependence investigated in the previous chapter holds in addition to the mirroring claim is not at stake here. None of the observations crucial to the argument rely on this additional specification of the direction of the dependence of one temporal structure on the other.

Analysis of Phillips' argument

Premise 1, 'transparency' consists of an observation made frequently in the philosophy of perception concerning the content of introspection. The central observation is that when we turn our mind's eye inward in introspection, towards experiences themselves, we encounter only the objects of experience: 'external' properties and relations rather than intrinsic properties of the experience itself. Metaphorically speaking, experience is 'transparent' because none of its properties are given in introspection: we see right

through it, and apprehend properties of the objects of experience instead. As Harman puts it:

“When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experiences ... Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree”. (Harman 1999, p.251).

Here, Harman construes transparency as evidenced by our inability to pick out any intrinsic features of experience when we introspect: when we describe a certain experience, for example of a red tomato, we will inevitably use words such as ‘red’ and ‘round’ to describe the experience despite the fact that these are properties not of the experience itself, but of the object of the experience. If there are intrinsic features of experience (such as qualia etc), we cannot apprehend them or learn about them via introspection. Martin describes the notion in a similar manner:

“...introspection of one’s perceptual experiences reveals only mind-independent objects, qualities and relations that one learns about through perception.” (Martin 2002 p. 378)

As with Harman, the passage from Martin suggests transparency as a largely negative thesis concerning what we are able to apprehend via the process of introspection: it is *only* mind-independent properties we can apprehend, since we are not able to specify any intrinsic features of the experience itself. Based on this construal, we might take transparency considerations to highlight an important constraint on our ability to come to know facts about our own experience; whilst we might think of our own experiences as immediately and intimately knowable, a moment’s reflection tells us that we cannot

pick out any of their intrinsic features by simply turning our attention towards them. What is revealed by this process are external properties and relations. Transparency construed in this way can be seen as a negative claim, since it entails that we cannot apprehend ‘internal’ properties (i.e. properties of our experience itself) that are distinct from the properties of the external object/event we are experiencing.

The manner in which Phillips presents the transparency claim, however, can be seen as a development of the conception of transparency advocated by the theorists quoted above. Whilst acknowledging that introspection acquaints us directly only with external properties and relations, Phillips supplements this thesis with a further claim that is crucial to his argument: that we can nevertheless have indirect access to some of the properties of our experiences (i.e. an experience’s temporal properties) through introspection, at least in the temporal case (Frischhut 2014). In explicitly introducing a contrast between direct and indirect perception, Phillips concurs with Harman and Martin that our introspection is restricted in a significant sense to external properties and relations of objects, but supplements the negative thesis with a further claim about experiential access. For Phillips, despite the fact that when we attempt to describe our experiences we find ourselves attending to / describing properties of the objects of those experiences, we nevertheless have a kind of *indirect* access to the temporal structure of our experiences through perceptually attending to the temporal structure of their objects. This represents a positive thesis, and a somewhat weaker interpretation of the original transparency observation. Whilst the manner in which Harman and Martin present transparency is largely negative, Phillips explicitly makes the claim that the act of introspection, or rational reflection on one’s own experience, does reveal something about experience; its temporal structure, albeit indirectly via apprehension of the temporal structure of the objects of the experience:

‘...when we set out to describe our experience itself, we find ourselves doing so, at least partly, by attending to its objects.’ (Phillips 2014a p. 132).

The positive formulation of the transparency claim is taken by Phillips to simply follow from reflection on our own experience, rather than from any theoretical considerations about perception. However, premise 1 is also implicitly committed to a ‘uniqueness’ claim concerning the perception of temporal properties; that the kind of indirect access to experience required for premise 1 only holds in the temporal case. Recall the well-worn observation made earlier that it doesn’t seem as if experience has a colour or spatial structure, but it does seem that experience has a real temporal structure. The transparency claim in premise 1 presupposes that temporal properties are unique in this regard, since in the spatial and colour cases, experience simply has no such properties to be indirectly apprehended on introspection (or so we might think). Phillips does not want to hold that experiences are spatially transparent, for example, since he takes it to be implausible that experience has any spatial properties at all that are indirectly acquainted with on introspection. Considered in isolation, this uniqueness claim is plausible, however it is important to consider it as it will play a key role in the criticism of some of the other premises of Phillips’ argument. Premise 2, concerning ‘rational judgement’, is taken by Phillips to simply follow from premise 1, given a couple of background assumptions concerning our ability to determine the temporal structure of our own experiences.

More interesting, and controversial, are the next steps in the argument, from the claim concerning rational judgment to the conclusion concerning not our judgment of experience, but the nature of experience itself. Having established via transparency that we can determine the temporal structure of our experiences indirectly by taking it to match or mirror the temporal structure of the objects of experience, and that we make rational judgments concerning experience on this basis, Phillips introduces the notion of ‘seeming’. For Phillips, judging that x (where x is a fact about our own experience) entails seeming that x ; that is, if S judges that x , then it seems to S that x . Premises 2 and 3 jointly entail that it always seems to us as if the previously mentioned mirroring constraint on temporal experience holds; that is, we can determine the temporal structure of experience only by taking it to mirror the temporal structure of the objects

of experience. Premise 5 states that we are not systematically deceived when it comes to experience; if experience seems a certain way to us, then it really is this way; introspection is 'infallible' when it comes to the nature of experience (Lee 2014c). The conclusion, that the temporal structure of experience and the temporal structure of its objects must match, follows.

Different views on access

I will now consider ways the argument might be challenged, and rejected. Lee (2014c) notably challenges the 1st, and by extension 2nd, premise. In his critique, Lee formulates two claims contained within Phillips' premise 1: the 'transparency' premise that introspection only reveals external properties and relations, and the 'reflection' premise that we can nevertheless apprehend the temporal features of experience *via* apprehending the properties of their objects in introspection (Lee 2014c p.156). I take this distinction to be the same as I made above, consisting of the negative thesis that we cannot directly apprehend 'inner' properties in introspection, but we can indirectly encounter them, though apprehending the (temporal) properties of their objects. In critiquing the transparency premise, we could either deny that we have any introspective access to the temporal features of our experience at all, or we could reject the claim that introspection is transparent in the sense that the *only* access we have to the temporal features of experience is through reflection on their objects. I formulate these two options as follows:

- **No access:** we have no access to the temporal structure of our experience.

- **Other access:** we have access to the temporal structure of our experience that does not depend on apprehending the properties of its objects.

Note that both these claims are ways of rejecting premise 1 of the argument, and thus premise 2 by extension. The 'no access' claim constitutes a rejection of a presumption

in the transparency premise: that we can, in fact, determine the temporal structure of our own experiences. The ‘other access’ claim is more a straightforward denial that there is a transparency constraint on our temporal experience. Are these claims plausible, and how might they be substantiated?

It would be a neat and satisfactory way of rejecting Phillips’ argument for the advocate of divergence to simply deny that we have any access to the temporal structure of experience at all; there would be no need to engage with the details of the argument, which could just be rejected as containing a false presupposition about experience. A spatial analogy might be pushed here: just as we have no introspective access to the *spatial* properties of experience (if there are any), why not too with time? This line of thought would also constitute an explicit denial of the ‘uniqueness’ assumption that I mentioned earlier, and for some might be too radical a denial of what is generally taken to be an obvious fact about experience: that experiences change, succeed each other etc in time. Faced with a view on which we have no access in this manner, a supporter of Phillips’ argument could just make a ‘Moorean’ assertion that this view is obviously false: commonsense notions of the stream of consciousness, the unity of consciousness over time, and the continuity of conscious experience presuppose that the basic ‘no access’ view is false.

However, a more promising line can be taken on this point when we consider a weaker ‘no access’ view: a view that holds that we have no access to the *boundaries* between experiences (Frischhut 2014), and thus no access to the duration of experience. On most views party to the debate about temporal properties, the basic picture is one on which individual experiences are unified over time to form the stream of consciousness (Dainton 2010a); the manner in which they are unified is the matter of some debate. The unity of experience over time (‘diachronic unity’) is generally taken to be necessitated by the fact that our conscious experience seems seamless and continuous; or rather, the seamless nature of the stream of consciousness, within which experiences seem to flow one into the other, is seen by many theorists to indicate that experiences are

diachronically unified. Reflection on what this unity consists in points to an alternative, restricted interpretation of ‘no access’ that would serve as a more plausible rejection of Phillips’ premise 1 in the argument. To see this, we need only note that if experiences are unified to form the stream of consciousness, we are not aware in our experience of the joins, or ‘boundaries’ between each experience. As Frischhut states, this is just what the apparent continuity of conscious experience may consist in; an inability to phenomenally individuate experiences when we introspect (Frischhut 2014 p.53-54). For James (1890), the subjective continuity of experience consisted in experiences *‘being without breach, crack or division’* (p.237): as Dainton (2010a) elucidates, *‘a succession of stream-phases belonging to S could be said to be continuous if S is incapable of discerning any gaps between them’* (section 3). However, once this observation is on the table, an obvious question is the question as to how we are able to have access to the *duration* of experience if we have no introspective access to the boundaries between experiences. Though it seems as if we represent the temporal property of duration in experience in the sense that when we hear a long note, for example, we can hear something of the duration of the note itself, it is hard to see how the duration of any particular experience could be specified based on introspection without first specifying when the experience begins and ends. Specification of the beginning and the end of an experience would seem to naturally require apprehension of the boundaries of the experience, but as we have seen we have no access to these boundaries in introspection. If this reasoning is correct, and we have no introspective access to the duration of experience, then premise 1 of the argument fails, since we have given an example of a temporal property that doesn’t abide by Phillips’ transparency premise. Whilst not being a decisive criticism of the argument, this reasoning if sound places the burden on the supporter of the argument to specify exactly how we have access to the duration of an experience given that we cannot discern its boundaries. And if it is conceded that we don’t in fact have introspective access to the duration of an experience, then a non-arbitrary reason must be given for why the transparency premise does not apply to this temporal property, but does apply to others, such as succession.

There is of course an ambiguity in the notion of ‘access to duration’ that is being employed in the above argument. It is one thing to apprehend the property of duration with regard to experience; that is, seeming to one that one’s experience has duration. However, it is another thing to have access to the *specific* duration of one’s experience; that is, seeming to one that one’s experience lasted, say, 2 seconds. It is plausibly only the second reading of duration access that requires access to the boundaries between experiences; it is hard to see how we could specify the length of an experience in seconds without being able to specify when that experience begins and ends (i.e, where its boundaries are). However, having access to the fact that our experiences have *some* duration does not seem to be precluded by the fact that we have no access to the boundaries of experience: it seems possible that one could discern that an entity has duration or extension in time without being able to discern how much duration or extension the entity has. For this reason, we might wish to cast doubt on the more developed version of the ‘no access’ approach advocated by Frischuut (2014) above, or at least call for a better argument to the conclusion that we have no access to the temporal structure of our own experiences.

