EXPERIENCE AND TIME

IAN B. PHILLIPS

DEPT. OF PHILOSOPHY

UNIVERSITY COLLEGE LONDON

2009

DECLARATION

I, Ian B. Phillips, confirm that the work presented in this thesis is my own. Where material has been derived from other sources I confirm that this has been indicated in the thesis.

I.B.P.

January, 2009

ABSTRACT

We are no less directly acquainted with the temporal structure of the world than with its spatial structure. We hear one word succeeding another; feel two taps as simultaneous; or see the glow of a firework persisting, before it finally fizzles and fades. However, time is special, for we not only experience temporal properties; experience itself is structured in time.

Part One articulates a natural framework for thinking about experience *in* time. I claim (i) that experience in its experiential aspect has a realistically conceived temporal structure; (ii) that our judgements about that structure always go via judgements about the temporal structure of the apparent objects of perception; and (iii) that a subject undergoing perceptual experience of a given experiential kind is always in a position to know that they are undergoing experience of that kind simply in virtue of so undergoing. On this basis, I argue that the temporal structure of experience cannot systematically come apart from the temporal structure of its objects.

Part Two treats four puzzles relating to our experience *of* time. The first is Dennett's notorious discussion of masking and apparent motion phenomena. The second is the traditional debate regarding the very possibility of perceiving temporal properties. The third is Fara's recent contention that standard explanations of our experience of slow changes preclude us from perceiving constant motion. A common reaction to these three puzzles is to reject some element of the naïve picture of temporal experience developed in Part One. I resolve them instead by showing how each arises from mistakenly thinking that experience is *homoeomerous* down to very short durations or instants. That is, thinking that we can analyse experience into a series of independent short slices, and explain the nature of the stream of consciousness in terms of those slices. The final chapter discusses a fourth puzzle about visual motion perception which I diagnose as driven by a rather different, but equally misguided way of thinking about vision.

CONTENTS

PART I – EXPERIENCE IN TIME

Chapter One:

The Temporal Structure of Experience

> *Time is the thing I am made of. Time is a river that sweeps me along, but I am the river; it is a tiger that tears me apart, but I am the tiger; it is a fire that consumes me, but I am the fire. The world, unfortunately, is real; I, unfortunately, am Borges.*[1]

1. The Datum

Myriad different temporal properties and relations can be made manifest in perceptual awareness. We experience objects enduring, and their doings and modifications as sequential, simultaneous, and persistent. We can hear one word succeeding another; feel two taps as simultaneous; or see the glow of a firework persisting, before it finally fizzles and fades.[2] Our experience is also replete with aspects which, whilst not strictly temporal, bear a logical connection to time. We can feel a raindrop run down our cheek; watch a ball trace its parabola into our cupped hands; or hear the crescendo of the horn section as the symphony reaches its climax. Since motion and change in general necessarily take place over time, I use the phrase 'temporal properties' so as to include these aspects of our experience in addition to the strictly temporal aspects already mentioned.

---

[1] Borges (1964) 'A New Refutation of Time' from *Labyrinths: Selected Stories and Other Writings*.

[2] Cf. Block and Zakay who note, "There are several qualitatively different kinds of temporal experiences: simultaneity, successiveness, temporal order, duration, and temporal perspective" (2001: 59). My primary interest is in the first three kinds of experience, and also our experience of relative duration – see §1.1. Some argue that we also experience 'A-properties' (presentness, pastness etc.) or even the 'flow' or 'passage' of time itself. I do not discuss these possibilities.

Just as with properties such as shape, colour, and pitch, when it comes to describing what it is like for us experientially, we naturally appeal to temporal properties in order to give an adequate account of our conscious lives. Strawson famously pointed out that, when asked to describe their experience, naïve subjects do so in terms of ordinary, apparent objects of perception and their features. Thus, his imagined subject reports his visual experience by saying that he sees "the red light of the setting sun filtering through the black and thickly clustered branches of the elms … [and] the dappled deer grazing in groups on the vivid green grass…" (1979: 43). Such descriptions already refer to processes which essentially take place over time, in particular *grazing*. But in any case, Strawson's naïve subject might equally have answered that he could see the deer gambolling in the dusk-lit grass. He might have noted how his attention was drawn to the branches of the elms swaying in the wind and casting a play of successive shadow patterns on the ground. From this perspective, we seem to be no less directly acquainted with the temporal structure of the world around us than with its spatial structure. As John Foster puts it,

> duration and change through time seem to be presented to us with the same phenomenal immediacy as homogeneity and variation of colour through space. (1982: 255)[3]

Part Two of this thesis considers ways in which philosophers have struggled to account for these two aspects of our perceptual experience: both the strictly temporal aspects of experience (e.g., perceived succession, persistence and simultaneity) and those aspects, like perceived change and motion, which whilst not strictly temporal bear a logical connection to time. There I try to understand and respond to the difficulties that philosophers have had in accounting for temporal experience, difficulties felt to be so large that some have even denied that experience really does make manifest temporal properties. For now I take our experience of such aspects as a datum in order to explore issues to do with the temporal structure of our *inner* lives and the relation between our inner experience and its apparent objects. This may seem back-to-front. However, as will become evident in Part Two, one can only achieve clarity regarding the answers to questions concerning how our experience can present or represent the world as it unfolds in time, if one is first clear about how experience itself is structured in time. The

---

[3] See also Dainton 2000: 114-5, Broad 1923: 287, 351, O'Shaughnessy 2000: Ch.1, §3 and many other places; though contrast Le Poidevin 2007: Ch.6.

fact that experience itself has temporal structure imposes constraints on the way experience can make manifest the temporal structure of the world; we must understand experience *in* time before we can understand experience *of* time.

Throughout, my primary concern will be experience in the narrow sense of perceptual experience. Other kinds of mental episode (imagining, recollecting, thinking and so forth) have temporal shape and are no less a part of the stream of consciousness than perceptual experience. There is, for example, something it is like to undergo an imaginative episode and its temporal shape contributes to that being so. As such, many of my claims will in fact apply to such episodes.[4] Nonetheless, the problems of time and experience arise most forcefully in the perceptual case and that case will be my burden.

1.1 Metrical Properties

Amongst the temporal properties we should distinguish the *metrical* properties, in particular the metrical durations of events and processes, and also the metrical rates of change of events and processes. The examples of temporal perception above leave it open whether we can, for example, experience a movement from A to B as taking two seconds, or a thunder clap as lasting three; and likewise whether we can hear the period of a siren as being 500ms or see a sports car as accelerating at seven metres per second per second. Certainly we can perceive *that* such things are the case – we can be positioned to know these facts in virtue of undergoing perceptual experience. But should we think of such properties as being genuine aspects of the content of our experience? As being, strictly speaking, perceptual? There are good reasons to think not.

Before motivating that claim, a caveat. Some theorists are inclined to think of vision as possessed of one kind of content which determines phenomenal character and a further content which is independent of phenomenology.[5] My focus here is purely on the phenomenological. Nothing I argue conflicts with the view that, in some non-

[4] Especially those in Chapter Three – though see discussion there also. The notion of transparency developed in Chapter Two highlights a clear contrast between perceptual experience and these other episode kinds.

[5] Some will want to think of this as the narrow/wide distinction but the distinction can be drawn independently of one's view of externalism. For helpful discussion of the various views one might take, at least on a representationalist approach, see Chalmers 2004.

phenomenological sense, experience has content with metrical properties. I say more about the phenomenological in Chapter Three. In terms there elucidated, my claim is that metrical properties are not amongst the *experiential* properties of experience.

Why deny that metrical temporal properties are amongst the experiential? Firstly, consider the spatial analogy. Can we experience objects as being a metre long? Or as growing by a centimetre a second? The cartoon below is evidence that the intuitive answer is, 'No.'

If our cartoon-Crusoe's vision actually *presented* his island as barely a few metres across, or the miniscule palm tree as only centimetres high, the pretence engaged in would be impossible to relate to – or at least not sufficiently plausible for us to find the joke amusing. One could hardly pretend to be a giant if you saw the things around you as only centimetres large. Rather, the pretence takes place against the tacit background understanding that vision is neutral as regards the metrical size of the island.

Pretence would also make sense if vision presented the island illusorily. But the idea that our cartoon-Crusoe has been enjoying visual *illusions* of the island as many metres across, is not how we ordinarily react to the conceit. We regard visual illusions as robust in the face of countervailing beliefs. Yet we take the pretence to be extremely tenuous and

easily abandoned.[6] In other words, we take it that the man-as-giant pretence is not grounded in experience; rather, experience leaves the pretence open to Crusoe. As the novelty wears off, the island continues to look, as it always has done. How is that? My suggestion is that the island continues to look small *relative to his* [Crusoe's] *body-size*. In this way, Crusoe's visual experience is neutral as between giant-pretence and small island-reality.

Similarly, consider how Lewis Carroll relates the Wonderland story after Alice has begun to shrink again. "[Alice] looked down at her hands, and was surprised to see that she had put on one of the Rabbit's little white kid gloves while she was talking. 'How *can* I have done that?' she thought. 'I must be growing small again.' She got up and went to the table to measure herself by it, and found that, as nearly as she could guess, she was now about two feet high" (2001: Ch.2). Again, it is crucial to the narrative here that Alice does not perceive metrical properties. If perception presented or apparently presented her with metrical size-properties, then Alice's experience (whether or not veridical) would not be neutral as to her size, and so she would immediately see herself as small again. Instead, Alice seeks comparison with the table to establish that she is small again (as opposed to the glove having grown). The table is precisely used to gain metrical knowledge, since Alice (correctly) takes it that the table has remained constant in size. Alice's behaviour seems quite rational to us and is best explained by the fact that she, like us, does not perceive metrical spatial properties. The table is not being used to rule on whether her experience is veridical or not. That could be its only role if her hand was apparently presented as only a few centimetres in size.

I now spell out the argument against metrical spatial properties figuring in the content of visual experience more rigorously. The argument focuses on two cases.

> Case 1: Alice, who has led an entirely normal life until the current time, awakes from a dreamless sleep. Everything in her visible environment, except for her own body, is a quarter of its normal size.[7]

---

[6] Likewise, we do not think that entertaining this kind of pretence (or indeed convincing ourselves of a false belief) can alter the nature of our visual experience, dramatically shifting its content.
[7] The shrinking must be limited to avoid metaphysical issues concerning absolute and relative space.

> Case 2: Alice, who has led an entirely normal life until the current time, awakes from a dreamless sleep. She is four-times her normal size. Nothing else has changed in size.

Clearly in both cases, the sizes of visible things relative to Alice's own size have changed. However, the question is whether visually speaking it looks to Alice as if she is in one case or the other. Firstly, I claim that Alice cannot distinguish the cases *on the basis of her sensory experience*, specifically visual experience. They are visually indiscriminable from Alice's point of view.[8] If we think that phenomenal aspects of experience make themselves manifest (a claim I defend in Chapter Three), the experiences cannot present *different* properties and yet be indiscriminable. Thus, we have two basic options. Either vision presents metrical properties to Alice veridically in at most one of the cases. Or vision is neutral and does not present metrical properties at all.

If vision is not neutral, at most one case is veridical. The question is, which? Of course, the answer might be neither. But this simply raises the further question, what are the veridicality conditions of Alice's experience in both cases? These questions raise a clear challenge to someone who thinks that we perceive metrical properties: provide some basis on which to prefer one metrical content as opposed to another. Conversely, the denial that we perceive metrical properties allows one to offer the following simple and attractive account of the situation. In both cases, Alice perceives the world as being much smaller than her own body size, she is aware of the relative change in size of things. However, her visual experience is quite neutral as to what has actually happened; vision does not speak to the issue of which situation she is in.

We can put the point another way. One way of introducing veridicality conditions into experience is as follows.[9] At least with respect to those aspects of experience which are *transparent* (see Chapter Two), a request to describe our experience will be answered by describing the world apparently encountered in experience. This description of the world

[8] At least in terms of size. I leave aside various distracting complications. For example, if Alice's powers of discrimination remained the same, at least on one way of understanding discriminative powers, then she would be able to discriminate more detail on a larger hand and so there would be certain differences in how things looked to her.

[9] This way of putting things attempts to remain neutral as to whether we think of experience as having propositional content with veridicality conditions or alternatively we think of experience in relational, or naïve realist terms, but nonetheless wish to allow for veridicality conditions. For a challenge, at the very least to the propositional content approach here, see Travis 2004.

as apparently encountered provides a source of veridicality conditions for we can say that the experience is veridical to the extent that the description applies to the world *actually* encountered.[10] Now, if we consider Alice's ability to make judgements about the world encountered on the basis of her experience, she is in no position to describe the metrical properties of the world. She cannot judge which case she is in. As a result, employing the above method of introducing veridicality conditions, these will at most include the condition that the visible world be a fourth of its usual size relative to her body. In other words, both experiences will be veridical.[11]

One worry about this argument is that it is too powerful. We certainly allow that we can see things as being pea-sized or golf-ball-sized. Does this not amount to seeing things as having metrical properties, where the unit is a pea or football? If we deny this, what sizes *can* we see things as having? I suggest quite generally that we see things as having *relative* sizes, in particular that we see things as having sizes relative to our own bodies.[12] Seeing things as golf-ball-sized simply means that we see things as having the size-relative-to-us that a 'normal' pea or golf-ball does. Similarly, we can allow that things have the size-relative-to-us that a 'normal' metre-rule has. Now consider the two cases above. In both we can say that things look smaller-relative-to-Alice. For her, golf-balls look pea-sized, where this means the size that peas normally have relative to her. And, likewise, there is a sense in which things look to be much smaller than a metre, viz., they look much smaller than a metre-long object normally does to her. But this is not to say that things look a metre long *simpliciter*.