Sensory deprivation and covert attention shifts

Leading on from this, I will now consider the prospects of rejecting transparency by arguing for the possibility of discerning the temporal structure of experience by some way other than specifying the temporal structure of its objects: an ‘**other access**’ approach. Specifically, I will consider two kinds of cases that suggest that we can have access to the temporal properties of experience *independently* of the temporal properties of their objects. Consider first the idea of a *sensory deprivation tank*. When placed in the sensory deprivation tank, a subject is unable to experience anything by means of the senses; she is unable to apprehend any sensory data, with only her thoughts for company. However, in such a case, the subject might well have access to the pattern of her thoughts if she is conscious, and in particular be able to discern an ordering, or succession, of thoughts in her stream of consciousness. For example, on waking up in

the tank, she might first think ‘who put me in this tank?’, but then remembering that she signed up for the experiment (whatever it is), think ‘oh yes, I volunteered to be in here.’ In thinking the first thought and then the second, the subject would be aware of a succession of her thoughts; the feeling of one thought succeeding another in time. In doing so, she would be apprehending a temporal property: the temporal property of succession. However, no sensory stimulus was required in order for the subject to have this temporal experience; in fact, no sensory stimulus was present at all. The sensory deprivation tank case therefore seems to be an example, (conceivable, possible and actual), of how we can be aware of the temporal structure of our own experiences without being aware of the temporal structure of worldly objects: a conclusion that is in tension with Phillips’ initial transparency claim. Phillips holds that we can’t help but access and describe the temporal structure of experience via the temporal structure of worldly objects and events; however, the sensory deprivation tank case seems to pose a situation where we can access and describe the temporal structure of experience *without* having access to or describing the temporal structure of the world.

The sensory deprivation tank case prompts reflection on the kind of experience we have of our own stream of consciousness, whether being aware of a succession of thoughts is enough to be properly experiencing succession as I have currently construed it (as an immediate presentation or representation similar to when we experience colour in the visual case), and what notion of transparency is required in order for Phillips’ claim to stand. In his account, Phillips (2010) is not at all clear what *kind* of transparency thesis he wishes to uphold, and more specifically, whether transparency is intended to be restricted to the *sensory* case, as opposed to the general *experiential* case. If Phillips’ transparency point (that ‘when we set out to describe our experience itself, we find ourselves doing so, at least partly, by attending to its objects’) is unrestricted and so applies to non-perceptual experience such as that had by the subject in the sensory deprivation tank, then it seems that the case provides a straightforward counterexample to transparency, and so Phillips’ argument can be rejected. However, we might also consider the possibility that Phillips wishes to restrict his transparency claim to the case

of sensory perception; that is, *sensory perceptual experience* is transparent in the manner specified, to say nothing of other kinds of perceptual experiences which are not sensory by nature. Much of what Phillips says elsewhere in his work suggests this restriction, especially considering the close analogue focussed on between temporal and spatial perception, at least when considering the *immediateness* of the perception of temporal properties with other non-temporal properties. If we grant the restriction of the transparency premise to sensory perception, then the sensory deprivation tank case provides no challenge: it is not a case of sensory perception. However, an alternative case is posed by Lee (2014c), who presents a putative case of ‘other access’ that does not involve sensory deprivation, but a genuine case of sensory perception. In pursuit of this strategy, Lee focuses on proposed ‘covert attention shifts’, purported cases where we can be aware of our experience changing in a way that is not tied to or dependent upon external changes in the way that transparency would dictate. Lee gives the example of a case where an observer’s visual and auditory scene is identical, but nevertheless where one’s ‘inner’ attention is shifted:

‘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’re looking at might not be changing at all. Still, you experience a change as happening - a psychological change....it is not merely that your experience changes when you shift your attention: you have an experience of change happening....for example if you shift attention back and forth between two objects at a certain rate...you can be aware of this rate.’ [Lee 2014c p.158].

The way in which the possibility of these cases is supposed to challenge transparency is as follows. If the external, worldly scene we perceive is qualitative identical (in the sense that all the same external properties and relations are fixed), but we can still shift our attention to aspects of the scene and be aware of the change involved in the attention shift, then we have a proposed counterexample to transparency: a case where

we have introspective access to a temporal property of our experience (change or succession) which cannot have come via the properties of the worldly scene we are perceiving. The possibility of such covert attention shifts does not simply come from reflection on how we seem to perceive the world, but also from psychological research. For example, Hoffman (1998) concludes that *'eye movements directed to a location in space are preceded by a shift of visual attention to the same location'* (p.119). Furthermore, not only do covert attention shifts precede and play a vital role in eye movement, but covert attention can vary whilst the eye position remains fixed (e.g. Eriksen and Hoffman 1972). As Hoffman states, *'The relationship between attention and eye movements is one of partial interdependence. Attention is free to move independent of the eyes, but eye movements require visual attention to precede them to their goal'* (Hoffman 1998 p.120). Such research lends plausibility to the claim that if the visual scene is fixed (as well as scenes in other sense modalities), and thus there is no change in the external properties perceived, introspective perception of experiential properties can vary, leading to apprehension of change and succession without a corresponding perceived change in objects of the experience. Though it is not obvious how it is that such psychological discernments are made (which theory of attention we go for will be a matter for psychology), all that is needed to resist Phillips' first premise, even when it is restricted in the manner I suggested above, is some evidence for the possibility of covert attention shifts.

The two examples presented in response to the transparency claim in Phillips above differ as follows. In the sensory deprivation tank case, the subject experiences temporal phenomena (the succession of his/her thoughts in the stream of consciousness) in a purely non-sensory way. Such an example can only hold force against the transparency point if the transparency point is intended to apply to non-sensory perception; in Phillips' account, the non-restriction of transparency in this manner is not made clear. Supposing that the transparency point, and therefore the mirroring thesis that it is introduced to support, is intended to be restricted to sensory perception only (i.e. experience only temporally mirrors its objects in cases of sensory perception), the

example of covert attention shifts can be introduced as a more compelling counterexample to this restricted transparency thesis.

Different kinds of seeming

However, suppose these cases are met with an adequate response, and the supporter of mirroring continues to claim that perceptual experience is transparent; how might we critique the argument based on the remaining premises? The crucial claims are contained within premises 3 and 5; premise 4 is simply a consequence of premises 2 and 3. Premise 3 states that rational judgment about experience entails experiential seeming; premise 5 states that we are not systematically deceived when it comes to how our own experience seems. I consider now both these premises beginning with premise 5, the ‘seems->is’ premise. This premise states that seeming is a good guide to the nature of experience: if perceptual experience seems a certain way (in a sense to be specified), then it really is that way. Lee (2014c) terms this an ‘infallible’ conception of awareness: the essential claim to Phillips’ argument is that there is ‘no appearance-reality’ distinction for experience (Lee 2014c. p.157). If our experience seems a certain way to us, then it really is that way; this claim is consistent with cases of illusion and hallucination since it only concerns what we find in introspection rather than what is actually the case in the world perceived. However, it is hard to conclusively accept or reject this principle without having a clearer notion of what sense of ‘seeming’ is being employed here. I argue, following Frischhut (2014), that each of the two alternative readings of ‘seeming’ land the argument in trouble; only one kind of ‘seeming’ makes premise 5 plausibly true, yet this reading of ‘seeming’ raises concerns about premise 3. Conversely, premise 3 is only made plausible by interpreting ‘seeming’ in a way that brings premise 5 into doubt.

- **Phenomenal seeming:** a ‘what-it’s-like’ notion of seeming concerning the qualitative aspects of experience, or ‘qualia’. For experience to seem a certain way is for it to have a distinctive ‘raw feel’.

- **Cognitive seeming:** if it cognitively seems that P, then S endorses P (Frischhut 2014 p.51). A notion of seeming concerning assenting to certain propositions.

Premise 5, the 'seems-is' premise, states that if experience seems a certain way, then it really is that way. This claim is only plausibly true if we take seeming in the phenomenal sense; indeed, it is very hard to see how things could be otherwise on this reading. If we conceive of experiences as having to a large degree 'phenomenal character' as well as representational content, as they are commonly conceived, then Phillips' seems-is premise is almost trivially true: it is very hard to see how we could err about the nature of our experiences on this conception of seeming, given the intimate acquaintance with our experience that it implies. However, if seeming is taken in this intuitive sense, then premise 3, the judging-seeming premise, looks to be in doubt. This premise states that if S judges that P, then it seems to S that P; however, it is possible to conceive of all sorts of situations where a person might judge something about experience, but lack the vital phenomenological component required for this reading of 'seeming'. Rational judgement and phenomenal seeming simply appear to be independent modes of coming to believe a proposition about an experience. For example, I might rationally judge that I'm experiencing a note of 2 seconds, and on that basis rationally judge that my experience itself lasts 2 seconds. However, the experience might not phenomenally seem to be 2 seconds long: it might not have any phenomenal component, though it also might not phenomenally seem to be of a different duration (Frischhut 2014 p.50). The phenomenal aspects of experience can appear to come apart from rational judgments about experience in a way that is inconsistent with the latter entailing the former.

In order to get around this worry, the proponent of Phillips' argument may wish to employ some other notion of seeming: 'cognitive seeming'. On this reading, if it seems to S that P, then this just entails that S is disposed to assent to P, or endorse it. This notion lacks the phenomenological component, and as such appears to be a concept much closer to rational judgment. Indeed, it is very hard to specify a situation where rationally judging proposition p about experience to be the case does not entail

assenting to p; perhaps no case exists. However, once this alternative notion of seeming is employed, premise 5 loses its force. Whilst in the case of phenomenal seeming, the seems-is claim was plausible, cases can be specified where cognitive seeming that p does not entail that p; an appearance-reality distinction can indeed be made with this new conception of seeming. We might assent (on the basis of a good argument, perhaps) to all sorts of things about our own experience if they cognitively seem to be the case; however, that they aren't in fact the case is a real possibility. For example, I might be convinced by some compelling arguments (e.g. in Chalmers 1996) that property dualism is true, and so my experiences involve an irreducibly mental component; however, it is also possible that property dualism is false. In this case, various propositions about experience are assented to on the basis of cognitive seeming, but this fact alone does not entail that those propositions are true of experience. And if the seems-is premise fails, then the conclusion that the temporal properties of experience and its objects match doesn't follow.

Conclusion

The above reflections on Phillips' argument should not be taken to be conclusive rejections of the argument for temporal mirroring. With regard to the 'no access' and 'other access' rejections of premise 1, possible responses are available to the mirroring theorist, who might spell out ways in which we needn't be able to discern the boundaries between experiences in order to discern their duration, or who might deny that altering one's inner attention on a scene is sufficient for the experience of temporal properties such as change. Furthermore, it might be disputed that a distinction can be made between rational judgment and phenomenal seeming in the manner I have suggested above. However, the force of the argument in this chapter is as follows. The conclusion Phillips presents for temporal mirroring was put forward as following from ordinary reflections concerning experience; that is, when we reflect on our own experience, it turns out that mirroring is true, and therefore the divergence of the temporal structures of experience and object is impossible. I have given several reasons

why one might resist this strong claim, and show how it is far from obvious that mere reflection on how experience seems to us can be used to derive metaphysical conclusions concerning the nature of experience. If the points made above are good ones, then the burden is on the mirroring theorist and supporter of Phillips' argument to clarify notions of introspection and seeming that make the argument's conclusion plausibly follow. Until this is done, he/she who rejects metrical and topological mirroring can feel safe that Phillips' argument does not demonstrate their position to be impossible when attending to ordinary experience.