Clearly the aforementioned aspects of vision depend on the relative sizes things have during a 'normal' environmental upbringing. So one might argue that in a normal environment we can see things as a metre long after all. Now this certainly is a plausible line to take in certain cases. For example, imagine that I have grow up in a 'normal' environment often encountering oranges in the way we ordinarily do. Consider also the

---

[10] Of course, I assume an ideally rational and attentive describer.

[11] Once again this argument assumes that we have access to the veridicality conditions of our experience; again see Chapter Three below.

[12] I take it that giving the precise details of the relativity here is an empirical question. I also assume that we only have a grip on the relative size of our own bodies. As Pickard puts it, "…what is hard to see is how [one's] awareness of one's body size [relative to other bodies] could be made available from the inside alone, and not in conjunction with one's awareness, through the outer senses, of the sizes of these other things. Whatever awareness of one's size amounts to, it seems to require an external measure" (2001: 27). Thus, in my view, all size perception is relative.

schmorange, a fruit that has precisely the look of a normal orange but, being the creation of a hare-brained geneticist, is in fact not of that natural kind. If I do have a chance encounter with a schmorange, it is not wholly outlandish to suggest that I enjoy a visual illusion of the schmorange as an orange. Given the context and my environmental upbringing, possession of an orange-look might be thought sufficient for appearing to me as an orange. The fruit may also have the look of a schmorange, but I have not been brought up around schmoranges.

Such a story is not plausible in the metrical case. For it provides no answer to the initial challenge: what are the veridicality conditions of Alice's experience? Which case, if either, is her experience veridical in? Moreover, whereas when I am presented with an schmorange, I naturally describe my experience in terms of *oranges* (hence the grounds for positing illusion), Alice does not naturally describe her experience in terms of things being metrically different sizes. It is not as if a table that she would have described on the basis of vision as being two metres long, she now describes as being half a metre long. Rather her description is neutral. In sum then, the most natural view to take of Alice's predicament is that her visual experience is neutral across cases and does not involve illusions. More generally, the most natural view to take is that our experience presents us with the relative sizes of objects and not their metrical sizes.

Turning to the temporal case, we can perform a similar thought experiment to probe whether the temporal features of our experience include metrical properties.[13] Take a subject, Rose, and record her neurological activity over an hour. Then replicate the same pattern of activity in her brain but speeded up so that it is all compressed into the space of a few seconds.[14] This will plausibly induce a massive hallucination whose apparent objects will, I presume, match the objects seen in the recorded hour at least in structural properties.[15] Now, if we genuinely perceive temporal metrical properties in experience, this kind of case generates a question: how long do the events in the hallucinatory experience appear to last to Rose? Do they seem to last an hour (or more precisely: to

---

[13] Conversations with Geoff Lee inspired this discussion. In forthcoming work, Lee uses reflection on experience consequent on ingesting a drug which speeds up neurological metabolism to argue against transparency claims (such as those in Chapter Two below). My own view is that transparency claims are immune to such objections once we recognise that metrical temporal properties do not feature in the content of experience as I argue here. Again, see also Chapter Five, §1.

[14] Not that this is nomologically feasible but it is, I take it, metaphysically or conceptually possible.

[15] The use of hallucination is inessential to the example but helps simplify discussion.

have just those durations which were apparent in the original experience) or do they appear to unfold over merely a few seconds (or more precisely: to have vastly briefer durations than those apparent in the original experience)?

There is a complication in thinking about how to answer this question. In trying to imagine Rose's experience, we cannot ignore the temporal structure of Rose's non-perceptual inner life, her stream of thought in the broadest sense, in contrast to her stream of perceptual experience. In the story as told above, this too will, I presume, be speeded up. In other words, Rose will think very much faster than normal, her stream of imagery will go past very much faster and so forth. Given that this is so, it is hard to see how Rose could possibly discover 'from the inside' that she was not simply 'reliving' her prior experience. That is, it is extremely hard to see how Rose could, by reflection on her inner life alone, distinguish her new situation from one in which everything in the world apparently perceived went on at the same pace and over the same period of time as ever. To know that something was radically different would require some kind of 'meta-time' against which to measure her current experience and thought in order to determine whether it lasted the same amount of time as her past experience. I see no reason to believe in such a 'meta-time'.[16] The default and by far the simplest hypothesis should be that no such time exists.

To clarify the situation we should distinguish two possible kinds of metrical property: subjective and objective. Objective metrical properties include the objective amount of change per unit time. Thus, in the above scenario Rose's inner life changes objectively much faster than normal. Rose, however, is quite blind to this. Subjective metrical properties include the rates of change that things seem to have. Insofar as these properties exist, I suggest that they do not change in Rose's case: everything in the world appears to her to go on at the same rate as ever. Objective metrical change isn't something that Rose notices (or can notice) in the envisaged scenario. But the question remains open whether we should think of her experience as involving subjective metrical properties.

[16] Cf. Dainton (2001: 21-3) in his discussion of meta-time in the context of McTaggart's argument.

If experiential content *does* include subjective temporal metrical properties, the fact that Rose will not be able to distinguish her new situation suggests – assuming the principle defended in Chapter Three that phenomenal aspects of experience are manifest – that these are the same in both cases; in both things seem to take, say, an hour.

One reason to hesitate here is that Rose's neural activity seems consistent with an entirely veridical experience in a creature with a very fast brain, living in a very fast part of the world. It seems deeply implausible to suggest that this creature would always be subject to illusions of temporal metrical properties. However, this consideration is not decisive because of the role that might plausibly be assigned to the 'normal' environment in determining content ascriptions. Nonetheless, there are reasons to think that an alterative account is more plausible, reasons which mirror the Alice cases above.

Assuming that our stream of perceptual experience and our stream of thought can be manipulated separately, we can imagine that Rose's internal mental activity – her thinking, imagining, deployment of attention and so forth – is dramatically slowed down whilst leaving her perceptual experience to flow along its normal time-course. How would things seem to Rose? From the inside it is hard to see how Rose could distinguish her situation from one in which the experienced world had suddenly speeded up, her thinking failing to keep pace. Again, to distinguish scenarios would seem to require a 'meta-time' with respect to which the competing hypotheses could be adjudicated, and it is extremely difficult to believe that we have any grasp of such a meta-time. If that is right, then the most plausible thing to say is simply that Rose feels as if her thinking were immensely sluggish *relative* to the experienced world. In contrast, if subjective metrical properties were presented in experience, it would have to be the case that things did feel one way or the other to Rose (either that her thinking had become incredibly sluggish or that the world had speeded up). But again, there seems no reason to think that this should be so given the intuitive indiscriminability of the scenarios.

If this is right the temporal content of experience is relative, on analogy with spatial content. Temporal content is relative to the flow of 'internal' conscious processes; slow down her 'internal' conscious processes without affecting her perceptual experience and Rose will have the impression that the surrounding world had speeded up in relation to that internal stream.

A view one sometimes encountered is that we do, in fact, have a sense of a more fundamental 'personal' time against which *all* inner processes can be measured.[17] On such a view it would make sense to think of our entire stream of consciousness slowing down or speeding up relative to this more fundamental flow. I am sceptical. Of course, the rate at which a certain kind of image or thought train recurs or proceeds may increase in speed or slow down relative to other mental goings on, just as with perceptual experience. But if we try to think through the idea that *every* aspect of our conscious life was slowed down in objective time, it is hard to gain any grip on how we could be aware of that change from the inside. If all such scenarios are mutually indiscriminable, that is powerful reason for thinking that *phenomenologically* they are all alike and that there is no further aspect to our mental life which involves a sense of this 'absolute flow'. Consequently, it seems to me that the burden of proof is squarely on the shoulders of the proponent of a fundamental personal time to evidence its phenomenal reality.

In a recent article, Dainton comments as follows.

> Since we lack clocks for directly measuring phenomenal durations, we have no option but to correlate them with objective processes whose durations can be measured using normal clocks. This is less than ideal, but it does give us a reasonably reliable handle on subjective durations: we all know roughly the quantity of auditory change which can occur during one second of objective time. There are of course complications, e.g. when in a hurry, a few (objective) minutes spent waiting in a queue can seem to take (subjective) eternity" (2008: 19, fn.9).

The first claim here strikes me as quite correct and, I suggest, further highlights the difficulties with the notion of a fundamental personal time. On the other hand, I suggest that once metrical properties are excluded from the content of experience, the every day phenomena that we talk about as involving 'time slowing down' or 'taking an eternity' take on a quite different guise. I return to these complications at the start of Chapter Five (§1). For now I simply want to flag that our starting datum should be understood narrowly as the claim that we are aware of certain temporal properties in experience –

[17] Barry Dainton suggested to me that we might think of Husserl's notion of absolute time-consciousness (and correlatively, his notion of absolute flow) in these terms. Forthcoming work by Geoff Lee considers such a notion independently of Husserl's view.

succession, simultaneity, and relative, non-metrical persistence – and need not involve the controversial idea that metrical properties are so presented. When I talk of (perceived) temporal properties in what follows I mean to exclude metrical properties.

2. Only a Matter of Time

In order to begin our investigation of the temporal structure of experience itself, consider the following passage from Carnap's *Aufbau*.

> The psychological objects have in common with the physical ones that they can be temporally determined. In other respects, a sharp distinction must be drawn between the two types. A psychological object does not have colour or any other sensory quality and furthermore, no spatial determination. (1967: §18, 33)

By 'objects' Carnap simply means "something about which a statement can be made"; psychological objects, we are told, "begin with, the acts of consciousness: perceptions, representations, feelings, thoughts, acts of will and so on" (ibid.: 32). Thus, Carnap puts forward two claims: (i) that temporal properties are common to psychological and physical objects (in Carnap's thin sense); and (ii) that time is unique in this respect: no other interesting properties are common to psychological and physical objects.

Carnap's first claim here seems clearly true. As we have just seen temporal properties form an utterly pervasive aspect of our experience of the world; these temporal properties are properties of physical objects and events – the raindrop's motion, the relative duration of the flash of light. Thus it is not just that physical objects have temporal properties; amongst the objects of perception are found temporally determined physical objects. The second conjunct of the claim, viz., that psychological objects also have temporal properties, seems equally unobjectionable. Our mental lives change and develop over time. We fall in and out of love. We hope to increase our store of knowledge even whilst we forget much that we once knew. Sharp, irregular stabbing pains abate, leaving dull throbbing ones in their stead. Last but not least, we can chart the course of our stream of consciousness over time and describe its flow. Experiences are not just *of* unfolding events and processes, they are in addition *themselves* dynamic in

character. More precisely, experience is itself composed of events and/or processes which persist through time and occur before and after one another.

What though about Carnap's second claim concerning the uniqueness of time in this respect?[18] At first blush the assertion seems far more tendentious. If certain identity theories are true, experience possesses whatever physical properties the neuro-chemical events or processes with which it is identified possess.[19] Thus if Carnap's claim is read transparently, it follows by Leibniz's law that psychological objects also have such properties. If identity theorists are right, is Carnap straightforwardly wrong? Is there nothing to be said for a Carnapian uniqueness claim? In the rest of this chapter I argue that much remains to be said in favour of uniqueness. In particular, I argue that temporal properties are experiential or phenomenal properties of experience. Thus, even if the identity theorist is right, temporal properties are plausibly the only interesting *experiential* properties of experience which are also properties of the objects of experience. Time *is* uniquely common to experience *in its experiential aspect* and the objects of experience.

3. Temporal Properties as Experiential Properties

Temporal properties contrast with other perceptible properties such as shape, colour and pitch in being *both experiential and* instantiated by experience itself. This is why time is distinctive. Insofar as the truth of identity theories is an open question, so is the attribution of spatial properties to experience itself. However, even according to such a theory, the locations of experiences are not experiential properties. In contrast, experiences are episodes in our stream of consciousness – our own subjective, experiential lives unfold as we perceive the world unfolding around us. That is to say, experience is not just composed of events and processes occurring in and persisting through time. It is composed of *experiential* events and processes occurring in and

---

[18] A claim also made by Mellor who draws attention to "the striking fact … namely that perceptions of temporal order need temporally ordered perceptions. No other property or relation has to be thus embodied in perceptions of it: perceptions of shape and colour, for example, need not themselves be correspondingly shaped or coloured" (1981: 8).

[19] At the very least such events will have physical locations. What other properties the identity theorist will want to ascribe to neural *events* or *processes* is a nice question. There may also be reasons to think that experiences have locations independent of identity theories. See, for example, Russell (1927), Lockwood (1989) and Lee (2007), who draw on special relativity to argue that any event located in time must also be located in space.

persisting through time. The time in which experience is structured is experiential time; its temporal structure *matters* from the point of view of its counting as experience.

3.1 Experiential Properties

What exactly does it amount to when we claim that a property is an *experiential* property? Experiential or phenomenal properties (I use the terms interchangeably) are those properties of an episode in virtue of which there is something it is like to undergo that episode from the subject's point of view. Thus, a property is experiential just if it is a property of an episode which contributes to what it is like, subjectively, for the subject of the episode to be undergoing that episode.[20] The phenomenal or experiential character of some course of experience is then the totality of that course of experience's phenomenal properties.[21] A consequence of this is that experiential properties will figure essentially in providing a complete personal-level psychological description of a person's conscious life.[22]

Some think this is as far as we can go.[23] However, even if the notion is not analysable in other terms, there are clearly informative things that can be said about phenomenal properties. For example, many hold that certain phenomenal properties are constituted by the presentation or representation of certain kinds of objects and their properties (for more on this see Chapter Two). Can we say anything about experiential properties in

---

[20] As Peacocke puts it following Nagel 1974, "Perceptual experiences and sensations, on the one hand, and so-called occurrent conscious propositional attitudes, on the other, differ in many respects. But there is one property they share. They both contribute to what, subjectively, it is like for the person who enjoys them" (2003: 83). Obviously we will want to distinguish between what it is like for a subject *when* they have a perceptual experience – something which might be taken to include the whole range of background phenomenology and attendant feelings which are occurring at the time of that event – and what it is like *to* undergo said experience. My focus is on the latter notion.

[21] This should not be read as committing to any kind of atomism about phenomenal character.

[22] Note in this regard that 'being an experience of an *x*' does not straightforwardly count as an experiential property. One can have an experience of a sparrow-hawk, say, but at such a distance and flying at such a speed that one perceptually discriminates only a dark blur high in the sky. In such a case, in contrast to simply seeing a sparrow-hawk in the aviary, one does not see the sparrow-hawk *as such*. Only the as-such properties will be experiential. Many philosophers have claimed that, strictly speaking, we never see sparrow-hawks as such since natural kind properties are not perceptual properties. See McGinn 1982, Tye 1995, Clark 2000, Foster 2000, Millar 2000, Burge 2003 and Price 2005. Contrast Siewert 1998 and Siegel 2006. The considerations for and against this view are complex and would take us far beyond the scope of this thesis. All I claim here is that temporal properties are frequently very obviously a part of the content of experience. The last thing we should do is deny that.

[23] Perhaps this is one thing that Block is suggesting with his notorious borrowing from Louis Armstrong: "If you got to ask, you ain't never gonna get to know" (1980: 281).

general? The most obvious connection to draw upon is their connection to self-knowledge. For it seems that the instantiation of phenomenal properties is something that is at least typically available for privileged self-conscious reflection. I explore this idea in depth in Chapter Three where I defend the idea that (roughly speaking) experiential properties are, in fact, *just* those properties of an episode that we are positioned to know about solely in virtue of undergoing said episode. For now, the idea of an intimate connection between phenomenal properties and availability is sufficient to allow us to motivate the idea that temporal properties of experience are experiential.

3.2 The Accessibility of Temporal Structure

Let us say that a property of experience is available for introspective self-reflection if and only if we can be in a position to know that our experience has such a property just in virtue of undergoing the experience. If a property is available in this way, then this is at least strong evidence that it is an experiential property in the relevant sense.

Consider supposed neural properties of experience. Introspection offers us no inkling of the neural goings on within our skulls (unless we are actually peering into our own skulls whilst reflecting on our experience). We have no inclination, for example, to think that the spatial location of experience (if any) is accessible to introspective reflection just in virtue of undergoing experience. In contrast, it is overwhelmingly natural to hold that at least certain aspects of the temporal structure of experience (its ordering for example) *are* accessible to reflection on one's experience.[24] And regardless of whether accessibility is a necessary condition for being experiential, accessibility does seem indicative of being experiential.

Why do I say that it is overwhelmingly natural to insist on the accessibility of some aspects of temporal structure? For two reasons. Firstly, we can raise questions concerning the temporal structure of experience *itself* as opposed to the temporal structure of the *objects* of experience. Such questions are in good conceptual order and have answers which are objectively true or false depending on the facts. Secondly, we can

[24] Cf. Lee who contrasting time and space writes, "Your experience might make you aware of spatial relations between external items, but experiences don't seem subjectively to be themselves spatially structured" (2007: 372, fn.9).

answer these questions knowledgably merely by reflecting on our experience, i.e., just in virtue of having undergone a certain course of experience. Jointly, these commitments evidence the accessibility and so experiential nature of experience's own temporal properties.

One common way in which we raise questions concerning the temporal properties of experience in contrast to the temporal properties of the object experienced is in cases of partial viewing and hearing. For example, Ruby might, in many contexts, respond truly to the question, 'Did you see the opera last night?' even if she walked out before the final act. Likewise, in many contexts, Bertie will count as having seen the football match even if he dozed off and missed various periods of play. Thus, even when we are fully informed concerning the temporal features of the event witnessed, we can still learn new information by asking about the temporal features of the witnessing: 'How long did Ruby stay and listen to that dreadful *Tosca*?' 'How long, in total, did Bertie actually watch the match?'

Similarly, imagine Tom who finds himself in the predicament of the imagined subject in (Grice 1961: 142). Tom has an apparently seamless course of visual experience as of a clock on the shelf. However, through certain periods of this time, his clock experience is not explained by the presence of the clock but rather by an apparatus which causes suitable stimulation of the visual cortex. In such a set up, although Tom will enjoy apparently seamless experience of a clock, we can enquire as regards the durations of periods where he is in perceptual contact with the clock and those where he is hallucinating: 'How long was Tom perceptually aware of the clock?', 'He was conscious of it for about a minute before the apparatus was turned on and he started to hallucinate a clock instead. Then, after another minute, the apparatus was turned off and he saw the clock once more.'

Judgements of partial viewing have as their subject matter the temporal structure of experience itself. Since such judgements are commonplace, this bolsters the above picture. However, such cases are not of central interest. This is because, although they evidence our willingness to ask about the temporal properties of experience as opposed to its objects, questions raised in the context of partial viewings can be reposed so that their subject matter is the objects of experience: 'How much of the match did Bertie

see?', 'How many acts of *Tosca* did Ruby endure?'. Even in the hallucination case we can ask, somewhat clumsily, 'How much of the clock's biography was witnessed by Tom?'

What is of real interest is the distinction between experiential temporal structure and the temporal structure of the apparent objects of experience which can be drawn even in cases where viewings are not partial, cases in which, in the limit, for every temporal part of the event in question, that temporal part is perceived. Even in such cases of maximally complete viewing one can, I suggest, consider theoretically whether the structural ordering of the experience of each of the temporal parts need match the ordering that they are presented or represented as having. Likewise, we can consider whether the relative duration of the experience of a given collection of parts is the same as the apparent relative duration of that collection of parts. For reasons which will become evident in Chapter Two, we rarely ask such questions in a way that contrasts with asking about the apparent objects of perception. Nonetheless, such questions are in perfectly good order.

Even assuming that all perceptual experiences are complete or non-partial, I am suggesting that the following are perfectly good questions. 'How long did your experience of the singing last?' in contrast to the more usual, 'How long did the singing seem to you to go on?' Similarly, 'Did your experience (as) of flash A come before or after your experience (as) of flash B?' in contrast to the more usual 'Did it appear to you that flash A came before flash B?'

These questions are not only perfectly coherent, but amongst them are questions that we feel able to answer simply in virtue of undergoing the experience. They may seem *strange* to us but that is simply because we are not used to contrasting act-time with object-time. Look at your left-hand and then your right-hand. Now consider: 'Which experience came first, your experience of your right-hand or your experience of your left-hand?' If you did as indicated, you can answer knowledgably and with ease: 'My experience of my left-hand occurred before my experience of my right-hand.' Whistle two notes in quick succession, sustaining the second. Now consider: 'Did your experience of the first note last longer or shorter than that of the second?' If you followed the instruction, you are in a position to answer knowledgably: 'Shorter.' The accessibility revealed in both these cases evidences the experiential nature of experience's temporal properties.

As Martin puts it, acknowledging Strawson's (1966: 100-1) insistence that experience must make room for the thought of experience itself, "Any of us, suitably linguistically sophisticated, can move back from judgements about the environment surrounding us to judgements which simply concern our experiential position" (2006: 395).[25] Our experiential position is a position that unfolds over and is structured in time. With suitable linguistic sophistication we can move back from judgements about the environment's temporal structure and the way it unfolds over time to make judgements which simply concern the temporal properties of experience itself.

3.3 Temporal Structure and Episodic Memory

When I sensorily or *episodically* remember an event or process as opposed to merely remembering *that* an event took place, or *that* I saw some particular process going on, I remember the event or process *through* remembering consciously experiencing it.[26] This dependency of memory on conscious experience arguably explains the distinctive nature of episodic memory. Episodic memory has as its primary objects events or processes – my first kiss, moving house, Osipova's performance in *The Upper Room* – but when we recall those events or processes we are best thought of as doing so through recalling our own past experience of them (our experience of kissing, moving or watching). Reflection on memory is thus probative of experience since experience is built into the structure of episodic memory.

According to the above picture, episodic memory tracks experiential, conscious properties of experience. Contrast supposed neural properties of experience such as their locations. Such properties are not retained in episodic memory; we cannot recall the location of the experiential event; this is explained by noting that experiential locations (if

---

[25] Cf. also Evans 1980: esp. 277. An even clearer statement which alludes to the temporality of experience is this remark of R.W. Sellars: "The concrete individual does, and can, distinguish between his consciousness (*his changing field of experience*, not the psychologist's artifact) and a realm which his achieved ideas can intend as their object" (1917: 674, my emphasis).

[26] See Martin 2001, and also Foster 1982, Peacocke 1985, and Martin 2002 on sensory imagination. Some care is needed here because of cases where people remember events (e.g. wedding and dinner parties) but misremember their own locations relative to the remembered scene or cases where people remember seeing someone but misremember the time and place. See here Schacter 2001 on the sin of misattribution. Alive to the fact that remembering is often in part misremembering, we should more cautiously hold that we remember or misremember by remembering or misremembering experiencing.

they exist) are not experiential properties of experience. At most our location as the *subject* of the experience is recorded as it is given in relation to what is perceived. Insofar as a subject's location *is* so recorded, this evidences that part of what it is like to undergo experience is to be aware of objects located egocentrically.

Episodic memory does, however, reacquaint us with temporally structured experiential events and processes. Many can remember in vivid detail watching Ayrton Senna's fatal crash at the 1994 San Marino Grand Prix at Imola. In less than two seconds, Senna's car, traveling at 193 m.p.h., careered off the track at the Tamburello corner throwing up puffs of dust and struck an unprotected concrete wall, tearing off the right front wheel and nose-cone, and sending them into the air. Those who can graphically recall experiencing this whole event have as the object of their episodic memory an original experience with evident temporal structure, a beginning, middle and end. Firstly, their seeing the car leaving the track, immediately followed by their seeing each puff of dust, and finally watching its collision with the barrier and the various car parts flying into the air. The same is true across the modalities. Dwelling on those fateful moments in court, a prisoner may recall *hearing* the sound of the foreman's cough, *listening* carefully to his verdict, and *being deafened* by the silence that followed. In other words, his recollection involves recalling a section of experience with temporal structure. Granting that when I episodically remember an event or process, I remember the event or process through remembering consciously experiencing it, the cases of recollection just sketched evidence the conscious, experiential nature of experience's own temporal properties.

Episodic memory does not typically preserve the structure relating *different* episodes in our stream of consciousness. When we recall longer periods of our lives, we must temporally locate and order episodic memories (with their internal temporal structure) in an overarching narrative. Thus, autobiographical memory involves both semantic and episodic memory. Nevertheless, this much remains true: we think of the narratives that we construct in autobiographical memory as being correct or incorrect depending on, at least in part, the structure that our past experience actually had. In this way, autobiographical memory aims to reflect accurately the temporal structure of our experience. Thus, the way we think of autobiographical memory also evinces our manifest commitment to experience as consisting of events and/or processes which

persist (i.e., occur in temporal intervals) and have an objective ordering (i.e., occur before and after each other).