4. Postdiction cases and temporal experience

Introduction

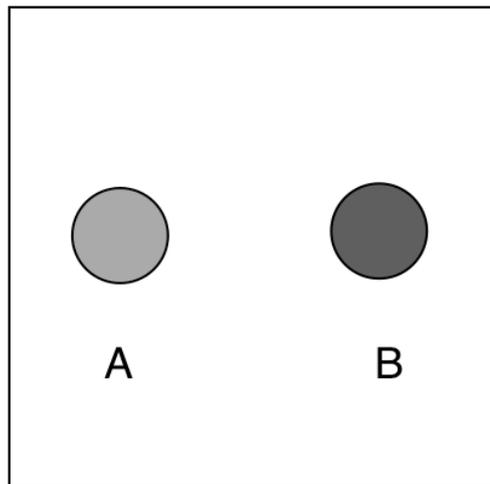
In chapters 1 and 2, I introduced the terms ‘inheritance’ and ‘mirroring’ in characterizing a major response to what I take to be a central question concerning temporal experience: what is the relationship between the temporal properties of objects, represented in experience, and the temporal properties of experiences themselves? Phillips (2009, 2010, 2014a, 2014b) advocates an ‘inheritance’ thesis: the view that experiences inherit their temporal structure from their (apparent) objects. In Phillips, inheritance is characterized as a claim of asymmetric dependence (of experiences on their objects), but also entails what I, following Lee (2014a) termed ‘mirroring’, which is the claim that the temporal structure of experience must match the temporal structure of its (apparent) objects. The central question of this thesis is whether ‘mirroring’ is a genuine constraint on temporal experience, or alternatively whether the temporal structure of experience can *diverge* from the temporal structure of its objects. In chapter 3, I assessed a central argument from Phillips (2010) for the mirroring thesis, and argued that it is unsuccessful; in this chapter, my aim is to develop a more positive argument against metrical and topological mirroring, and for the divergence of the two temporal structures under consideration which focusses on some empirical findings from neuroscience. I argue that the most natural interpretation of these ‘postdiction’ cases is one that implies that the above forms of mirroring are not genuine constraints on temporal experience; I then argue against a possible response from a supporter of mirroring that offers a different interpretation of the empirical results.

In the way that I have characterized the debate thus far, ‘inheritance’ entails ‘mirroring’, and ‘mirroring’ is false if it is shown that an experience can lack a temporal property (in our case, either duration or succession) whilst that same property is instantiated by its (apparent) object. I will here focus on one kind of empirical case that is often interpreted as fulfilling this role: the kind of cases that have been termed ‘postdiction’

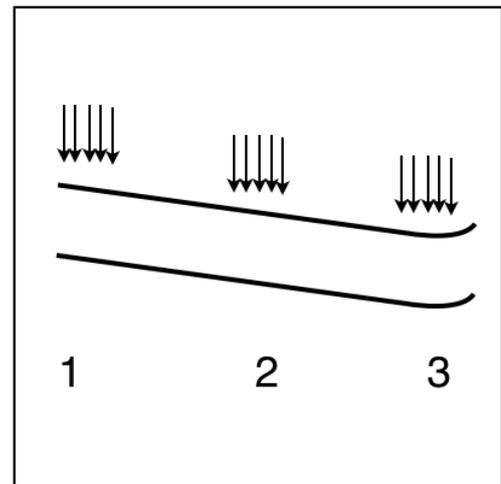
cases. After outlining the sort of challenge these cases have been taken to pose for the supporter of theories that require ‘mirroring’, I will consider a prominent response to this interpretation (Phillips 2014a), and argue that we should reject the response. I grant that back-and-forth about disputed empirical counterexamples doesn’t reasonably rule out the possibility that inheritance may be true of temporal properties nonetheless; however, I do conclude that the burden remains on the supporter of a mirroring constraint on temporal experience to deal with the empirical findings in a more convincing manner as a result.

Postdiction cases

The ‘color-phi’ case has been well discussed in recent debates, with Dennett and Kinsbourne (1992) taking it as a central example in motivating their ‘multiple drafts’ model of consciousness, a view that has the divergence claim at its very heart. When shown two flashing dots, one following another on opposite sides of a screen, it is generally reported that the dots are perceived as moving back and forth between the positions where it was actually depicted; some subjects in various experiments concerning color-phi also reported perceiving a gradual colour change from red to green and back in the course of this motion. In other words, when confronted with the flashing dots, illusory motion and change are represented in our experience. In the ‘cutaneous rabbit’ case, subjects were presented with a set of 5 short pulses at the wrist, followed by 5 short taps at the mid-forearm, followed by 5 short taps at the elbow (Geldard and Sherrick 1972). Though there were only three actual locations which were presented with the stimuli, subjects reported perceiving a greater series of evenly spaced taps ascending up the arm from wrist to elbow. In other words, illusory succession was represented in experience. The two cases are visually represented in the following pair of images; A and B, and 1, 2 and 3, represent different times at which each of the associated stimuli were administered in each case.



Color-phi



Cutaneous rabbit

The pertinent fact about cases of this form is not simply that false presentations, or non-veridical representations, of temporal properties were induced in subjects; the mirroring thesis as described above is perfectly compatible with such cases of illusion. The supporter of mirroring may hold that properties (re)presented in experience are mirrored by the experience itself, regardless of whether those representations are veridical or not. What's intriguing about both cases is that the experience of motion in the Color-phi case is seemingly dependent on the experience of the second stimulus (cutaneous rabbit: second set of stimuli), despite the fact that the subject reports experiencing motion *before* the second stimulus/set of stimuli are received. The dot does not appear to move towards the right from the left if there is not a second stimulus on the right in the color-phi case, and there is no illusory rabbit if only the first set of pulses are administered in the cutaneous rabbit case. The crucial stage of the experience, using the example of the color-phi case, is the experience enjoyed by the subject at the flash of the second dot. It seems plausible to observe that this is a momentary (or very short) experience, since the stimulus (the flashing green dot) is momentary (or very short). However, what is represented in experiencing this stimulus is an extended temporal interval, which must be the case due to the perceived motion of the dot, and the fact that representing temporal properties such as motion in experience requires representing a temporal interval. In the cutaneous rabbit case, the point is the same: for the subject to represent the illusory rabbit hopping up the arm, she must at the point of the sixth stimulus at the

mid-forearm represent an interval of time as preceding this stimulus, containing representations of the rabbit at evenly spaced points on the arm. As with before, the experience of this stimulus is instantaneous, yet again represents the temporal property of succession and thus a duration of time. These cases seem to suggest that, *prima facie*, the temporal structure of experience can come apart from the temporal structure of its objects; that the forms of mirroring we have been considering are false. An instantaneous experience does not inherit its temporal structure from its objects if it also (re)presents a temporal interval.

Interpreting postdiction cases: Orwellian vs Stalinesque

Phillips (2014a) considers three possible interpretations of the above cases, only one of which (he claims) supports the divergence thesis. He first argues that the presence of an alternative interpretation of the empirical data significantly weakens the argument against mirroring-based theories, and furthermore that his own distinct view concerning how experiences are realized allows us to reconcile the data with temporal mirroring. The two main interpretations of the data Phillips describes can be described metaphorically as ‘amnesia’ vs ‘blindness’ (Phillips 2014a p.135). On the ‘amnesia’ interpretation (e.g. Grush 2007), we experience a static red dot, then we experience a static green dot, and following this (through sub-personal neural processing) the brain eradicates these experiences, replacing them with a representation of motion. Dennett and Kinsbourne (1992) refer to interpretations of this kind as ‘Orwellian’; in Orwell’s ‘Nineteen Eighty-Four’, the ruling Party *retrospectively* modifies past news items in order to produce a representation of the truth that suits them better. On the blindness interpretation, by contrast, we do not experience anything before the apparent motion; our brain ‘holds off’ until the second stimulus has been detected, and only then is the representation produced in conscious experience; the representation is adjusted at the sub-personal level before this point. Dennett and Kinsbourne refer to such interpretations as ‘Stalinesque’; the brain produces false representations *before* entering in conscious experience, similar to Stalin’s propaganda machine during the Cold War.

This interpretation, Phillips contends, is compatible with mirroring, since it allows the experience of the motion itself to be extended in time (i.e. possess duration and change) at the point of the second stimulus. It could be argued, for example, that the experience of motion is extended, but that the experience is had only after a delay, due to neural processing (Dainton 2010b).

The matter of which of these interpretations of the data we should choose is to a large extent an empirical matter; that is, presumably one of the interpretations is the correct one, and we might expect that further work on the issue by neuroscientists could in principle yield the answer. However, this does not preclude us saying something useful about the relative merits of each interpretation. A central question in interpreting cases such as color-phi and cutaneous rabbit is at what point representations are produced in the brain; whether representations are produced in conscious experience and then ‘overwritten’, whether representations are produced pre-consciously and modified before being introduced into conscious experience, or whether representation in experience at a time depends constitutively on experiences at later times. Though the question should plausibly be interpreted as a broadly empirical one, we can still note that if one of these interpretations entails a view about experience that is philosophically or intuitively problematic, then that might be a reason to reasonably reject that view. Dennett and Kinsbourne (1992) argue that the question ‘Orwellian or Stalinesque?’ with regard to these cases has no answer, since *‘the boundary between perception and memory, like most boundaries between categories, is not perfectly sharp’* (p.192). For Dennett and Kinsbourne, there need be no true distinction between the two rival explanations of how the brain creates the false representations involved in the color-phi and cutaneous rabbit cases, since to pose the distinction is to presume that in one case, memory revision is involved (post-experiential revision), whereas in the other case, modification of the representation is pre-experiential, before the representation is committed to memory. The two supposedly rival explanations, in Dennett’s view, seek to differ over whether the neural processing of the false representation occurred ‘before or after the fact’ of actual conscious experience, which (Dennett claims) presupposes a

ultimately misguided ‘Cartesian materialist’ conception of consciousness. Dennett and Kinsbourne’s key contention is that this model, which holds that there is a place in the brain ‘where it all comes together’ in terms of conscious representation such that there is an absolute representation ‘in’ actual consciousness at a time, is false. The Orwellian and Stalinesque interpretations of the color-phi and cutaneous rabbit cases are therefore interpretations that presuppose a false view of consciousness; thus, the question as to which of them is correct has no meaningful interpretation. To postulate a ‘*mythical Great Divide*’ (p.192) between ‘actual’ consciousness and other cognitive processes such as memory is simply misguided according to Dennett, as they presuppose a kind of central observer in the brain where neurally produced representations are presented at a definitive point in time for subjective judgment. The alternative view provided by Dennett and Kinsbourne is a ‘multiple drafts’ model in which neural discriminations are distributed in both space and time in the brain. Instead of a single, well-defined stream of consciousness which explains exactly when neural representations enter ‘actual’ conscious experience, there is simply a series of constantly revised and updated representations in the brain, cognitively produced. Conscious experience has no real phenomenal character that is taken to be conceptually distinct from the representational content of experience, in the manner Chalmers (1996) and others suggest.

It is customary in the literature to dismiss Dennett’s approach to consciousness on the basis that it neglects certain obvious and key aspects of our inner mental lives, aspects of our conscious experience which are manifestly the case when we reflect on our awareness of the world and its properties. Whilst Dennett’s position is worth considering and raises some interesting questions such as to the relationship between consciousness and other cognitive processes, I agree with this orthodox rejection. I have already argued in chapter 1 that the assumption of unity of consciousness and the fact that we represent temporal properties in consciousness with their associated phenomenal ‘feel’ is incompatible with Dennett’s approach, and we should reject it on this basis. Furthermore, we needn’t throw the baby out with the bathwater in order to deal with colour-phi and cutaneous rabbit cases; the Orwellian and Stalinesque

distinction is indeed an instructive one if we make the basic phenomenal assumptions concerning the unity and raw feel of conscious experience. I argue that the Orwellian, amnesia-based interpretation advocated by Grush looks more promising than the Stalinesque interpretation advocated by Dainton and others. The Orwellian interpretation is, it seems, the most natural interpretation that takes the data at face value; it does seem as if we have an instantaneous experience of motion at or shortly after the flash of the second dot, and this is manifestly incompatible with mirroring. By contrast, the Stalinesque ‘blindness’ interpretation constitutes an *ad hoc* addition to the properties of experience that is poorly motivated by any other empirical considerations. Dainton points out, in support of this interpretation, that since we already know that experience necessarily involves some delay between the stimulus and the time of the resulting representation in consciousness, why could this delay not simply be long enough to compensate for the sort of perceptual/experiential modification suggested by the postdiction cases? *‘Might it not be that our visual systems take some time (perhaps 80–100 msec) before producing experience in response to a given stimulus? And might they not use this time to work out a single coherent version of events before committing it to experience?’* (Dainton 2010b). The problem with this response is that it is poorly motivated by any independent concerns, and though Dainton cites one scientific paper that supports his contention claiming that unless this interpretation can be ruled out, the ‘extensionalist’ (for these purposes, the supporter of ‘mirroring’) has nothing to fear from postdiction cases, this response misses a key methodological point. Rather than taking mirroring to be the case and then simply choosing an interpretation of the postdiction data to match, we should focus on the data first and note what can be inferred about experience based on the best interpretation. And as I argue, Grush’s Orwellian interpretation is superior in that it does not involve making the little-supported empirical claim that experiences are subject to much more delay than has commonly been thought.

A new conception of experience as anhomoeomerous

Putting Dainton's interpretation to one side, I now address in more detail a better-motivated and more plausible solution to the problem that is *prima facie* posed for mirroring. Phillips (2009, 2014a) advocates this third response, which is close in structure to Dainton's blindness interpretation in that relies on the notion of an experience at a time being logically dependent on experience at a later time. On this view, on which experiences are taken to be 'anhomoeomerous' processes, we cannot exactly say what a subject experiences at a specific time without taking into account facts about later experience; for example in the color-phi case, whether a second stimulus is presented in experience or not. Phillips claims that this sort of alternative view about how experiences are realized by their underlying neural processes allows us to preserve mirroring in light of postdiction cases without the need for experiential overwriting (Grush 2007) or experiential delay (Dainton 2010b).

Like the 'blindness' approach of Dainton, this new 'holistic' approach of Phillips seeks to preserve mirroring in light of the cases, but I hold that it leads to problematic metaphysical commitments about the nature of experience. It is worth spelling Phillips' interpretation out in more detail, beginning with introducing the Aristotelian notion of 'homoeomery.' In Aristotle, we find the now well-worn distinction between 'things' and 'stuffs'; in Aristotelian terminology, 'primary' and 'secondary' substances respectively. For Aristotle, primary substances were concrete individuals of a particular kind, and countable, such as a certain block of gold, or a particular human being (say, 'Socrates'). 'Secondary' substances are general kinds to which the primary substances, or things, belong; for example, 'gold' would be the secondary substance, or 'stuff', to which the individual block of gold belongs. It is this second use of the term 'substance' that we are most familiar with in contemporary discourse: the notion of 'substance' as stuff, to be contrasted with individuals by the fact that unlike individuals, it does not admit of counting. For example, there can be many gold bars, but not many golds. For Aristotle, a primary substance is an object of predication; it cannot be properly predicated of any

object. As Hacker puts it, “*Socrates*’ cannot be said to qualify anything, or be true of anything.’ (Hacker 2004 p.41). It is the secondary substance terms that, when predicated of individuals, categorize an individual thing as something of a given kind; for example, ‘this is gold.’ The notion of a secondary substance is close to the Aristotelian notion of ‘form’, and as such can be used according to Aristotle to individuate individuals of a particular kind, or provide a criterion for counting them. So for example, to state ‘this is gold’ of a particular object, i.e. making a specification of an object’s secondary substance, is to specify the nature of the individual object.

The key characteristic of stuffs, or secondary substances, is that they are homogenous, or homologous, or *homoeomerous*; that is, almost any piece of portion taken of a substance will also be that substance kind. It is important to add ‘almost’ here, since modern science has since revealed that many substances or stuffs commonly thought to be fully homoeomerous in Aristotle’s period have since been revealed to be homoeomerous only down to the atomic level. Gold is more homoeomerous than wood, for example, since one can subdivide gold all the way down to atomic level and still have gold (though this fails to hold at divisions smaller than the atomic level); wood is clearly not homoeomerous to this extent. Based on these observations, we should take the notion of homoeomery, or ‘like-parted-ness’ (Phillips 2009 p.97), to admit of degree, rather than being a binary notion that is either predicable or not of a substance.

Homoeomery: a substance S is homoeomerous to a degree x if for any portion s of S up to x , s necessarily is S .

If a substance is homoeomerous to a degree x , then it is *anhomoeomerous* past x ; that is, beyond the point of x , s can fail to be S , hence the ‘necessarily’ clause in the definition of homoeomery. This means that past x , it is an open question whether s is S ; we need to know more before we can identify this portion as being S . It may well be the case that all stuffs are *homoeomerous* to some degree; the point at which homoeomery breaks down for a substance will vary depending on the substance kind in question. We have

seen that gold is significantly homoeomerous; for gold, x is the molecular level; below this level, gold is said to be *anhomoeomerous*. But how about fruitcake (Taylor 1985)? Fruitcake is plausibly a substance for which homeomery breaks down at the fully observable level; if we subdivide fruitcake continually, it will not be long before we come upon a portion whose nature, or status, we cannot at first specify, for example a sultana (Phillips 2009 p.98). A sultana taken by itself does not belong to the substance kind ‘fruitcake’: it is simply a sultana. A piece of fruitcake that is a sultana can fail to be fruitcake, and so we might say that at the sultana-level, fruitcake as a substance kind is *anhomoeomerous*. However, taken in the context of a volume full of fruitcake, it is plausibly true to predicate an individual sultana to be of the substance kind ‘fruitcake’. This observation shows that the status of some portions of substances can depend on the nature of the surrounding substance: a sultana in a bowl of sultanas does not fall under ‘fruitcake’, but a sultana in a bowl of fruitcake *does*. The fact that fruitcake as a substance is significantly anhomoeomerous even at the observable level is explained by the fact that fruitcake is a ‘bitty’ substance, having many different kinds of object as large parts, for example nuts and fruits (including sultanas).

The key point relevant to Phillips’ argument with reference to the fruitcake example and the definition of *homeomery* above is an epistemological point; at the point at which a mass of stuff is anhomoeomerous, we cannot specify the nature of a part of that stuff without taking into account the nature of the surrounding volume. A sultana will fall under ‘fruitcake’ if it is surrounded by fruitcake, and not if it is surrounded by sultanas. This epistemological observation is the consequence of an stuff being anhomoeomerous at some level; the problem of specification applies at the level of anhomoeomery. In order to specify the substance kind of an object or portion of stuff, we need to know about the context in which the stuff occurs.

Up to this point, we have been concerned only with substances; we have seen that different substances can be said to be homoeomerous down to different levels. Phillips builds the point about homeomery into his point concerning the nature of experience

via a well-worn analogy between stuffs and processes, discussed in Mourelatos (1978) and developed by Taylor (1985). It has commonly been noted that the event/process distinction and the individual/substance distinction are closely analogous, with events and individuals sharing many features, and processes and substances likewise. One example is countability. As Gill (1993) points out, events are intrinsically countable in a way processes are not, and the same is true with things and stuffs; the way in which these entities are individuated is largely the same. In the specific case of homoeomery, processes are plausibly homoeomerous to some degree, whereas events are not; events do not have a structure of parts of which we can predicate things that are true of the events themselves. Phillips (2009) gives the example of the process of walking: if person S walks for a period, then at each moment in this process, it is true to say that S is walking. However, if we single out S's position at any moment in the process and consider it in isolation, it would in many cases be impossible to specify whether S is walking, or instead standing in some odd position. The epistemological point again applies: to know whether S is at that moment walking, we need to be able to specify whether or not S is in the process of walking, and we cannot do this without knowing facts about S that obtain at earlier and/or later times; the analogue of the surrounding volume in the spatial case. The epistemological claim can be thought of as a thesis of explanatory dependence: if a process is anhomoeomerous, the status of the temporal parts of the process is to be explained in terms of the status of the whole process.

Both processes and experiences are necessarily temporal entities; both occur over time. The shared temporal aspect is what allows Phillips to develop his account of experience as *anhomoeomerous* process. Phillips (2014a) states, quoting Dainton (2000):

'When we come to explain the nature of the stream of consciousness there are significant, extended periods over which we must explain the properties of sub-parts in terms of the properties of the whole duration and not vice-versa. That is, the stream is structured such that over short periods the explanatory direction runs from temporal whole to temporal parts.' (Phillips 2014a p.96).

As we have seen, if a process is *homoeomerous* then it can be analysed down to a series of independent short chunks or slices, which are taken to have explanatory value for the whole that they constitute: the entity itself. Phillips' claim is that experience is an anhomoeomerous process; in that the temporal parts of an experience do not explain the existence of the experience itself, but vice-versa; the existence of the experience itself is explanatorily prior to the existence of its parts. We might refer to this notion as a kind of holism; the whole is taken to be prior to the parts in some sense to be specified. We might also refer to the contrasting view as a kind of atomism; the 'atoms' of an experience, or its parts, are taken to be in some sense prior to the experience as a whole (Lee 2014a, 2014b).