3.4 Realism and Uniqueness

I have argued that the temporal properties of experience are experiential. That claim is evidenced by the fact that we are in a position to know of the temporal structure of our own experience simply in virtue of undergoing a given course of experience. Moreover, reflecting on episodic memory provides grounds to count certain temporal properties of experience as genuinely conscious or experiential. A consequence of this is that time is distinctive because of its unique combination of properties; it is arguably the only interesting *experiential* determination that is instantiated both by experience itself and by its apparent objects. Even if an identity theory is true, that is no reason to attribute location *as an experiential property* to experience itself. As noted above, perhaps our location as the *subject* of the experience as it is given in relation to what is perceived does count as an experiential property (i.e., part of what it is like to undergo experience is to be aware of objects located egocentrically). Beyond this, however, we have no reason think that the location of experience itself contributes to what it is like to be a subject of said experience (at least not directly).[27]

Are there any other properties which might be thought to be common to experience in its experiential aspect and the objects of experience? Causal relations are one possible candidate. It is increasingly accepted (though very far from uncontroversial) that we can perceive causal relations.[28] We can see a bat causing a ball to fly into the stand or feel the earth causing us to shake. Furthermore, our mental events stand in causal relations with one another. Hearing his vicious remark might cause me to become angry. Seeing paw-prints might cause me to believe that a dog has passed this way. It is far less clear whether we should accept the following two claims, however: firstly, that experiences cause each other (or causally explain features of other experiences); and secondly, that

---

[27] Thanks to Barry Dainton for pressing me here. Dainton also objected by noting that a doctor can ask a patient where exactly the pain in their leg is. Reply: it is plausible to think that pain is a certain kind of awareness of one's body and that the pain's location in this context is the bodily location of the *object* not the act of the pain-experience. See, for example, Crane 1998: §2 for discussion.
[28] For discussion see Siegel forthcoming.

such causal relations, if there are such, are amongst the experiential features of our experience. In particular, neither of the considerations above lend strong support to the second claim here. Introspection does not obviously provide access to the causal relations amongst our experiences (or within our stream of experience) if there are any. We can introspect our experience changing. For example, when we experience a chameleon changing colour we are positioned to be aware of our experience changing character. But this is not obviously a causal matter. That said, nothing in what follows will turn on resolving this issue. Temporal properties are plausibly unique in the above sense, but if other properties must be allowed to be common, so be it.[29]

Let me summarize the key claims made so far in this chapter under the label Realism.

> **Realism** Experience itself in its experiential aspect consists of events and/or processes of which it is true that they persist through time and occur before and after one another.

Realism is, I submit, a part of our ordinary conception of things. It is simply taken for granted by most writers on the perception of time. Moreover, that it is part of our ordinary picture of things is something acknowledged even by those who ultimately reject it. Consider Merleau-Ponty who begins his discussion of temporality by declaring that "all our experiences, inasmuch as they are ours, arrange themselves in terms of before and after … temporality … is the most general characteristic of 'psychic facts'" (2002: 476). Here Merleau-Ponty appears to endorse Realism. As such he evidences that Realism is the natural or default view.

Of course, Merleau-Ponty ultimately follows Kant in rejecting Realism and giving a transcendental idealist account of this datum. For him, the reason our experiences have a temporal structure is "because temporality, in Kantian language, is the form taken by our inner sense" (ibid.: 476); that is to say, (self-)consciousness "deploys or constitutes time" (ibid.: 481), at least the time in which we are aware of experience as being structured, and not the other way around. Nonetheless, the starting point for Merleau-Ponty (and Kant)

---

[29] If causal relations are present in the stream of consciousness itself, we can ask questions like those pursued below. In particular, we can ask whether causal relations in the perceived world are represented by causal relations amongst experiences. I am sceptical.

is the datum of experience's own temporal structure which is granted to be naïvely conceived as something concerning which there are objective facts that hold prior to and explain our self-conscious awareness of them.

Likewise, consider Kant's more recent heir, Daniel Dennett, who writes as follows.

> Our intuitions suggest that our streams of consciousness consist of events occurring in sequence, and that at any instant every element in that sequence can be classified as either having occurred "in consciousness" or as having not occurred "there" yet. (1991a: 144)

I shall come to Dennett's hostility to this picture below where I consider his challenge to Realism. For now his agreement that our intuitive picture involves a form of realism with respect to the temporal structure of our streams of consciousness is all that is important.

4. Mental Events and Physical Properties

Given that Realism is the most natural way of thinking about our perceptual experience and is implicit in our ordinary ways of thought and talk about experience, it is striking to discover apparent rejections of the core claim that mental events have temporal properties from within the philosophically engaged scientific community. Let me give two examples.

Firstly, consider the behavioural psychologist John Staddon who, in a paper discussing the supposedly "fruitless attempt" of philosophers to answer a certain question about motion perception, propounds the following.

> Dennett and Kinsbourne [1992b] write "Conscious experiences are real events occurring in the real time and space of the brain, and hence they are clockable and locatable within the appropriate limits of precision for real phenomena of their type" [235]. Well, no, not really. What can be clocked and located are *reports* of conscious experiences and *measurements* of physiological events. Conscious experiences are *Domain 1* [that is, the 'domain of felt experience, the phenomenological domain' (Staddon 2000: 874)], which has neither time nor space, but only ineffable *qualia.* (ibid.: 879)

According to Staddon then, temporality is simply not the right kind of property to be applied within the phenomenological domain.

Likewise, consider the cognitive scientist Roger Shepherd who, again in responding to Dennett and Kinsbourne, discusses "the problematic attribution of physical relations (spatial orientations or temporal orders) to non-physical objects (mental images or subjective events)" (1992: 223). Shepherd notes that just as *mental* images cannot literally be rotated in space, neither can *mental* events like experiences literally be ascribed temporal properties.

These passages are puzzling, all the more so given that neither writer provides any argument for their claim other than an implicit appeal to the assumption that physical predicates can never be applied to mental states and events. One influential psychologist does, however, attempt to argue for this puzzling view. Zenon Pylyshyn in his brief but influential paper, 'Do mental events have durations?', calls into question the claim that temporal structure is common to experience and its objects, suggesting that it is a mistake to think of mental events as having temporal properties. Interestingly, Dennett (1991a: 148, fn.6) refers to Pylyshyn's paper with approval and, more recently, Treanor (ms.) avers that the paper betters Dennett in offering "a non-verificationist argument against the view that conscious experiences have duration".[30] In this section I briefly discuss this argument, in part for its own interest but also because it introduces issues which receive fuller treatment in the next chapter.

Pylyshyn writes as follows.

> [N]o one is disposed to speak *literally* of such physical properties of a mental event as its colour, size, mass and so on – though we *do* speak of them as *representing* (or having the experiential content of) such properties. For instance, one would not properly say of a thought (or image) that it was large or red, but only that it was a thought *about* something large or red (or that it was an image *of* something large or red). Physical predicates are never appropriate as literal descriptions of *mental* representations or events because the

[30] I discuss Dennett's argument at length in Chapter Five below including the notorious charge of verificationism.

> contents of the latter depend only on their *form* and are intended to reflect what they are *about*, not how they are materially instantiated. (1979: 277)

Pylyshyn here begins by pointing out (surely correctly) that common usage is in accord with Carnap's second claim that psychological objects lack colours and spatial determinations. However, Pylyshyn then makes the much stronger claim that "[p]hysical predicates are never appropriate as literal descriptions of *mental* representations or events". Indeed, earlier, Pylyshyn suggests that "a general covering distinction" can be drawn between mental and physical events "by noting that the two types of events require *distinct* vocabularies for their description" (ibid.: my emphasis). And the implication seems to be that the two vocabularies are mutually exclusive.

Pylyshyn concludes:

> It ought to strike one as curious, therefore, that we speak so freely of the *duration* of a mental event. The temporal dimension has exactly the same status as size or mass: It is, in fact, one of the primary *physical* magnitudes. … Why, then, do we not speak of mental representations as *representing* duration, rather than as *having* duration, in the same way that we speak of mental entities as representing size or mass? (ibid.)

The question, I take it, is rhetorical; Pylyshyn's view is that we should not speak of mental events as having temporal properties, only as representing them.[31]

Pylyshyn's mode of argument appears, on one reading, rather strange. He begins by noting the "curious" fact that we "speak so freely" of the temporal properties of mental events. He then proceeds to use another fact about our thought and talk in relation to other kinds of properties (viz., that we don't talk about the mass or colour of experiences etc.) as support for his general principle of mutually exclusive vocabularies. Yet if our ordinary ways of thinking are probative in this domain, why do they not tell against the principle itself? Surely, the most plausible explanation of the way we speak about different properties in relation to our mental lives is that Carnap is right: time is *uniquely*

[31] Though I think Pylyshyn's view is quite clear, it should be noted that he never explicitly answers his title question.

common to experience and its objects. Certainly that is an interesting fact about time, but it is not curious if 'curious' here implies dubiousness or ineptitude.

Perhaps this is not the best way of reading Pylyshyn. Perhaps what he is saying is that time belongs to the same family of properties as mass, colour and size, viz., the *physical magnitudes*, and that this family comes as a package; objects either do or do not have this whole set of properties. That would ground the move from the claim that experiences lack mass or colour properties to the claim that they lack temporal properties. Unfortunately, the grounding principle cannot be right. First, electrons are not coloured but they do have other physical properties. Furthermore, even if we exclude colour from the *primary* physical magnitudes, modern particle physics posits massless bosons (e.g., photons, gravitons and gluons); again, these may lack mass but they have other physical properties (e.g., location, spin).[32]

If we leave aside these considerations, how should we best understand Pylyshyn's view? An obvious thing to consider is Pylyshyn's metaphysics of mind on which he has the following to say.

> [O]f course each instance of a … mental event is realized as some particular (token) physical event which itself has the full range of physical properties. Thus one could speak of such things as the location, size, colour, and duration of these instantiating events. By extension, one might even speak loosely of the location or duration of the mental event that particular token physical event in fact realized on that occasion. This might be an innocent enough slip, providing that what one was referring to was some particular *token occurrence* of the mental event. (ibid.)

Pylyshyn here implies a distinction between realization and identity. But he could as well think of realization as identity, agree with the materialist point above, but insist that we should distinguish one and the same event viewed *as* mental from that same event viewed *as* physical. His rule of distinct vocabularies would then neatly carve our predicates into the mental ones and physical ones. In any case, it is clearly *not* Pylyshyn's position that because materialism is true, mental events should be thought to have spatial

[32] Not to mention that *processes* and *events* have spatio-temporal locations but are not naturally thought of as having *masses*. Arguably, sounds, though perfectly real physical phenomena, have temporal and pitch properties but only derivatively locations (the locations of their sources).

properties as well as temporal ones. Quite the reverse – Pylyshyn thinks that once we understand the distinction between the mental and the physical properly we will realise than mental events have *neither* spatial nor temporal properties strictly speaking.

The question remains: whence Pylyshyn's hostility towards mental events having temporal properties? As we have seen the key principle Pylyshyn invokes is the following.

> Physical predicates are never appropriate as literal descriptions of *mental* representations or events because the contents of the latter depend only on their *form* and are intended to reflect what they are *about*, not how they are materially instantiated. (ibid.)

As I understand him, Pylyshyn here asserts two things. Firstly, that all descriptions of mental representations *qua* mental will be descriptions of their content (what they are about). Secondly, that representational or semantic content supervenes on syntactic properties of representations (i.e., their form). These are both substantial commitments. However, it is the first claim that is most relevant for present purposes. Indeed, it is essential if Pylyshyn's argument is to be successful, whereas Pylyshyn's claim about the supervenience of the semantic on the syntactic is arguably dispensable.

The first claim is controversial because it is not at all obvious that we cannot describe aspects of mental representations which go beyond their contents, whilst still being genuinely mental. Perhaps temporal aspects are just such aspects. Given this, the immediate worry for Pylyshyn is that, in the absence of a compelling case for treating mental events as exhaustively describable in terms of their representational content, any supposed consequence of that claim is just as liable to be a reason to give the claim up. Thus, even granting that it can be shown to follow from his claim that mental events lack duration, we are very distant from endorsing any such conclusion.

This is not to dismiss the claim that we can exhaustively characterize mental events in terms of their representational content. In the next chapter I consider in more depth a key line of thought which has been invoked to justify it – *the transparency of experience*. However, to anticipate the discussion there, I argue that the temporal properties of experience are indeed transparent but that this is quite in keeping with Realism.

In Pylyshyn's defence it might be noted that his use of the term 'mental' is at least sometimes explicitly technical. In the technical sense, "A *mental* process is … a process viewed under a symbol-processing or algorithmic description" (277). According to Pylyshyn mental processes (in this sense) are instantiated in the neural architecture but are held to be multiply-realizable such that details of the physical (and so temporal) properties of one architecture leave open the 'mental' (i.e., algorithmic or syntactic) properties realized. However, this poses a dilemma for Pylyshyn. On the first horn, the technical use of 'mental' really is technical and thus Pylyshyn need not be taken as denying that mental events in the non-technical, everyday sense have temporal properties. On the second horn, the technical use of 'mental' *is* a genuine elucidation of our ordinary concept and, as such, claims about mental events having durations apply also to such events as perceptual experiences. This second horn would be of genuine interest. However, it faces precisely the difficulty raised above, namely, its reliance on the controversial joint assumptions which underlie the claim that mental processes in the non-technical sense are processes "viewed under a symbol-processing or algorithmic description," i.e., are wholly understandable in representational and so syntactic terms. I take up the first of these assumptions in the next chapter where I consider the transparency of experience in relation to its temporal aspects.

Chapter Two:

Time and Transparency

As the bird trims her to the gale,
I trim myself to the storm of time …[33]

1. Act Time and Object Time

Time is uniquely common to experience and its objects. But this fact leaves open the relation between time of representation and time represented or, in older terminology, between act time and object time. The contemporary common wisdom on this issue is that it would be a serious mistake to think that the two sets of temporal properties (those of experience and those of its apparent objects) must match up. Rather, so the wisdom goes, just like other cases of representation, the properties of the represented must be sharply distinguished from those of the representation since it is quite possible for the two series to vary dramatically from one another. My experience as of a protracted silence might itself last only a moment; my experience as of a flash followed by a bang might itself be structured in the opposite order.

Certainly, it is true that all that we can determine *from a theoretical perspective* is that there is a conceptual distinction between the temporal structure of experience and the temporal structure of its apparent objects. This leaves it an open question as to how the two are related. In this respect, I disagree with Roache who, commenting on a debate between Mellor (who argues that act time and object time must match) and Dennett (who demurs), claims, "it would be contradictory to assert that the order of perceptions may differ from the perceived order" (1999: 237).[34] In the abstract, the possibility of the order

---

[33] Emerson (1867) 'Terminus' from *May-Day and Other Pieces*.
[34] Mellor argues as follows, "Suppose for example I see one event e precede another e*. I must first see e and then e*, my seeing of e being somehow recollected in my seeing of e*. That is, my seeing of e affects my seeing of e*: this is what makes me – rightly or wrongly – see e precede e* rather than the other way

of conscious perceptual experience coming apart from the apparent order of its objects is in no way contradictory. Indeed, I see no purely theoretical reason for thinking that the two times must stand in any particular relation to one another.

On the other hand, this concedes a lack of theoretical grounds for positing a relation between act and object time series, not a lack of grounds. There may be other grounds for doing so. In the first part of this chapter, I raise some doubts about the contemporary wisdom that the two series can in fact come apart. In the second part, I argue that introspective reflection on our own experience provides precisely the grounds to reject that wisdom and to insist that act and object time are yoked together. Roache is right to this extent: we find the possibility of the order of perceptions diverging from the perceived order of their objects incoherent *when we reflect upon our experience*.

Throughout I talk in terms of representation and represented, since in this context (even more than elsewhere) it is assumed that experience has a representational content. It remains possible to construe the debate in other terms, however. As I have mentioned, there is an older distinction which can be invoked, that between act and object. Thus, in *The Problems of Philosophy*, Russell warns us, in strikingly contemporary terms, that we must take great care to distinguish "on the one hand the thing of which we are aware – say the colour of my table – and on the other hand the actual awareness itself, the mental act of apprehending the thing" (1912: 21), emphasising that this "distinction between act and object in our apprehending of things is vitally important" (22). This traditional distinction between sensation or psychical occurrence or mental act or experience on the one side, and sense-datum or sensible object or mental content or object of experience on the other, is found in most views of perception bar adverbialist approaches that deny that experience has this structure. I take its application (and so the inadequacy of adverbialist approaches) for granted in this thesis.

---

round. But seeing e precede e* means seeing e first. So the causal order of my perceptions of these events, by fixing the temporal order I perceive them to have, fixes the temporal order of the perceptions themselves" (Mellor 1981: 8; discussed in Dennett 1991a: 149). The argument of this chapter sides with Mellor with respect to the claim that "seeing e precede e* means seeing e first," i.e., entails that my experience of e must come before my experience of e*. However, nothing I say commits me to Mellor's views on the roles that memory and/or causation play in allowing for temporal experience of succession. On the role of memory see Chapter Six below.

The act/object distinction is, perhaps, wont to be overlooked because of a general grammatical ambiguity. John MacFarlane glosses the ambiguity as follows.

> **Act-object ambiguity** Ambiguity of many common words (e.g. "utterance", "experience") between an "act" reading ("the act of uttering something," "the act of experiencing something") and an "object" reading ("that which is uttered," "that which is experienced"). Also sometimes called the ing-ed ambiguity.[35]

In what follows, talk of representation and represented can be freely translated into the idiom of act and object. Indeed, as I argue in Chapter Six, the possibility of experiential act-time and object-time coming apart is essential to most theories of temporal awareness, and can be seen very clearly in the diagrams of early theorists such as Husserl and Broad as well as more recent writers such as Tye and Gallagher.[36]

2. Confusing Representation and Represented?

Daniel Dennett, along with his collaborator Marcel Kinsbourne, is the most famous contemporary exponent of the view that a coming apart of time represented from time of representation is not only theoretically possible but actual.[37] Dennett contends that we must "distinguish time represented from time of representing" (1991a: 161) and goes on to suggest that the two can readily come apart. According to Dennett, it is quite possible for my experience as of a long silence to itself last only a moment, or for my experiences as of a flash followed by a bang to occur in the opposite order to which the events are represented as occurring in (cf. his 1991a: 144). For, on Dennett's view, all that matters for determining the order of the apparent objects of experience is the temporal *content* of the experience and, he insists, there is no reason why such contents cannot be tokened in temporal orders quite distinct from the temporal orders of the events that they represent. To think otherwise is to confuse time represented with time of representing.

---

[35] See MacFarlane's philosophical glossary at http://johnmacfarlane.net/135/glossary.html. Calling it the "ing-ed" distinction is helpful since it removes the misleading suggestion that the distinction has anything to do with *actions per se*. In particular, talk of the act of experiencing something has no implication that experiencing is or has any close connection to acting.

[36] See Husserl 1964, Broad 1923, Tye 2003 and Gallagher 1998.

[37] See in particular Dennett and Kinsbourne 1992a, material from which is incorporated into and expanded upon in Dennett 1991a.

His central arguments for this claim focus on psychological experiments which involve short timescale phenomena. Dennett's exploitation of these experiments (and what they in fact show) is discussed at length in Chapter Five below. For present purposes, however, I want to discuss some intuitive analogies that Dennett uses to warm us up for his main argument. In part this is because I want to make sure we remain stone cold; the analogies that Dennett pumps us with should be set aside because the supposedly analogous cases are in no way plausible models for temporal experience[38] except in one case where the analogy in fact tells *against* Dennett's contention that time of representation can come apart from time represented. It is also because other writers who are impressed by Dennett's 'representation/represented confusion' critique often do rely upon such analogies.

First, consider Susan Hurley (1998: Ch.1, *passim*) who, citing Dennett, urges that we should not "suppose that temporal representations represent their own times, or project temporal relations between vehicles of content into the content of temporal representations" (ibid.: 29). Hurley argues as follows.

> In general representations do not have to resemble what they represent. This is clear enough in some cases: no one thinks the representation of something green must itself be green, or that the representation of something round must itself be round. But it is easy to slip into this confusion for more complex, abstract or relational properties, such as simultaneity. (1998: 29-30)

Here it is not clear what kind of representation of green is meant. But, if as seems most likely it is the salient experiential case that she intends, then Hurley is right. Experiences which represent green trees do not themselves have to be green. Nor of course is Emily Dickinson's reference to Spring – "This whole experiment of green" – itself green! Yet it is hard to see how this bears on the case in point. Time is *uniquely* common to experience and its objects. Consequently, there is no clear analogy at all between the representation of green in experience (or poetry) and the representation of simultaneity or relative

---

[38] Reingold writes: "[Dennett & Kinsbourne] not only expose brilliantly an important confusion between the temporal properties of the process of representing and the temporal content of the representations themselves, but they also provide powerful metaphors that may help one sliding back into this ingrained confusion" (1992: 218). Precisely not.

duration in experience. Experiences do not have colour properties; they do have temporal properties.

Second, consider Michael Tye, who writes as follows.

> It seems to me that there is a serious confusion here.[39] Granted, I experience the red flash as being before the green one. But it need not be true that my experience or awareness of the red flash is before my experience or awareness of the green one. If I utter the sentence
>
> The green flash is after the red flash,
>
> I represent the red flash as being before the green one; but my representation of the red flash is not before my representation of the green flash. In general, represented order has no obvious link with the order of representations. Why suppose that there is such a link for experiential representations? (2003: 90; see also his 2006: 511)

The same elements are in play here. Firstly, the assertion that writers who have thought that experiential temporal structure matches the temporal structure of its objects are seriously confused. Secondly, that they are making the same kind of mistake as the (I hope fictitious!) person who holds that the order of events has to match the order of referring terms in the English sentences used to refer to them. Finally, obliviousness to the possibility that time might be special – witness the last sentence which does not mention time at all.[40]

---

[39] Tye is criticising Dainton (2000: 134) who picks up on Miller's *Principle of Presentational Concurrence* (PPC) according to which "the duration of a *content* being presented is *concurrent* with the duration of the *act* of presenting it. That is, the time interval occupied by a content which is before the mind is the very same time interval which is occupied by the act of presenting that very content before the mind" (Miller 1984: 107). See §6.2 below.

[40] Hurley and Tye are far from alone. Treisman writes, "The time represented by an element of phenomenological experience is not the time at which that element is generated – as the naïve realist with respect to time would suppose – but the time to which it refers, just as beige is not the colour of a conscious sensation but of my word-processor" (1992: 225). Shepherd backs up his agreement that "*the experience of* (*A before B*) does not entail (*the experience of A*) *before* (*the experience of B*)" by noting Kubovy's observation that "*to imagine* (*a rotation of* (*an object*)) is not necessarily to *rotate* (*an image of* (*an object*))". On this he comments, "The second statement in each of these … italicised pairs entails the problematic attribution of physical relations (spatial orientations or temporal orders) to non-physical objects (mental images or subjective events" (1992: 223). Perhaps so in the case of imagination but, as discussed in Chapter One, I cannot understand why Shepherd thinks it is problematic to attribute temporal properties to subjective events. Block who in other respects trenchantly criticises Dennett and Kinsbourne simply remarks, "[Dennett and Kinsbourne] correctly point out that the temporal order of outside events needn't be represented by the temporal order of inside events. This Kantian point (Kant distinguished apprehension

It would not be fair to say that Hurley and Tye *simply* rely on these analogies. They also allude to Dennett's more sophisticated argumentation.[41] However, like Dennett, they all begin by making the quite general point that in general representations do not have themselves to possess the properties they represent and draw the conclusion that there is no reason why things should be any different in the case of experience and time.[42] If we agree, Dennett's discussion and diagnosis of various psychological phenomena may seem quite reasonable – just another case where the properties of a representation come apart from what is represented.

Dennett plays this game by making repeated and vivid appeals to the way properties including time are represented in other media such as novels, letters, pictures and films followed by the implicit suggestion that our mental lives should be expected to operate along similar lines. However, as already noted, for such appeals to be probative the ways in which such media represent must be analogous to the experiential representation of time. In the rest of this section I consider the cases Dennett appeals to. I have two aims in doing so. Firstly, to bring out that, in most of the cases appealed to, there is a clear disanalogy between the experiential representation of time and the kind of representation in question. Secondly, to consider in more depth a case where the analogy does appear compelling: cinematic depiction. I argue that cinematic depiction is, on reflection, very plausibly a case in which, at least within scenes, the temporal structure of depiction does mirror the temporal structure depicted. As such, the case tells *against* Dennett, Hurley, Tye and others. That said, reflection on the analogy also helps to introduce, and to some extent motivate, a form of anti-realism about depictive temporal structure and, by analogy, the anti-realism about temporal experiential structure which is the focus of Chapter Five.

---

of succession from succession of apprehension) is certainly correct" (1992: 206). Block here gives no indication of why he thinks Dennett and Kinsbourne are "certainly correct". Indeed he *rejects* their Stalinist/Orwellian arguments; he does not appeal to analogies as they and others do; and the appeal to Kant only establishes a conceptual distinction and not a genuine possibility. Remarkably, of the twenty-eight original commentators on Dennett & Kinsbourne (1992a), Rollins (1992) is the only one who questions the view that time of representing might come apart from time represented. Others who explicitly endorse the view include Clark (1992: 207), Farah (1992: 209), and Lloyd (1992: 215).

[41] Though Tye (1993) is extremely hostile towards Dennett's arguments so it is not clear what weight he thinks that they can be given in this respect.

[42] Hurley (1992: 211) also makes some suggestive diagnostic remarks as to why we might be tempted to 'confuse' vehicle with content in the temporal case. One can read the second part of this chapter (in combination with the chapter that follows) as explaining why we are right to be so tempted.

3. Relating Representation and Represented

3.1 Novels

Dennett illustrates his claim that time represented can come apart from time of representing by comparing experience to a novel (1991a: 148). He notes, obviously correctly, that a novel is not constrained to represent events in the order that they are represented as occurring in. A later paragraph can refer to events which are put forward (in the fiction) as occurring well before events described in the previous paragraph. This shows that there is a distinction here between the order of representation and the order represented.

However, Dennett's analogy is poor. Experience not only presents temporally structured events but is itself a dynamic event with temporal parts. In contrast, though a novel has an ordering (either spatial or simply structural) amongst its constituent chapters, paragraphs, and sentences, there is no sense in which a novel has a *temporal* structure of its own. A novel is not an event with dynamic properties, nor is it composed of such. It does not have temporal parts. As Gennette says, "The narrative text, like every other text, has no other temporality than what it borrows, metanymically [sic.], from its own reading" (1979: 34; quoted in Currie 1995: 94, fn.14).

Now, as this remark suggests, novels are intended to be read or listened to, and reading and listening take time. However, this would only begin to introduce an analogy with experience if we held that the reading of a novel *had* to have a certain temporal structure to be correct. The temporal structure might then be thought of as *the novel's* in some sense. But it is not plausible to insist that a novel must be read line by line (word by word?) starting at the beginning and working through at constant pace. No-one is forcing you to read the text in any particular order, nor to spend any particular length of time in doing so. Indeed, in theory, one is at liberty to take in a whole novel at once – speed readers do take in whole pages (and so presumably very short stories) at once. There is nothing incorrect about such a reading; there is nothing about the narrative that one cannot appreciate from non-standard readings. What is important in a reading of whatever sort is that the novel's (non-temporal) structure is grasped by the reader; that

grasp does not require a reading with any particular temporal structure. Thus, there is no good sense in which a novel has its own temporal structure and so no analogy with experience.