We are now in a position to see how Phillips' conception of experience allows a response to the aforementioned postdiction cases to be formulated. The epistemological point was that we (often) cannot say at particular times what a subject is experiencing, as experiences are extended processes that are connected to form the stream of consciousness. Therefore, insofar as taking an Orwellian or Stalinesque interpretation of the postdiction data involves specifying the nature of an experience at a time (for example, claiming with the Orwellian that 'at t_1 , a duration d is represented in experience'), we should reject such interpretations. The epistemological consequence of experience being anhomoeomerous is that we *cannot specify* the nature of experience at a time in this manner. We might see the manner in which Phillips' strategy as avoiding the problem as similar in form to that of Dennett and Kinsbourne (1992), albeit with a very different conception of experience at hand. Both accounts hold that the question 'Orwellian or Stalinesque?' in response to the problem posed by the postdiction cases falsely presupposes a model of experience that should be rejected. In Dennett's case, it is the so-called 'Cartesian materialist' position that falsely has it that there is a time and a place in the brain where actual consciousness takes place; and in Phillips' case it is the assumption that we can specify the nature of experience at a time, where in actual fact (it is claimed), we cannot; all experiences are anhomoeomerous in that they are dependent on past and future experience for their nature.

The temporal correlation principle

Having already rejected the approach of Dennett and Kinsbourne, I now pose a problem for the account of Phillips. Phillips' account of experience as an anhomoeomeric process seeks to allow the supporter of temporal mirroring to reject the Orwellian interpretation of the postdiction cases; indeed, it seeks to preclude the need for any mirroring-endangering interpretations at all, since such interpretations must involve stating facts about experience at a time, which is something we cannot do if experience has this nature. Phillips' account rests on three key claims: the claim that processes are like stuffs, and the claim that experiences are processes, and finally the claim that experiences are significantly anhomoeomeric processes. The first of these claims I will leave alone. The second claim, that experience is a process, is a significant claim about the nature of experience. Since the question as to the nature of experience isn't the central question of this essay, I will refrain from criticizing this claim, though it is a claim that those who reject mirroring will also likely wish to reject; for example, Lee mounts a strong case against a 'process' view of experience in favour of an 'atomic' approach (Lee 2014a). The key claim is that experience is anhomoeomeric, and it is this claim I argue is implausible, at least on a popular view in the broader philosophy of mind, when we consider the relationship between experiences and the neural processes that realize them.

The majority of the discussion of experience in Phillips (2009, 2010, 2014a, 2014b) focusses on the relationship between the structure of experience and the structure of its objects; no mention is made of the equally important relationship between the structure of experiences and the structure of their physical realizers. The most straightforward way to think of the relationship between experiences and neural processes is, put simplistically, a straightforward correspondence in which the timing of experiences is intimately related to the timings of their underlying neural processes. In the broader philosophy of mind, this kind of view has been associated with physicalism. On this picture, a token experience of, e.g., pain is realized by a specific neural process; for

example, C-fibres firing. Furthermore, we can individuate the neural states or events which are taken to be the realizers of the experiences as finely or as broadly as we like on this picture for the important point to remain: individual experiences are realized by individual neural events and each experience has a neural ‘realization base’ (Kim 1992). Contrary to this however, the view of experience Phillips seems to be presenting is one on which experiences can ‘float free’ of their neural realizers. The neural state of a subject’s brain immediately after *A* and before *B* in the color-phi case is presumably identical to the neural state of the same subject’s brain if she were to just experience *A*, without *B* occurring, yet Phillips’ account suggests that despite this fact, the subject will have different experiences in both cases, due to the fact that in the first case, a further experience at *B* takes place. Consider the following brief argument:

- 1) If experience is anhomoeomeric, we cannot specify facts about its nature at a time without taking into account facts about experiences at earlier or later times.
- 2) Facts about the timing of experiences (for example, ‘experience *e* begins at time *t1*’ or ‘part of experience *e* occurs at *t2*’) are facts about the nature of experience at a time.
- 3) We can specify facts about the timing of experiences.
- 4) Therefore, experience is not anhomoeomeric.

The premise doing most of the work in this argument is premise 3; why should we accept it? Lee (2014a) formulates the following plausible ‘temporal correlation principle’:

TCP: *‘if two experiences are realized over the same interval or moment, then they themselves occupy the same moment or interval’* (Lee 2014a p.4).

TCP is a principle concerning the relationship between the timings of experiences and their neural realizers; more specifically, the claim that the timings of experiences and their realizers are systematically correlated. We needn’t read this principle as

entailing that experiences have the *very same* timings as their realizers (Lee: the stronger ‘temporal identity principle’); the weaker temporal correlation principle is consistent with an experience happening, for example, at the end of its neural realization process. But as we have seen, Phillips’ view of experiences as anhomoeomeric processes that are dependent for their specification on later experiences falls foul of this principle; an experience of the red dot in the color-phi scenario and an experience of a red dot outside this scenario are realized over the same moment, but on Phillips’ view, the experiences could differ in terms of their timings. That is, *‘the parts of an experience can happen at different times, despite being physically realized over the same temporal interval’* (Lee 2014a p.7). TCP entails premise 3 in the above argument, *given that* we can specify temporal facts about the neurological states that realize our experiences, and I take this to be an uncontroversial claim. All that is needed is a reason to support TCP, and the above argument can be made.

What reasons do we have for thinking that the temporal correlation principle TCP is a plausible principle? Lee (2014a) points out that it is very hard to make sense of the view that experience is anhomoeomeric in this way when we specifically consider the relationship between experiences and their neural realizers. Lee considers a spatial analogue to Phillips’ proposal: consider the claim that experiences of different regions of space *themselves* occupy different regions of space, for example the claim that an experience of a banana is banana-shaped. If we consider this claim alongside the claim that experiences are anhomoeomeric, then we are forced into the conclusion that two experiences could be located in different regions of the brain whilst being physically realized in the same brain-shaped region. It is very hard to make sense of this claim; however, the supporter of Phillips might simply respond that a fundamental difference in the perception of space and time allows us to rule out the idea that experiences could be predicated with spatial properties in this regard. A rejection of Lee’s spatial analogy would be motivated by a claim similar to the ‘uniqueness’ claim central to Phillips’ argument for mirroring I described previously in chapter 3; if we reject that spatial properties can be predicated of temporal properties in the manner Lee suggests, then the

analogy will not hold. Nevertheless, the main point in favour of the temporal correlation principle is that it is entailed by any form of physicalism that holds that mental states and physical states are systematically correlated; and this is true on both type- and -token identity theories. On the type-identity theory, for every mental property instantiated there is a (strictly) identical physical property; on the token-identity theory for every actual object, event or process, there is a (strictly) identical physical object, event or process. Both these physicalist theses entail the TCP, since on both theses, token neural states are taken to be strictly identical to token mental states; the token-identity theory leaves open the possibility of particulars also having contingent mental properties, whilst the type-identity theory does not. If every mental state is identical to a physical state, and if experiences are mental states, then it follows that experiences will have both the same intrinsic and extrinsic properties as the neural states that realize them, and this will include properties of experiences such as ‘beginning at t_1 ’, and other properties concerning the timing of experiences.

For our present purposes, the key claim to note is that any account that involves a denial of the temporal correlation principle will necessarily involve a denial of both type and token identity theory more generally as an option when it comes to explaining the nature of mental properties. This of course does not prove a view to be false if we don’t already assume some specific forms of physicalism to be true. It does however, place serious metaphysical constraints on an account which we would like not to be burdened with. If it is indeed the case that holding experience to be anhomoeomeric entails the denial of some popular forms of physicalism, then this should make us (and more specifically, the advocate of mirroring) think very carefully before adopting an account of experience along these lines. Grush’s original Orwellian interpretation provides just such an account without entailing any thesis that might force us into wider commitments when it comes to the nature of mind. Unfortunately, this option is unavailable to the advocate of mirroring, whose account of the relation between the temporal structure of experience and the temporal structure of its apparent objects is not consistent with an Orwellian interpretation. Therefore, without arguing explicitly for a

physicalist thesis, we should still conclude that Phillips' response to the postdiction cases is unsatisfactory; the thesis provides a view of experience that is far from intuitive, and one that is incompatible with a major and popular thesis in the broader philosophy of mind.

Russell worlds

The conception of experience under discussion is a conception of experience as anhomoeomeric: a process extended in time, and dependent for its nature at a time on earlier or later experience-phases, such that we cannot specify the nature of experience at a time without knowing more about these earlier and later phases. Even leaving concerns about the temporal correlation principle aside, and focussing not on the physical realizers of experiences but the experiences themselves, another objection can be levelled against this picture. This is the objection that the claim about specifying the nature of experience is false: if we can after all specify the nature of experience at a time, then the claim that experience is anhomoeomeric is false. If it could be shown that this specification could (at least in principle) be made, then further doubt is cast on Phillips' conception of experience. Consider the following argument, closely related to the previous argument, but more general in form:

- 1) If experience is anhomoeomeric, we cannot specify facts about its nature at a time without taking into account facts about experiences at earlier or later times.
- 2) We can specify facts about the nature of experience at a time without taking into account facts about experiences at earlier or later times.
- 3) Therefore, experience is not anhomoeomeric.

What reasons can be advanced in favour of premise 2? Firstly, we might wish to make a common-sense argument for this premise, based on observations about how we ordinarily take experience to be. At this current moment, for example, I seem to have no problem specifying various facts about my conscious experience: that I am

concentrating on writing, that I am hearing the hum of my laptop, or that I am looking at a certain visual scene. However, these specifications are of experience at a ‘high level’ of observation; with regard to the postdiction cases we have been discussing (and which the view that experience is anhomoeomeric is trying to account for), experience at the level of seconds and nanoseconds is the focus. A better argument is needed for the ability to specify facts about experience at a time at the more fine-grained temporal level than simply making assertions about what I am now experiencing. A more promising way to argue for premise 2 in the above argument is, I argue, to consider a notion that has been called a ‘Russell world’. In ‘*The Analysis Of Mind*’ (Russell 1921), Russell invites us to consider a world identical to the actual world, but that sprang into existence only moments ago; this world is identical to the actual world but has a different history, one that began just before the present moment. If it is true that the actual world *could* be this world, then the Russell-world is conceivable. Though the example of a Russell world has been used to illustrate a skeptical scenario concerning our knowledge of the actual world, the conclusion Russell himself draws (and a conclusion we should consider too) is that ‘*there is no necessary connection between events at different times*’ (Russell 1921, lecture 10). If the state of the world as it is at the present moment could have emerged a moment ago, rather than as a product of long causal history, then the state of the world at a time cannot depend on the state of the world at a past time. We might also consider a closely-related case, a Russell2-world: a world identical to the actual world, but suddenly goes out of existence immediately following the present moment. If a Russell2-world is a real possibility, there is no necessary connection between present events and future events; the state of the world at the present moment cannot depend on the state of the world at a future time. Experiences, conceived of as either events or processes, cannot be anhomoeomeric if the Russell-world and the Russell2-worlds are real possibilities (i.e. could be the actual world), since as we have seen for experience to be anhomoeomeric is for it to depend constitutively on past or future times. The ‘necessary connection’ between events at the present moment and past and future events cannot exist if either of these worlds are possible, but premise 2 of the argument relies upon this necessary connection. If the

Russell-world (and the Russell2-world) are genuine possibilities, then we *can* specify facts at a time without taking into account facts at other times, since there is no necessary connection between these facts and present facts. As Phillips himself concedes (Phillips 2010 p.181), the case from Russell worlds could be strengthened by making a similar claim specific to *physical* facts. If we restrict the class of facts that do not depend on facts at other times to physical facts, and then hold to a physicalist thesis on which mental facts are identified with, or supervene on, physical facts, then it follows that mental facts (including facts about experience) do not depend on facts at other times. Here is a more explicit contrast with the thesis that experience is anhomoeomeric: experience cannot be anhomoeomeric if this argument is sound, though it clearly faces the same partial issue as the previous argument against anhomoeomery in that it assumes a popular view in the philosophy of mind (some form of physicalism) that nevertheless some may be reluctant to subscribe to.