3.2 Letters

Elsewhere Dennett (1991a: 146-7) appeals to an analogy with letters and messages arriving from different locations but with postmarks to determine their correct temporal order. Here things are a little more complex than with novels for there are further features of the letter or message potentially analogous to the time of occurrence of an experience. In particular, leaving aside internal structure, we can consider the time of arrival or reading of the letter or message as analogous to the time of occurrence of the experience, and the post-mark of the letter as representing the time of the content represented.[43] Postmarks are useful because they allow one to determine time of sending without reference to time of arrival. Consider a very simple SMS text message (to avoid problems arising from internal temporal structuring) which simply says, 'I am leaving now.' Imagine one is without reception and receives this message sometime later. The time of reading is not the time of sending. To interpret the message one looks at the digital postmark and determines what time 'now' refers to. In this way time of representing (receipt) comes apart from time represented (the reference of 'now').

Our perceptual experience *could* work this way. We might, as Guyer imagines, find "a digital time-stamp on every one of our perceptions" (1987: 421). But we evidently find no such aspect to our actual experience immediately casting the analogy into doubt.[44] Indeed, in contrast to the message case, it seems we are licensed immediately to make present tense judgements on the basis of our perceptual representation: 'That is happening now,' where 'that' is a visually grounded demonstrative.[45] We do not have to check the post-mark!

---

[43] Of course, postmark and letter content timings can come apart. But let us assume that the present tense is in use throughout the letter.

[44] As Bennett nicely puts it, "Events don't come with their 'when' written on their faces" (2004: 173).

[45] As Broad puts it, "It is of the essence of a perceptual situation that it claims to reveal an object as it *is* at the time when the situation is going on" (1929: 145). It is commonly objected to this idea that it takes time for light to reach the eye and for processing to occur. But this would only show that experience could not license present tense judgements if the present introduced by the present tense were strictly momentary. It

3.3 Paintings

A case more plausibly analogous to the experiential case is pictorial representation. Here *space* is a common medium to most paintings and their represented scenes. Moreover, it is not difficult to imagine a painting (cf. Harman 1990: 34; also Dennett 1991a: 131, 147) in which the spatial relations between objects represented and the spatial relations between representings of those objects diverge in a manner analogous to that which is being suggested is possible in the case of time and experience. For example, consider Brueghel's 'Landscape with the Fall of Icarus' (below).

De val van Icarus, Pieter Brueghel the Old, Wikimedia Commons.

Here the shepherd on the promontory is clearly represented as much higher up than the two rocks in the centre of the picture, the smaller ships and all the sea. Yet the painted figure representing the shepherd is lower on the canvas than much of the sea, those rocks and ships. Likewise, the painted figure of the tilling farmer is right up against the

---

is far more plausible to think that the present tense introduces a contextually determined *interval* of time including the present moment.

shepherd on canvas and yet the figures are clearly represented as being some distance apart. (Dozens more examples can be pointed to.)

Unfortunately, we really have no reason to think the analogy holds of temporal representation in experience. For one, the discrepancies noted above are due to the representation in two dimensions of a three-dimensional scene. No such gap is present in the uni-dimensional time case. Considering sculptures overcomes this issue and still appears to raise the possibility of differences between representing and represented most obviously in terms of scale.[46] However, it is more fruitful at this point to look at the best case for an analogy, viz., film which evidently possesses its own temporal structure. The issues concerning representation versus depiction that cinematic representation raises can then be carried back to the sculpture case.

3.4 Films

The best candidate for an analogy with experience is film. Dennett mentions them several times in passing, for example (1991a: 137, 152), and implicitly in his talk of editing rooms, institial frames and tape delay. We standardly recognise three distinct temporal structures with respect to films.

(i) The temporal structure of the scene that is filmed.
(ii) The temporal structure of the narrative represented.
(iii) The temporal structure depicted by the film.[47]

These structures rarely line up. Scene structure and depicted structure almost always come apart. The scenes depicted in modern Hollywood films are typically made up of a large number of individual shots combined in the editing room. Thus, although a single scene is depicted, many different scenes have been filmed (perhaps some in the studio,

[46] Though given the points of Ch.1, §1.1 about metrical properties, this is not obviously relevant.
[47] How to understand depiction is, of course, a vexed issue. Nonetheless, it is widely agreed that there is a distinction between representation and depiction when it comes to pictorial or cinematic representation. Thus, Hopkins notes how much Western religious art merely represents the Holy Spirit by depicting a dove (1998: 9). Likewise, in Lynch's *Mulholland Drive* what the significance (and so representative significance) of a mysterious blue box is may be quite obscure, yet there is no doubt that a blue box is being presented or depicted cinematically. Hopkins also argues that it is only the representation (depiction) of the dove which counts as distinctively pictorial. I suggest the same applies in the cinematic case.

some on location). A famous exception which attests to the general rule is Sokurov's *Russian Ark* where the whole film comprises a single, stunningly choreographed, ninety-minute Steadicam shot. The same applies to scene structure and narrative structure.

Narrative and depicted structure also often come apart. An example is the use of reversal effects within the intended structure of a film such as Avary's *The Rules of Attraction* which makes liberal use of such techniques. For example, at one point in the narrative, a piece of film of a plane moving through the sky is played backwards within the intended structure to indicate a backwards shift in narrative. What is strictly speaking depicted, however, is a plane moving backwards through the sky (this despite our knowledge that the scene actually filmed involved a plane moving forwards through the sky). The depicted order has the plane first at point B, then at point A; the narrative and scene structure places the plane first at A, then at B.

In contrast to painting and photography, film is a paradigmatic temporal art. What exactly does this mean? The point is not that film uniquely represents motion or succession. Paintings and photographs can clearly represent motion. And it may be that Walton is right to claim that photographs *depict* motion also, for example, through blurring or unstable configurations (as in the photograph below).[48]

Currie suggests the following: "with painting, temporal properties of events are not represented by *temporal* properties of representations, and the reason is clear: painting is temporal$_r$ [i.e., it represents temporal properties] but not temporal$_w$ [i.e., it does not itself possess temporal structure], so it does not have the capacity to be temporal$_c$ [that is to represent time by means of its own temporal structure]" (1995: 99). Walton makes a similar point.[49] Photographs contrast with films in that they do not represent movement or change by themselves moving or changing. He also notes that "a change in what something depicts is not necessarily a depiction of change" (2008: 163). In contrast,

[48] See Walton 2008: 158f. for a variety of examples and discussion. I would not want to insist on this. It may well be that photographs merely represent movement in virtue of depicting multiple or blurred images, or unstable scenes. See also discussion in Le Poidevin 2007: Ch.7, esp. 134f..
[49] See also Yaffe 2003.

Walton *defines* motion pictures as "pictures whose temporal properties do contribute to their representational content" (164).[50]

Freak Out, Oblivion © Andrew Dunn, 2005, Wikimedia Commons.

One might object at this point that a film only inherits its temporal structure from its watching and is consequently not really any different from a novel. However, in contrast to a novel, a film is created such that there is an intended temporal structure that the viewing public is to observe in watching.[51] If the film is to be viewed properly, the audience must watch the film projected in this intended manner. One can, of course, *decide* to view a film in some perverse way – at a different speed or, indeed, backwards. If

[50] This should be borne in mind in considering the idea that a succession of representations is not a representation of succession in Chapter Six. Though Walton does not make this point, reflection on cinema suggests that in some circumstances a change in the depiction does amount to a depiction of change.

[51] Cf. Hamilton (ms.) on music: "Music is quintessentially an art of time … our experience of music is time-bound. Experience of painting, like all experience, takes place over time; but experience of musical performance is constrained by a specific intentional order of aural events."

one does so, one has taken a decision not to view the film as it was intended to be seen, and so deliberately to view it incorrectly (though perhaps for good reason when it comes to some films). No such choice is typically present with still representations of time such as paintings or photographs.[52] And, as noted, no-one is forcing you to read a novel in any particular order and, in theory, one might take it all in at once.

It is possible to build in the idea of a viewing time to a painting, insisting that the work be viewed, say, for precisely an hour to have its proper effect. The artist may explicitly intend his or her painting to be viewed in this manner. Or one might, more generally, hold it essential to appreciating the full meaning of a work that it be contemplated for several hours. Thus, Wollheim notoriously describes his practice of slow-looking as follows.

> I evolved a way of looking at paintings which was massively time-consuming and deeply rewarding. For I came to recognize that it often took the first hour or so in front of a painting for stray associations or motivated misperceptions to settle down, and it was only then, with the same amount of time or more to spend looking at it, that the picture could be relied upon to disclose itself as it was. (1987: 8)

Does this undercut the distinction above between the temporal and non-temporal arts? It does not. Whilst we should certainly grant that appreciating static art takes time, we should note the following key difference. Wollheim's slow-looking technique may well have been conducive (at least for him) to appreciating certain paintings. However, he notes that his method is merely a reliable means towards an end. The technique is not *essential* to attending to the work as it truly is. In contrast, watching a film as intended does introduce genuinely new possibilities for attention, possibilities which may have been simply unavailable during perverse viewings. One cannot attend to the backwards motion of the plane in *The Rules of Attraction* if one is viewing the film on rewind. Such novel possibilities are not essentially afforded by viewing a painting for an hour or three (or by attending in some particular order to different aspects of the work).[53] Certainly the effect a painting has on us can depend dramatically on our manner of engagement. And,

---

[52] One might choose to view a painting upside-down, a point which evidences the fact that a painting has an intended orientation. This fact is further evidenced by the *mistakes* that galleries sometimes make in hanging artworks.

[53] Though there are difficult cases here such as Bridget Riley's *Movement in Squares* and other Op. Art. Cf. Gombrich 1964: 306.

moreover, a certain amount of time may be necessary in order to see all of the painting (especially if the canvas is large).[54] Nonetheless, strictly speaking, no new element is made available for attention by viewing in any particular manner (such as slow-looking) as there is when a film is watched according to (or against) its intended structure.[55]

Films – unlike photographs and paintings – then essentially have their own temporal structure. Thus we can consider how facts about the temporal order of the film relate to facts about the depicted temporal structure, the temporal structure of the narrative, and that of the scene. Jean-Luc Godard once famously commented,

> *Indeed, naturally I think that a film should have a beginning, middle, and an end – but not necessarily in that order.*

One thing this indicates is that which events are represented as occurring first, last and in-between *in the narrative* can come apart from the order of their representation. A film may begin at the end and 'flashback' to events that happened earlier as exploited so effectively in Nolan's *Memento* or Bresson's *Une Femme Douce*. On a smaller scale, consider again the plane in *The Rules of Attraction* which in both actuality (scene) and in narrative is first at point A and then at point B. Contrast the intended viewing of the film in which we first see the plane at B and then at A.

It is depiction though that is distinctive of cinematic representation. So what then of the relation between depicted structure and order of depiction? As mentioned above, Currie holds that "[w]hat is distinctively temporal about film is not its portrayal of time, but the manner of its portrayal: its portrayal of time by means of time" (1995: 96). Currie goes on to discuss whether the time of representation and time represented need match.

---

[54] Something which casts into doubt a strict interpretation of Walton's view that "At each individual moment during my five-minute observation of [a] photograph, I see (in imagination) the momentary occurrence that the picture depicts, a short time slice of the moving bicycles or the running kids, and see it in a moment" (2008: 181).

[55] Brough (2000: 228) presses the opposing view discussing the way a painter such as Brueghel guides the eye around the painting, deliberately aiming to keep us from noticing Icarus until the last. The question for Brough, however, is: Have I viewed the painting *incorrectly* if I do happen to cast my eye upon Icarus first off, or, indeed, am somehow able to take the elements in 'all at once'? And, likewise, what have I *missed* in doing so?

> Film's representation of time by time can be automorphic or merely homomorphic. The represented fight lasts five minutes, and its lasting that long is represented by the relevant representation onscreen lasting just that long. It is the default setting for cinematic interpretation that the representation of duration in cinema is automorphic; it is the assumption we start with, and from which we move only when some aspect of the narration, some clash with real-world belief or some combination of the two suggests we should. In Pasolini's *Gospel According to St Matthew* the representation of the Sermon on the Mount lasts a few moments. But changes in the background and lighting suggest that the whole performance lasted much longer and took place at various locations. Here the context of narration and real-world belief conspire to shift our understanding: the changes of scene were meant to indicate something, otherwise they would not have been made; landscape and lighting are, and are commonly known to be, by and large locally stable; so what is probably being suggested here is discontinuous shifting of place and time. (1995: 99)

He then notes,

> You might argue that what is represented here is not the whole sermon, but just a few parts of it, and thus we can preserve the default setting. But to insist on this move in all cases would be, in effect, to identify what is represented with what is displayed, and that identification is hard to sustain. (ibid.: 99-100, fn.16)

There are I suggest two ways of developing this picture. The first and most natural development emphasises the distinction between depicted ("displayed") temporal structure and narrative ("represented") temporal structure. In that light, Currie's discussion of Pasolini reveals two things. Firstly, that although default in some sense, our interest is not always primarily in the temporal structure depicted, but rather that of the narrative. Secondly, that depictive representation, at least with respect to *parts* of the sermon, does appear to be consistently automorphic whereas narrative representation is merely homomorphic.[56] If this is right, cinema involves a distinctive and in some sense basic form of automorphic depictive representation.[57]

---

[56] Yaffe (2003: 117f.) correctly points out that automorphic representation cannot simply be understood as any case where a representation has the very same properties as what is represented as in 'black ink'. He suggests that representation is automorphic where the very properties represented are not just also possessed by the representation but also those in virtue of which said properties are represented.

[57] Compare also cases where films *represent* the passing of often lengthy periods in the narrative by exploiting the *depiction* of a rapidly sinking sun or a quickly turning clock hand.

Can we conclude on this picture that the temporal structure of the events depicted always mirrors the temporal order of depiction? Not quite. Godard is also said to have remarked, "Movies are a world of fragments." And it might be thought that what Currie's discussion indicates is that order of depiction and order depicted match *within fragments*. We do not, after all, think that a cut between two quite separate temporal perspectives in any sense depicts events as continuing on from earlier events. Much more plausibly, no connection is depicted. Nonetheless, it does seem right to hold that, within each fragment or scene, time of depiction and time depicted cannot come apart even if, at a grosser level, time of depiction need not map onto time depicted.[58]

Does the analogy with depiction in film then support the idea that the temporal ordering of experience itself can come apart from the temporal ordering of events presented? It does not. Rather, at least within fragments, we are inclined to hold that the temporal order of depiction mirrors the temporal order depicted. Hence, by analogy, act-time and object-time should match at least within unified stretches of experience.[59]

I mentioned a second way of developing Currie's discussion of Pasolini. And here we find an analogue for a form of anti-realism about experiential temporal structure which presages discussion in Chapter Five. This second account begins by noting how Currie points to the way in which various non-temporal features of the film's content (changes in background and lighting, for example) contribute to our grasp of the narrative temporal structure of the film. It is these non-temporal features which pull the narrative structure away from the film's own temporal structure. Above, this led us to distinguish narrative temporal structure from depicted structure, the latter continuing to match the film's temporal structure. However, the conclusion drawn here is that there is no

---

[58] Note how we cannot really make sense (at least as cinema) of the extreme case where we make scenes/fragments simply one frame long. In other words, there must always be scenes of some length within which mapping occurs.

[59] Lee (2007: 343) attacks a "confused tendency in our thinking about temporal experience" which he labels the "Cinematic" view of temporal perception. Lee's Cinematic View is the view that "experience presents temporal phenomenology in virtue of its own temporal layout" a claim which has the consequence that there is "a direct link between the temporal properties of perception, and its temporal content". We should separate out claim and consequence here. I endorse the consequence and its "content/vehicle confusion," regardless of the fact that "most would agree that the model is prima facie very dubious" (373, fn.9). On the other hand, I see no reason to accept the claim with its implication that experiential temporal content is possessed *in virtue of* the intrinsic temporal properties of experience. Why not think that the temporal structure of experience is what it is in virtue of the world's temporal structure? That is, why not think that the world induces experience's temporal structure (at least in the good case).

distinctive depiction of temporal structure in film apart from the temporal structure represented in the narrative.[60] We are misled into thinking that there is a distinctive kind of depiction because in simple, perhaps default cases, the narrative structure mirrors the film's structure. But, on this second account, it is suggested that there is no reason to continue to believe in a form of depiction which mirrors the film's own structure in cases where the narrative structure departs from the film's structure. There is no more to the representation of time in films than the temporal structure of the narrative.[61]

This is not the natural view of the matter. But we might be led to this kind of view by pressing cases where there is a tension between the narrative and depicted structure. Vivid examples come in letter scenes. We might imagine that one character receives and reads a letter sent by another character, the scene showing this by switching between the reader and writer.[62] We know full well that these switches are between quite disjoint times – the letter is not being read as it is written! Nonetheless, we are naturally inclined to think of such cases in terms of a single scene regardless of the fact that it involves two (or perhaps more) temporally separated perspectives. It can, in such cases, feel uncomfortable to insist either that there are in fact a series of separate scenes here spliced together and that only within each scene does depicted structure mirror order of depiction, or alternatively, that strictly speaking what is depicted is a temporally continuous sequence despite the striking contrast with the narrative (and what we know about letters) in this respect.

On the first and most natural way of developing the picture above we must insist on one of these options, plausibly by distinguishing between different conceptions of a 'scene'. However, if we place a great deal of weight on such cases it is possible to insist that what they really show is that there is no more to the supposedly distinctive depicted time structure than what is represented as being so in the narrative. The tension in the above cases is then resolved by discarding a false assumption, namely, the idea that there is such

---

[60] Something like this reading is suggested by Yaffe (2003: 121) who holds that often "considerations of narrative coherence" account for representation of order as opposed to the temporal ordering of the representations. Yaffe does not distinguish narrative representation from depiction leaving him with something like the account above.

[61] Cf. the final sentence of Dennett and Kinsbourne 1992a: "The temporal sequence *in consciousness* is, within the limits of whatever temporal control window bounds our investigation, purely a matter of the content represented, not the timing of the representing."

[62] Recall, for example, various adaptations of novels such as Austen's *Pride and Prejudice* where Lizzie receives a letter from Darcy or Jane. Or think of the email exchanges in Ephron's *You've Got Mail*.

a thing as distinctively depictive structure. As I say, this is not our natural view of things and the motivation just sketched is not in itself compelling. However, the picture is helpful in warming us up for Chapter Five where I explore a related, anti-realist view of the perceptual case. That view rejects the naïve idea that experience has an experiential temporal structure of its own which matches the temporal structure of the apparent objects of experience. Instead, it claims that there is nothing more to experiential temporal structure than the structure that experience is represented as having in autobiographical memory as grounded in aspects of experience's non-temporal content. Though initially hard to believe, there are in fact powerful arguments in favour of such a view.

For now we should conclude as follows. Firstly, cinema, the most plausible analogy for the experiential representation of time, in no way suggests that the temporal structure of experience and the temporal structure of the objects of experience easily come apart. At least within unified stretches of experience, it suggests that they must match. Secondly, reflection on cinema flags (though it does not significantly motivate) the availability of a form of anti-realism about experiential temporal structure which will be our central concern in Chapter Five.

4. Transparency

So far the discussion of how act-time and object-time might relate has been largely suggestive. I have argued that we should be sceptical of most supposed analogies with other representational media, and I have also noted that the best analogy available in fact suggests that act-time and object-time mirror one another. Might we argue for or against this conclusion directly? I think we can. Crucial in this regard is the idea that experience is *transparent*. In this section, I give a brief overview of that idea as I think it is best understood. In the next, I relate those general considerations to the case of time.

What exactly is the thesis of transparency as it is understood in contemporary debates in the philosophy of perception? A profusion of transparency theses and distinctions makes this a difficult question to answer. Here I will focus on two canonical statements of the

supposed transparency of experience – those of Gil Harman and Michael Tye.[63] How do they understand the notion?

The passage standardly cited from Harman's classic paper 'The Intrinsic Quality of Experience' is the following.

> When Eloise sees a tree before her, the colours 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 experience. And that is true of you too. There is nothing special about Eloise's visual experience. When you see a tree, you do not experience any features as intrinsic features of your experience. 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, including relational features of the tree "from here." (1990: 39)

And here is a typical passage from Michael Tye's influential work on transparency.

> Focus your attention on a square that has been painted blue. Intuitively, you are directly aware of blueness and squareness as out there in the world away from you … [as] features of an external surface. Now shift your gaze inward and try to become aware of your experience itself, inside you, apart from its objects. Try to focus your attention on some intrinsic feature of the experience that distinguishes it from other experiences, something other than what it is an experience of. The task seems impossible: one's awareness seems always to slip through the experience to blueness and squareness, as instantiated together in an external object. In turning one's mind inward to attend to the experience, one seems to end up concentrating on what is outside again, on external features or properties. (1995: 30; see also Tye 1992, 2002)

Both Harman and Tye instruct us to perform an experiment: "turn your attention to intrinsic features of your visual experience"; "shift your gaze inward to try to become aware of your experience itself, inside you, apart from its objects"; "focus your attention

---

[63] The term 'transparency' originates with Moore 1903. However, as Martin (2002) points out (see also Stoljar 2004), Moore's interest is in *diaphanousness*, the claim that sameness and difference in phenomenal character can be *exhaustively* understood in terms of sameness and difference amongst presented objects (of whatever kind). Martin proposes a more likely ancestor of transparency as here introduced in Grice's talk of intrinsic qualities of experience in his (1962); Strawson 1979 is also a clear influence.

on some intrinsic feature of the experience that distinguishes it from other experiences". Both think that, try as we might, we will fail in the attempt.

What is the nature of the experiment to be performed? Evidently, they are not to be read literally, as suggesting that we shift our outward perceptual attention or gaze inwards, so that the very same capacity is targeted on the "intrinsic features of your visual experience" (Harman), or "your experience itself … some intrinsic feature of the experience" (Tye). The way in which we can attend to a tree in the middle distance, or to the birdsong emanating from it, is not a way in which we can attend to our experience itself. Our experience is not itself an object of sense experience. Nonetheless, there is a clear sense in which we *can* attend to our experience; we can *think* about it. As Martin puts it,

> In general, whatever we are prepared to call an object of thought – be it the things thought about, what one thinks about them, or the proposition one thinks in thinking these things – we can also take to be an object of attention. Conscious, active thought is simply a mode of attending to the subject matter of such thoughts. … Arguably, it is part of the manifest image of the mind that we are aware of objects of sense experience in a different way from being aware of the objects of thought, and that this is reflected in the ways attention can relate one to an object of sense as opposed to thought. (1997: 78)[64]

Harman and Tye are suggesting that, when one engages one's *cognitive* attention in thinking about samenesses and differences amongst our experiences, "one seems to end up concentrating on what is outside again" (Tye); or, as Harman puts it, "the only features there to turn your [cognitive] attention [i.e., thought] to will be features of the presented tree".

When Tye and Harman instruct us to reflect upon what it is like to enjoy some particular episode of conscious experience, they claim that two things are revealed. First, that one of the things that one does in order to answer the question 'What, subjectively, is our experience like?' is to consider what the ordinary mind-independent objects of awareness are, and what features and properties they have (something that will evidently involve

[64] Cf. the *OED* entry for 'attention' and James' remark: "Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought" (1890: 403-4).

deploying perceptual attention). Second, that one can do nothing else in pursuing this question. No aspect of one's answer will involve reference to intrinsic features of experience or, more generally, features which are not aspects of the presented scene.

One can, of course, think about all sorts of aspects of one's experience (its possible neural correlate, whether it will prompt a dream tonight, whose birthday it might be occurring on) that are nothing to do with its objects. This is why we need to focus on answering the question, 'What, subjectively, is it like to be enjoying your current experience?' as opposed to 'What properties does your experience have?' We also need to focus on rationally appropriate strategies for answering this question. One might think that there are other ways of answering the question that do not go via making judgements about the world apparently presented, such as saying the first thing that pops into your head. This hardly generates counter-examples to transparency.

In the terms of earlier discussion we can put Harman and Tye's conclusion as follows: when one attempts to describe the *experiential aspects* of one's experience, one (a) rationally does and (b) can only rationally do so through making judgements about and, since those judgements typically require it, perceptually attending to the ordinary, apparent objects of experience and their properties. This helps makes explicit a distinction between positive and negative aspects of the transparency claim. It might be that some experiential properties can be attributed on a number of distinct grounds. As a result, the fact that it is rational to attribute an experiential property via attention to the objects of experience does not rule out distinct grounds for attribution. Harman and Tye evidently hold that such distinct grounds are ruled out – consideration of the presented scene is the *only* rational way of attributing properties to experience (on the basis of reflection alone).

We also need to flag a second key distinction between what we might call weak and strong transparency theses. It is helpful first to state what it is for a particular aspect of experience to be transparent according to Tye and Harman.[65]

> **Property Transparency** A phenomenal property is transparent iff for any particular experience with that phenomenal property, it is never the case that

[65] I here include the negative claim discussed in the last paragraph as part of the definition of transparency.

> when I consider how things are with me experientially, I am rationally positioned, (through reflection on my experience alone) to ascribe that property to my experience other than on the basis of a judgement concerning the apparent objects of perception or their features.

Strong transparency is the claim that *all* phenomenal properties are transparent in this sense; weak transparency is the claim that *some* are. In other words, weak transparency tells us that when we reflect upon our experience in an attempt to answer the question, 'What, subjectively, is our experience like? What are its experiential properties?' the rational answer to that question about our experience is *in part* and essentially determined by the rational answer to a question concerning the apparent objects (or content) of that experience, viz., 'What is the scene presented like?' Strong transparency tells us that when we reflect upon our experience in an attempt to answer the question, 'What, subjectively, is our experience like? What are its experiential properties?' the rational answer to that question about our experience is *wholly* and essentially determined by the rational answer to a question concerning the apparent objects (or content) of that experience, viz., 'What is the scene presented like?'[66]

Discussing transparency in terms of particular properties allows us to consider any particular experiential aspect of experience and ask whether *that* aspect is transparent, whilst holding in abeyance considerations regarding strong as opposed to weak general theses. All we need to consider is whether judgements attributing *that* property to experience are determined essentially by judgements concerning the apparent objects of perception.