It is clear to see that the Russell-world cases rely explicitly upon a move that has been controversial in contemporary metaphysics: the move from conceivability to metaphysical possibility. The reasoning in favour of the argument above, more explicitly stated, runs as follows: If Russell-worlds are metaphysically possible, then premise 2 of the argument is true. Russell-worlds are conceivable. What is conceivable is metaphysically possible. Therefore, Russell-worlds are metaphysically possible. Therefore premise 2 is true. Perhaps this move from conceivability to metaphysical possibility is the obvious step in the argument that the supporter of Phillips' view of experience would wish to challenge. However, this challenge would involve engaging in a wider and deeper debate concerning the metaphysics of modality. Given that the claim that Russell-worlds are conceivable is highly plausible (it is hard not to accept that a hypothetical God, if he had wanted, *could* have created the world as it is a moment ago, or obliterate it a moment hence), and given that a special reason or argument is needed to rule out the inference to possibility, I conclude that the Russell-world cases pose a strong *prima facie* challenge to the view that experience is anhomoeomeric.

Aside from these concerns, we might also advance a common-sense critique of Phillips' view that experience is anhomoeomerous. Rather than being a straightforward interpretation of the postdiction data, Phillips' response to the aforementioned cases is a full-blown thesis concerning the relationship between the stream of consciousness, and the individual stages, or experiences, that make the stream of consciousness. If experience is anhomoeomerous, then the nature of some experiences depends on later experiences in the stream of consciousness, and so it follows that it is the stream of consciousness from which the individual experiences derive, and not vice versa. Individual experiences do not in fact make up or constitute the stream of consciousness at all, but derive from it; their individuation depends on the stream of consciousness, which is metaphysically prior to the individual experiences. I think that it is hard to make sense of this view, at least without rejecting a common and compelling notion of experience, on which it is the stream of consciousness that is derivative from the individual experiences, and not vice versa (though of course, this has not explicitly been argued for here). Insofar as we are apt to think of the stream of consciousness as a succession of experiences, then it is perfectly natural to think of the stream of consciousness as nothing other than a collection of experiences, unified by some relation such that adjacent experiences will form the sort of continuity we generally take our conscious experience to exhibit. Phillips' conception of experience as anhomoeomerous involves rejecting this compelling picture, as well as rejecting the plausible temporal correlation principle mentioned above; furthermore, the view faces a special problem in having to rule out Russell-world cases. For these reasons, I think we are justified in rejecting Phillips' interpretation of these cases. The clear alternative, an Orwellian interpretation on which it is 'overwriting' in experience that gives rise to these illusions, looks most plausible in light of this. And if this is the correct interpretation, a significant challenge is posed for the supporter of mirroring: either explain how an Orwellian interpretation is *not* inconsistent with mirroring, or demonstrate more clearly why an Orwellian interpretation is the wrong interpretation of these cases.

Conclusion

The postdiction cases discussed above have been presented as a challenge to the theorist who holds that experiences inherit the temporal properties of the objects they represent. I have addressed one response to this proposed challenge, and argued that it is unconvincing; the supporter of mirroring must propose a better interpretation of the data to stave off the worry that the natural and best interpretation shows that the temporal properties of an experience and its objects can diverge. The cases as I have presented them certainly do not conclusively show that ‘mirroring’ is false; there is much more to be said about the merits and problems concerned with Grush’s model, and also that of Dainton’s approach. What I have tried to show, however, is that adopting an ‘anhomoeomeric’ model of experience and interpreting the postdiction cases in light of this in order to preserve the inheritance thesis is problematic. The burden is therefore on the supporter of mirroring to demonstrate in a more satisfactory manner how these cases are to be dealt with on his/her view; in the absence of such an explanation, I contend that we are justified in taking postdiction cases to pose a clear counterexample to metrical mirroring, and hold that the temporal structure of experience and the temporal structure of its objects can diverge.

5. Conclusion and further issues: the diachronic unity of experience

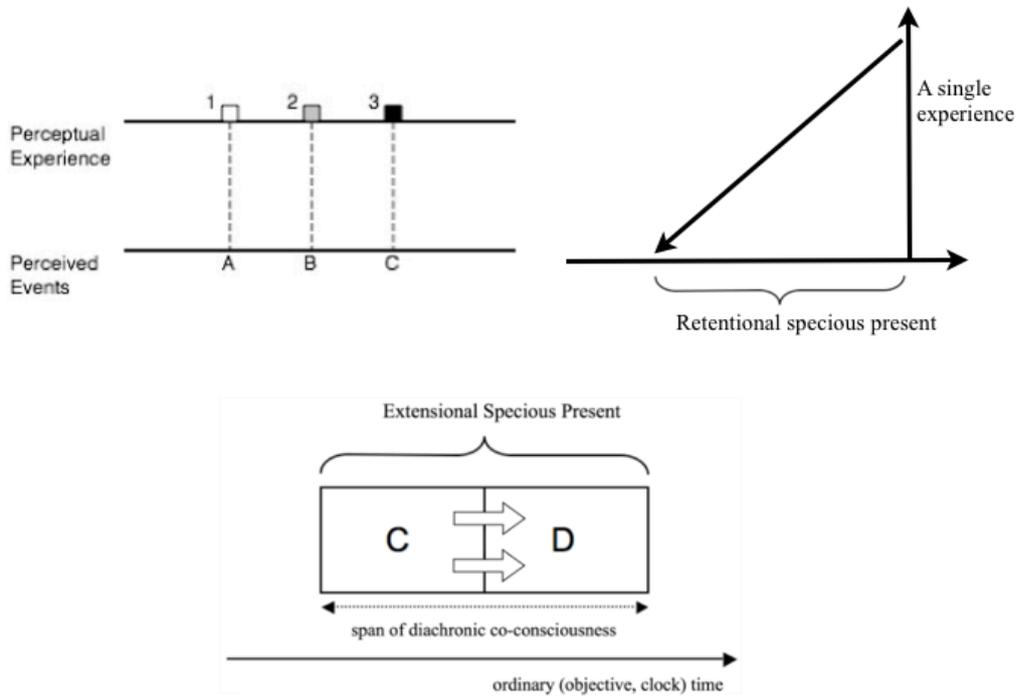
Conclusion

The topic of this essay has been the relationship between two temporal structures; that of experience itself, and that of the objects/events (re)presented in experience. After making clear why a bare ‘snapshot’ view of experience must be ruled out, with reference to three plausible constraints on temporal experience, I first set out Dainton’s approach to the debate, arguing that it does not adequately capture most clearly the issue to be discussed: the nature of the relationship between the two temporal structures. I then outlined a class of ‘mirroring’ views, with the aim of arguing that a strong form of mirroring, metrical mirroring, on which experiences inherit both ordering and duration relations from their objects, is false. In chapter 3 I presented several ways a major argument for metrical mirroring from transparency, that of Phillips (2010), could be resisted; either by calling for clarification concerning the specific notion of transparency required to make the argument, or casting doubt upon the notion of ‘seeming’ that plays a vital role in the argument. In chapter 4, I considered some well-worn empirical cases which I argued provide a further reason to reject metrical mirroring; cases where the two temporal structures appear to diverge in a manner not consistent with the mirroring claim. I then considered one possible response, Phillips’ view of experience as *anhomoeomerous*, and argued that this is not a plausible way to conceive of our experiences. Recall that the view under question, metrical mirroring, has it that both the succession and duration relations of events in experience are mirrored by experience itself. In Chapter 3, I considered cases where this thesis seems to fail for succession; the cases of the sensory deprivation tank and attention-shifting. In Chapter 4, I considered cases where this thesis seems to fail for duration; the cases of color-phi and cutaneous rabbit. In both cases, there seems to be (re)presentation of the temporal property in question without an associated experience possessing that temporal property; if the previous arguments are sound, then this constitutes a reasonable case against the mirroring principles previous outlined. Nothing in my critique rules out a *structural*

mirroring constraint on temporal experience (where only distinct stages of a perceived scene are mirrored by the experience itself); however, this is a very weak constraint, and one which is undoubtedly not robust enough to do the work required for either Phillips, or others who would wish to advocate his ontology of experience.

The above critique was intended to relate to the question as to the nature of the relationship between the two temporal structures. A central query along those lines might be: ‘do experiences mirror their objects in terms of temporal properties?’. I believe that we have reason to believe the answer to be ‘no’, based on the various considerations advanced in this essay. With this thought set out, it is instructive to return to the secondary question as to the nature of experience itself over time; the question that Dainton (2000, 2010a) seems most motivated by when formulating his own taxonomy of cinematic, retentionalist and extensionalist views. In this, the last chapter of the essay, I wish to briefly reconsider these views, and in particular the retentionalist view, which is most naturally suggested if the aforementioned claims concerning the viability of metrical mirroring are upheld.

Extensionalism and retentionalism revisited



Dainton's three models of temporal experience (clockwise from top left): cinematic (Grush 2007), retentionalist, and extensionalist (Dainton 2010a)

As we have already seen, the views in Dainton's taxonomy differ as to the extent to which either the content of the experience or the experience itself are held to be extended in time. The cinematic view is presented as having both the content and the experience to be non-extended, or at least very short if not momentary. I argued that two plausible constraints on temporal experience ruled this approach out: the phenomenological constraint (**PC**) that we perceive temporal properties themselves, and the relational constraint (**RC**) that in order to perceive temporal relations, we must perceive both relata in a single experience. The cinematic 'snapshot' picture of temporal experience simply does not have the resources to accommodate these constraints, for to satisfy them requires that at least the content of an experience be temporally extended, in order to perceive both relata in a single experience. Furthermore, the question I have

been considering as to the relationship between the two temporal structures would be irrelevant if the snapshot view were true, since neither the content of the experience nor its temporal structure are taken to have any interesting kind of temporal structure on this picture. This observation is true also for the view of Dennett (1991) and Dennett and Kinsbourne (1992), which I have also discussed; if we wish to uphold that the phenomenological constraint is genuine, and furthermore that experiences are unified over time in the manner in which they seem to be, we should reject Dennett's anti-realist picture of temporal experience.

With the snapshot view ruled out, we should consider which of the remaining views follows most naturally from the considerations that have been advanced against mirroring. The answer is clearly the retentionalist view, which envisions experience as momentary or short, but nevertheless apprehending a temporal spread of content. On the retentionalist view, there is a significant divergence between the duration (if any) of the experience, and the duration apprehended in the experience; and since it is metrical mirroring (inheritance of both duration and order) that is the target of my critique, it follows that metrical mirroring cannot be the case if experience is retentional in this way. By contrast, the metrical mirroring view finds its home most naturally in an extensionalist framework, where the duration and ordering relations of both experience and object are taken to align. In Dainton's framework, that means that experiences decompose into further experiences as parts, which are unified over time by the relation of co-consciousness; however, this is by no means the only way of filling out the extensionalist thesis, which may allow of other kinds of unity of time.

There is much more to be said than this concerning the relative merits and criticisms of these two competing views in the philosophy of temporal consciousness; Dainton (2000, 2010a) gives a full treatment of the debate between these two rival views of experience. Dainton himself subscribes to extensionalism, which he argues is the only view that does fill justice to the 'naive' considerations advanced by Phillips (2014b) and others than we *directly* perceive temporal properties. The retentionalist view, by

contrast, will clearly have to rely upon some additional cognitive ingredients to be a viable account of experience; since for the retentionalist it is possible for an experience to be very short but its content extended, some ‘filling in’ on the part of the brain must occur in order to create the presentation of a temporal interval, to avoid the very odd conclusion that the content of an experience increasingly ‘lags behind’ the experience itself. The most natural cognitive addition to make here may involve memory, specifically ‘retentions’, which are taken to be a kind of backwards-looking representation that creates the impression of an extended period of time despite the fact that the experience itself is shorter (e.g. Husserl 1964). In order to sustain their account, the retentionalist must provide detail concerning the precise way in which the brain is capable of producing this impression, whilst still allowing for the phenomenological constraint to be upheld. This requirement to both do justice to how temporal phenomena seem to us, but also reject the mirroring of temporal structures, is I contend the central challenge for a successful retentionalist account. The motivation provided for retentionalism by the arguments made in this essay is therefore only half the story: the retentionalist must provide this extra detail in order for their view to be upheld as more plausible than extensionalism, all-things-considered. Since there is not time for this more detailed inquiry here, I merely intend for the conclusion of this essay to be that we should cast doubt on the notion of mirroring; whether retentionalism is in fact the correct view of experience over time is a further matter that requires further work, though anybody who is doubtful of the kind of mirroring proposed by Phillips (2010) and others may well be naturally drawn to a retentionalist view of experience.

Unity and overlap

Aside from these concerns, the final issue I wish to discuss is that of the unity of experience over time, with regard to the retentionalism vs extensionalism debate. In motivating my inquiry in chapter 1, I set out what I called the diachronic unity constraint (UC); the view that any account of temporal experience should allow for experiences to be unified over time. The stream of consciousness seems unified when

we attend to our own experience; our conscious experience admits of being broken down into individual experiences such as an experience of a blue object, or a high-pitched note, but these individual experiences seem to flow one-into-the-other. How exactly is it that experiences are unified over time such that we are given this seamlessness in experience? Dainton (2000, 2010a) claims that the extensionalist view of experience can account for this constraint better than the retentionalist approach, and presses a concern against the retentionalist which he calls the ‘problem of repeated contents’. I wish here to outline this objection but also partially defend the retentionalist on this count, arguing that the extensionalist too faces equally serious worries concerning their ability to provide a coherent account of unified experience over time. It is instructive to first consider the notion of experiential ‘overlap’ with regard to the question at hand. To use common terminology chiefly associated with mereology, if two entities overlap, they must share proper parts, or have parts in common; the notion of overlap is thus intimately connected to the notion of parthood (Varzi 2015, section 2.2). Thus, when we speak of two experiences, or two contents, overlapping, we need to be able to make sense of the idea that experiences and contents can have proper parts. To enter into a full investigation on how experiential parts may be individuated is beyond the scope of this essay; it will suffice to say here that we can coherently make sense of the notion of part being applied both to experiences and their contents. That is to say, we can plausibly talk of ‘parts of an experience’ or ‘parts of the content of an experience’ without requiring more information as to how those parts are individuated. Based on the notion of experiences or their contents possessing parts, the notion of *overlap* can be used both by the retentionalist and the extensionalist to explain the unity and continuity of experience over time: the *diachronic unity* of experience. For the extensionalist (e.g. Dainton 2000), experiences themselves share proper parts by overlapping, guaranteeing a seamless series of experiences in the conscious mind of the perceiver. For the retentionalist (e.g. Broad 1923), it is only the content of the experiences that overlaps, whilst the experiences themselves do not, guaranteeing nonetheless a seamless stream of experiential content in the mind of the perceiver.

For the retentionalist, and perhaps also he/she who is skeptical of mirroring these based on the considerations I have advanced in this essay, talk of temporal parts of an experience will be unwelcome, especially if those parts are taken to match in some way the parts of the scene which is being experienced (as Dainton 2000, 2010a advocates under the banner of ‘extensionalism’). This is because for the retentionalist, experiences themselves have no significant temporal structure or extent; they are ‘atomistic’ (Lee 2014a, 2014b). Since for the retentionalist, experiences are not extended in time over any significant interval, and do not decompose into smaller experiences that are their parts, experiential overlaps (if such an idea can be made sense of on retentionalism) will not guarantee the required continuity of experience. Therefore, if the retentionalist is to make use of the notion of overlap to guarantee experiential continuity, it must be overlap at the level of experiential content, rather than experience itself. Just such an account is found in Broad (1923), whose application of this notion is depicted in Dainton (2010a) as follows:

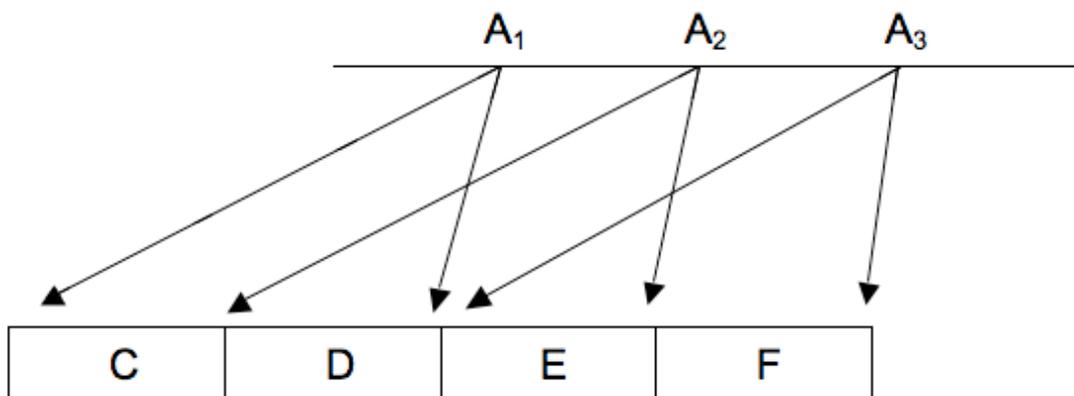


Diagram from Dainton (2010a)

The above diagram represents the retentionalist view supplemented with a notion of overlapping contents to create continuous experiencing by a subject over time. A_1 - A_3 represent individual experiences, (or ‘acts of awareness’ in Broad’s terminology), whilst C, D, E and F represent a sequence of entities experienced, for example musical tones. C and D fall within A_1 ’s ‘specious present’; they are the content of the subject’s experience A_1 . Furthermore, D and E fall within A_2 ’s specious present. To say that A_1

and A2 overlap in terms of their content is to say that the tone D is part of the content of A1 *and* A2; D is a proper part had by both contents. So too for E with regard to A2 and A3; the contents of these experiences overlap in virtue of both containing the tone E. In this way, the retentionalist (in this case, Broad 1923) is able to account for the diachronic unity and continuity of experience: individual experiences are unified over time to form a continuous stream of consciousness by possessing overlapping contents.

The chief concern levelled against such an attempt to satisfy the diachronic unity constraint on the part of the retentionalist has been made most forcefully by Dainton (2000, 2010a), who terms it the ‘problem of repeated contents’. Dainton argues that an implausible conclusion that follows from the above picture is that some contents will be repeated in experience; namely, the contents that figure in multiple acts of awareness; and this is manifestly at odds with everyday experience. Take, for example, the tone D, which is part of the content of A1 and A2; A1 and A2 ‘overlap’ by both containing D in their content. However, the fact that two successive experiences contain the note D in their content would seem to lead to that content being repeated in experience: thus, on Broad’s picture, the subject would experience C, D, D again, E, E again and then F. However, this is at odds with what the picture is meant to represent, which is an experience of C, D, E and F. That D and E appear to be repeated on Broad’s picture is a serious worry for Dainton, who claims that it shows that the retentionalist cannot appeal to overlap of contents and at the same time remain true to the phenomenological reality of experience. Dainton terms this problem ‘The problem of repeated contents’ (Dainton 2000).

I have no adequate reply on behalf of the retentionalist to the problem Dainton poses; but since it is not my aim in this essay to fully defend retentionalism as a theory of experience, I contend that this does not challenge my conclusion, which is that we have reason to doubt a mirroring constraint on temporal experience. However, since many of those who wish to deny mirroring in the way I have suggested might also wish to advocate something along the lines of retentionalism, it is worth saying something

instructive in response to Dainton’s argument. I argue here that Dainton’s solution to the problem on behalf of the extensionalist, also employing the notion of ‘overlap’, faces several worries as serious as the retentionalist’s problem of repeated contents. The notion Dainton proposes in order to avoid the problem of repeated contents is that of overlap at the level of experience itself. Since for the extensionalist, experiences *themselves* are extended in time, and furthermore mirror their objects in terms of duration and ordering, the extensionalist can allow that experiences themselves can overlap in order to ground continuity. The picture Dainton advocates on behalf of the extensionalist is as follows:

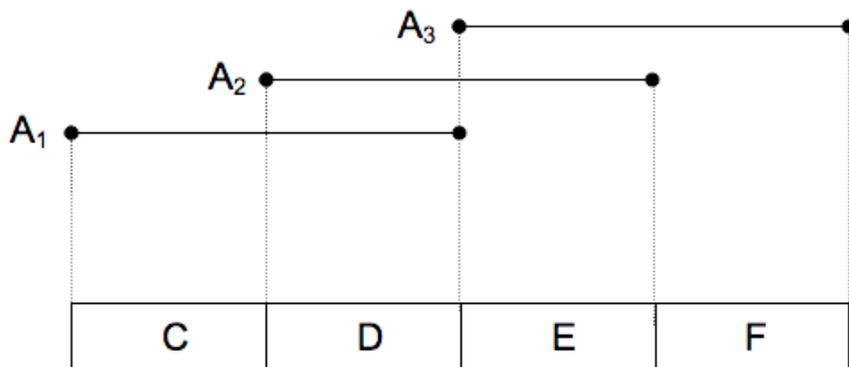


Diagram from Dainton (2010a)

On Dainton’s picture, the experiences (or ‘acts of awareness’) are themselves extended, and themselves overlap; there is overlap of both experiential content and the experiences themselves. Within one act of awareness A1, contents C and D are experienced together. The new picture is intended to avoid the problem of repeated contents as follows. The problem derived from having the very same content being reiterated in a successive experience. However, on Dainton’s picture, this reiteration does not occur: whilst D (for example) is part of two distinct experiences A1 and A2, these experiences themselves overlap, sharing a proper part which has D as its content. This shared proper part only occurs once, and so it follows that D is only experienced once, and not repeated. In this way, the phenomenological reality of experiencing C, D,

E and F once each is respected. As Dainton puts it, '*if A1 and A2 are discrete acts then D is experienced twice-over. But if A1 and A2 overlap in the manner indicated, D is experienced just once: by the episode of sensing [experiencing] that is common to A1 and A2.*' (Dainton 2010a, section 5.4).

It may be the case that Dainton's overlap model successfully avoids the problem of repeated contents, and in conceiving of experiences as extended phases rather than discrete acts gains an advantage over the rival theory. However, the notion of overlap at the level of experience itself raises some equally problematic questions. Both Strawson (2009) and Gallagher (2003) point out that the account does not do justice to our ordinary phenomenology when we reflect on our own experiences; we cannot discern the gaps between our experiences (an observation made previously in this essay in chapter 3), still less that our experiences overlap in the manner that Dainton suggests. But furthermore, and more seriously, Dainton's account faces another problem. In the diagram above, experiences A1, A2 and A3 overlap, which grounds their continuity in the stream of consciousness. However, it is possible to think of cases where what is experienced in succession is experienced as continuous, but *without* the experiential overlap; such cases may show that experiential overlap does not fully capture the continuity of conscious experience on Dainton's picture. One example, given in Bayne (2001), concerns a case where a subject enjoys *two* token 'D' experiences at a time (as opposed to one in the diagram above), D1 and D2. Suppose in this new case that D1 is experienced together with C, and that D2 is experienced together with E; that is, the act of awareness A1 has C and D1 as its contents, and the act of awareness A2 has D2 and E as its contents. For this case, Dainton's overlap model predicts that the subject in question will not experience D1 and D2 as continuous, as there is no experiential overlap between A1 and A2. However, as Bayne points out, it seems highly plausible that the subject will experience D1 and D2 as continuous, since she may well not be able to distinguish the two experiences on phenomenal introspection; the two token experiences occur at the same time (Bayne 2001 p. 88). This worry, and the worry that the proposed experiential overlaps are not given in our phenomenology as we might

expect, both give reason to doubt that experiential overlap is a viable model of experiential unity and continuity. Gallagher (2003) and others raise further worries against Dainton's view, some of which are countered in Dainton (2003), but which I shall not explore here.

Unity and continuity of experience over time

The question of overlap with regard to experience, and the broader question of how experiences are unified over time to produce continuity in conscious experience, both deserve a far more thorough treatment than I have given them here. The purpose of the above discussion was to highlight two key points that I take to be instructive to the debate that I have been considering in this essay. I have argued that we have reason to doubt that metrical mirroring is true with respect to experience and its objects. Those who reject mirroring may be drawn most naturally to a 'retentionalist' view of experience, in Dainton's terminology; conversely, those who advocate metrical mirroring may be drawn most naturally to an 'extensionalist' position (though neither stance on mirroring *entails* either retentionalism or extensionalism with regard to experience). I have tried to show that whilst the retentionalist faces a serious worry in the form of the 'problem of repeated contents', so too does the extensionalist with his notion of experiential overlap. We might generalize these worries for each theory to the conclusion that for any view of experience, a key question to be answered more adequately is the question of how individual experiences are unified over time to make up a continuous stream of consciousness that appears seamless to the perceiver. This is not the same question as the question I addressed in chapter 4, with regard to the 'homoeomery' of experience; the question is not whether the experiential part is prior to the whole or vice versa, but the relation individual experiences must stand in over time in order to account for the diachronic unity of conscious experience. Both the retentionalist and the extensionalist require a better theory of diachronic unity of experience. Just as much recent work in the philosophy of mind (for example, that of Bayne and Chalmers 2003) has focussed on the *synchronic* unity of conscious

experience (that is, unity of consciousness at a time), there is still much work to be done in providing an adequate account on the *diachronic* unity of conscious experience. Reflection on the retentionalism/extensionalism distinction and the shortcomings of existing attempts to provide such an account in terms of overlap indicates this area of the philosophy of mind and temporal consciousness that is undoubtedly ripe for future research.

Bibliography

- Bayne, T. 2001. 'Co-consciousness: review of Barry Dainton's 'Stream of Consciousness'', *Journal of Consciousness Studies* 8(3), pp.79-92.
- Bayne, T. and Chalmers, D. 2003. 'What is the unity of consciousness?' in Cleermans, A. (ed.), *The Unity of Consciousness: Binding, Integration, Dissociation*. Oxford University Press.
- Broad, C. 1923. *Scientific Thought*. New York: Harcourt, Brace and co.
- Chalmers, D. 1996. *The Conscious Mind*. Oxford University Press.
- Chalmers, D. 2005. 'The representational character of experience', reprinted in Chalmers, D. *The Character of Consciousness*. 2011. Oxford University Press.
- Chuard, P. 2012. 'Temporal experiences and their parts', *Philosophers' Imprint* 11(11), pp.1-28.
- Crane, T. 2011. 'The problem of perception' in Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*. Summer 2015.
- Dainton, B. 2000. *Stream of Consciousness*. Routledge.
- Dainton, B. 2003. 'Time in experience: a reply to Gallagher', *Psyche* 9(10).
- Dainton, B. 2010a. 'Temporal consciousness', in Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*. Spring 2014.
- Dainton, B. 2010b. 'Interpreting temporal illusions'. Supplementary document to 'Temporal consciousness' in Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*. Spring 2014.
- Dennett, D. 1991. *Consciousness Explained*. Penguin.

- Dennett, D. and Kinsbourne, M. 1992. 'Time and the observer: the where and when of consciousness in the brain.', *Behavioural and Brain Sciences* 15(2), pp.183-201.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge: MIT.
- Eriksen, C. and Hoffman, J. 1972. 'Temporal and spatial characteristics of select encoding from visual displays', *Perception and Psychophysics* 12, pp.201-204.
- Foster, J. 1979. 'In *self-defence*' in Macdonald, G. (ed.), *Perception and Identity*. London: Macmillan.
- Frischhut, A. 2014. 'Diachronic unity and temporal transparency', *Journal of Consciousness Studies* 21(7-8), pp.34-55.
- Gallagher, S. 2003. 'Sync-ing the stream of consciousness', *Psyche* 9 (10).
- Geldard, F. and Sherrick, C. 1972. 'The cutaneous "rabbit": a perceptual illusion', *Science* 178, pp.178-9.
- Gill, K. 1993. 'On the metaphysical distinction between processes and events', *Canadian Journal of Philosophy* 23(3), pp.365-384.
- Grush, R. 2007. 'Time and experience' in Muller, T. (ed.), *Philosophie der Zeit*. Frankfurt.
- Hacker, P. 2004. 'Substance: things and stuffs', *Aristotelian Society Supplementary Volume* 78(1), pp.41-63.
- Harman, G. 1999. *Reasoning, Meaning and Mind*. Oxford University Press.
- Hoffman, J.E. 1998. 'Visual attention and eye movements' in Pashler, H. (ed.) *Attention*. London: University College London Press.
- Hurley, S. 1998. 'Vehicles, contents, conceptual structure and externalism', *Analysis* 58(1) pp.1-6.

- Husserl, E. 1964. *The Phenomenology of Internal Time-Consciousness*, (Churchill, J. tr.). The Hague: Marinus Hijhoff.
- Jackson, F. 1977. *Perception: a Representative Theory*. Cambridge University Press.
- James, W. 1890. *The Principles of Psychology* (vol.1). Dover Publications.
- Kelly, S. 2005. 'The puzzle of temporal experience', in Brooks, A. (ed.), *Cognition and the Brain: The Philosophy and Neuroscience Movement*. Cambridge University Press.
- Kim, J. 1992. 'The non-reductivist's troubles with mental causation' in Heil, J. and Mele, A. (eds.) *Mental Causation*. Oxford University Press.
- Lee, G. 2014a. 'Temporal experience and the temporal structure of experience', *Philosophers' Imprint* 14(3), pp.1-21.
- Lee, G. 2014b. 'Experiences and their parts', in Bennett, D. and Hill, C. (eds), *Sensory Integration and the Unity of Consciousness*. MIT Press.
- Lee, G. 2014c. 'Extensionalism, atomism and continuity' in Oaklander, L.N (ed.) *Debates in the Metaphysics of Time*. Bloomsbury.
- Martin, M. 2002. 'The transparency of experience', *Mind and Language* 4(4) pp. 376-425.
- Mourelatos, A. 1978. 'Events, processes and states', *Linguistics and Philosophy* 2 pp.415-434.
- Phillips, I. 2009. *Experience and Time*. Thesis (Ph.D.). University College London.
- Phillips, I. 2010. 'Perceiving temporal properties', *European Journal of Philosophy* 18(2) pp.176 - 202.

- Phillips, I. 2014a. 'Experience of and in time', *Philosophy Compass* 9(2) pp.131-144.
- Phillips, I. 2014b. 'The temporal structure of experience', in Lloyd, D. and Arstila, V. (eds.) *Subjective Time: the Philosophy, Psychology, and Neuroscience of Temporality*. MIT Press.
- Russell, B. 1913. 'On the experience of time', in *The Collected Papers of Bertrand Russell* (vol. 7). London: Allen and Unwin.
- Russell, B. 1921. *The Analysis of Mind*. London: Allen and Unwin.
- Soteriou, M. 2013. *The Mind's Construction*. Oxford University Press.
- Strawson, G. 2009. *Selves*. Oxford University Press.
- Taylor, B. 1985. *Modes of Occurrence*. Oxford (Blackwell).
- Tye, M. 2003. *Consciousness and Persons*. Cambridge: MIT.
- Varzi, A. 2015. 'Mereology' in Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*. Winter 2015.