Reflection reveals that *some* properties are transparent in the sense of Property Transparency. When we reflect on the nature of our experience, we do so by making judgements concerning the (apparent) objects of perception. This observation was made forcefully by Strawson when he pointed out that, when asked to describe their experience, naïve subjects do so in terms of ordinary, apparent objects of perception and their features. Thus, as already mentioned, his imagined subject reports his visual experience by saying that he sees "the red light of the setting sun filtering through the

---

[66] Cf. Way's discussion of a different notion of transparency in his (2007: 223-4).

black and thickly clustered branches of the elms … [and] the dappled deer grazing in groups on the vivid green grass…" (1979: 43). This answering strategy will inevitably involve perceptual attention to the objects of experience – the more precise and detailed an answer we wish to give, the closer our attention (cf. Martin 1998b where he quotes the art historian Michael Baxandall). Moreover, with respect to such an aspect of experience this does appear to be the only strategy that we have for attributing this property to experience. Such a strategy is essential for attribution. Consequently, properties such as 'being a presentation as of a group of grazing deer' are transparent in the sense of Property Transparency. Henceforth that is what I mean by 'transparent,' unless the usage is otherwise qualified.

Strawsonian considerations justify an existential claim, the claim that we *do* attribute properties to our experience in certain ways, namely, via describing the world encountered. This thought alone has been held to have consequences for how we should understand the nature of perceptual experience.[67] But endorsing Strawson's claim does not commit one to strong transparency, viz., the claim that *all* experiential properties are transparent in the above sense, i.e., that there is no way of attributing experiential properties other than on the basis of making judgements about the represented scene.

Two possibilities face us at this point. On the one hand, Tye and Harman might be right to defend a strong thesis – not only does one describe one's experience through attention to and description of the encountered world, one cannot rationally make a judgement concerning the presence of an experiential feature on any other basis. Or, to paraphrase Tye, we can only rationally distinguish our experience from other experiences (in their experiential aspects) on the basis of differences in apparent objects and their apparent properties. Alternatively, one might think that there are ways in which we *can* rationally distinguish amongst our experiences in reflective awareness that do *not* turn solely on differences in the presented objects and properties of objects – differences in how things seem, despite sameness of presented objects and their properties. One issue here is whether introspection reveals not just properties represented or presented by experience, but also non-representational, subjective properties of experience. This is not the only way of construing the issue. A much-discussed example of difference in how things seem

[67] In particular, it places an explanatory burden on any potential theory of experience that certain theories struggle to discharge. See, for example, Martin 1998b, 2002.

despite sameness of presented objects and their properties is the example of blurry vision (see, e.g., Crane 2006: 130 and Pace 2007). We might well think of blurriness as a matter of the *way* objects and their properties are presented, where that need not be a matter of subjective qualia or the like.

Such concerns are not crucial to the argument of this chapter. For present purposes, it should simply be noted that Tye and Harman's strong thesis is much less obvious than the weak thesis. As just suggested, it may well be that the *way* in which a given object or property is perceived is an important part of the characterisation of what it is like to have an experience. Moreover, as Martin (2002: fn.11) points out, claiming that *all* experiential properties are transparent looks unlikely to be justifiable solely by appeal to a handful of examples (especially given the limited menu from which they are chosen).

Construing transparency as a thesis about particular properties of experience avoids these issues. It allows us to ask of any particular aspect of experience whether *it* is transparent. (I am, of course, specifically interested in experience's temporal properties.) This enables us to leave aside particular potential counter-examples to a general thesis. Moreover, as we have seen, this way of approaching things does not deprive the transparency claim of real force. If any aspects are transparent, a purely subjective view will, arguably, struggle to account for them. Conversely, we can lay down a challenge to those who believe that experience has subjective properties (sense-data, qualia): *find* a phenomenal property of experience – an aspect of how things seem – that does not conform to transparency in this sense.

One important issue remains. The claim that a property is transparent in the sense stated above is deliberately neutral with respect to what one might think of as the *metaphysics* of transparency.[68] To say that some experiential property is attributable to experience only on the basis of judgements concerning the represented or apparently presented scene is at most to establish that the phenomenal property attributed supervenes on that aspect of the experience's representational content, or the presentation of that particular scene. More generally, epistemic claims about how we attribute properties to experience can at most establish that there cannot be differences in phenomenal character without a

[68] It is possible that such neutrality is ultimately unsustainable but for present purposes I bracket that large issue.

difference in representational content/presented objects. This relationship is neutral on whether representational content/presented scene *determines* phenomenal character or vice-versa. Transparency theorists are inclined to discuss transparency in terms of determination. Harman and Tye want us to think of phenomenal character as determined by representational content, or, in less partisan terms, that the phenomenal properties of experience are determined by how the scene is presented or represented as being. Certainly, such a determination claim would explain transparency in the above sense. But so might alternative accounts.

Chalmers is sensitive to these issues when he writes the following.

> The plausible thesis that perceptual experiences have phenomenal content leaves many other questions open. For example, the thesis is neutral on whether phenomenal character is prior (in some sense) to representational content, or vice versa. It is compatible with the thesis that phenomenal character is grounded in representational content (as held by Dretske 1995 and Tye 1995, among others), and it is compatible with the thesis that representational content is grounded in phenomenal character (as held by Searle 1990 and Horgan and Tienson 2002, among others). (2006: 51)

I see no reason for thinking that Searle, Horgan and Tienson are straightforwardly refuted by transparency considerations. It may even be that some phenomenal properties are determined by aspects of the presented or represented scene, whereas for others the determination runs the other way around. For the main part, I do not take a stand regarding this metaphysical debate.

5. Time and Transparency

Debates about the transparency of experience have been dominated by examples of the visual perception of spatial and colour properties: seeing a blue square (Tye 1995: 30), or delighting in "the intense blue of the Pacific Ocean" (Tye 1992: 160); staring at "the variegated colours and shapes of [a lavender bush's] leaves and branches" (Martin 2002: 380); or, more mundanely, one's "experience of looking at a gray filing cabinet" (Stoljar 2004: 341).

Yet we also perceive temporal properties, as well as properties which are determinations of time. Temporal properties form part of the content of experience. There is a striking difference between perceiving temporal properties and cases of seeing shape and colour properties. The difference is due to the fact already much emphasised that unlike shape or colour, temporal structure is common to experience and its objects. Experience *itself* in its experiential aspect has a genuine temporal ordering – our streams of consciousness are made up of events with relative durations, which occur determinately before and after one another. Experience is not, in its experiential aspect, coloured or shaped. Thus, time is uniquely common to experience and its objects. This introduces an additional element into discussion of transparency claims.

My main aim in this section is to argue that the temporal properties of our experience such as succession and relative duration (the temporal structure of our stream of perceptual consciousness) are transparent in the quite general sense outlined above. Doing so will highlight the ways in which the transparency thesis looks rather different in the temporal case. As we saw above, Harman and Tye think that transparency is something that can be established by introspection. In particular, they think that, when we attempt to describe our experience, it turns out that our only rational way of doing so is by making a judgement about the apparent objects of experience and their apparent properties. So let us focus on a case where we set out to describe the temporal structure of our experience.

Imagine you hear two notes played on a piano, first a *G* above middle *C* and then the *C* a fourth above it. For example, think of the opening notes of the Rondo in Beethoven's *C* minor piano sonata op.13 (*Pathétique*) or (in a different key) the wedding march. Attending to the playing that is going on in the world, you judge the sound of the *G* to have occurred before the *C*. You note the temporal relation between the two sounds, how fast the *C* followed on from the *G*, you begin to have a sense of tempo, rhythm etc.

Now imagine instead attending to – that is reflecting on the nature of – one's experience as the notes are played. In particular, imagine reflecting on the temporal ordering of one's experience of the notes. More precisely, imagine attending with the aim of answering the question, 'What temporal relation does my experience of the *C* bear to my experience of the *G*?'

Considering the problem entirely in the abstract, the following possibility seems quite consistent with the scenario as sketched so far: when one comes to reflect on one's experience, one rationally judges that one's *experience* of the *C* came before one's *experience* of the *G*. Moreover, one judges that the time between the experience of a *C* and a *G* was significantly shorter than the time one judged there to be between the *C*'s sounding and the *G*'s sounding. The situation, theoretically hypothesized, is graphically depicted below.

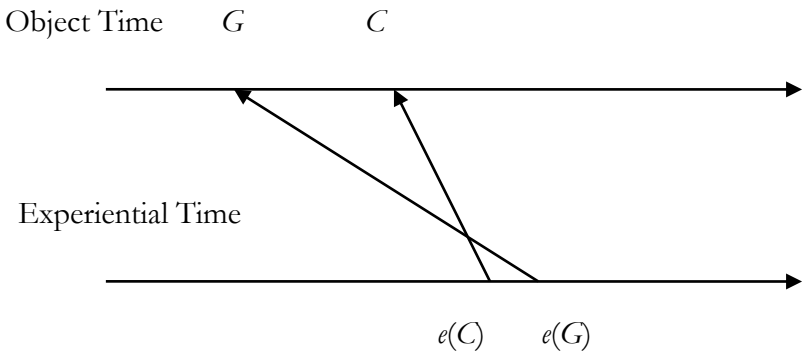

However, I submit, when we actually engage in reflection on our experience, this theoretical possibility is revealed to be incoherent.[69] The way we go about trying to answer the question concerning the temporal structure of our conscious experience is by making a judgement about the temporal structure of the apparent objects of consciousness (the sounds heard), *and then by taking our experience to be structured in the same way*. This is what explains the practical difficulty we have in trying to make sense of the idea that the objects could be represented as having a different ordering, or bearing different relations to each other (let alone such dramatically different orderings and relations) to the ordering and intra-relations of representing experience. When we reflect upon the nature of our experience, the musical scenario sketched above seems quite unimaginable, quite incoherent. It simply does not seem that, reflecting on our experience alone, we could come to rationally judge that our experience had a different ordering to the apparent ordering of the objects of experience. If my experience

[69] Note the implications for the status of the transparency claim. It is a claim about how perceptual experience is for us. It does not straight-forwardly rule out the existence of a radically different kind of perceptual experience which does not conform to transparency.

represents the *C* as succeeding on from the *G*, then we will rationally judge that my experience of the *G* came before my experience of the *C*.

The apparent matching here is a direct consequence of the transparency of temporal properties. As understood above, an experiential property is transparent just if its attribution essentially goes via our reflection on the apparent objects/content of experience. The attribution of temporal structure to our experience necessarily goes via judgements about (and typically attention to) the apparent temporal ordering of the objects of experience. Thus, temporal properties of experience are transparent. It is this transparency which explains why, restricting ourselves to how things seem to us on introspective reflection, we can make no sense of a gap between the apparent ordering of objects and the apparent ordering of experience itself. Call this claim Temporal Transparency.

> **Temporal Transparency** For any particular temporal determination of experience, it is never the case that when I consider how things are with me experientially, I am rationally positioned, through reflection on my experience alone, to judge my experience to be so determined other than on the basis of a judgement concerning (and typically perceptual attention to) the temporal structure of the apparent objects of perception.

6. Consequences

6.1 Time and the Metaphysics of Transparency

Sometimes it is suggested that the claim that all phenomenal properties are transparent establishes that there is no more to experience's phenomenal character than its representational content, no more to the conscious aspects of experience than *what is represented*.[70] Witness the following statements of this view.

> [T]here is nothing to your experience over and above the way it represents the world as being. (Carruthers 2005a: 40)

[70] Recall the discussion of Pylyshyn in Chapter One.

> [T]here *are no* non-relational properties of experience *qua* experience. (ibid.: 47)

> If we stipulate that something is a visual phenomenal quality or a quale only if it is a directly accessible quality of an experience, then there are no visual phenomenal qualities or qualia. Still there are qualities of which the subjects of visual experiences are directly aware via introspection. They are qualities of external surfaces (and volumes and films), if they are qualities of anything. (Tye 2002: 141-2)

> [We] are not aware of those intrinsic features of [our] experience by virtue of which it has [its] content. Indeed, I believe that [we have] no access at all to the intrinsic features of our mental representation that make it a mental representation of seeing a tree [or of another scene for that matter]. (Harman 1990: 39)

Finally, Byrne argues on the same basis that we should not be concerned with the mind-body problem, urging a move to a fictitious world w´ where the mind-body problem isn't taken seriously. Why is it not (and should it not be) taken seriously there?

> Well, the philosophers in w´ are much taken with the idea that experience is *transparent*. Introspection of one's experience of blue, for example, merely yields what the experience is *of* or *about* – the ostensible scene before the eyes. As to the intrinsic nature of the experience, we are completely in the dark. If we like, we can say experiences of blue have a "qualitative character", but that is simply because they *represent* that objects have a "qualitative" property – namely, blueness. The experiences are, in this respect, like the *words* 'blue', 'purple', 'yellow', and so forth. We may say that 'blue' is more similar in a salient qualitative respect to 'purple' than to 'yellow', but that can only mean that 'blue' represents a property that is more similar in a salient qualitative respect to the property represented by 'purple' than it is to the property represented by 'yellow'. Hence, the philosophers in w´ [think that], if we can provide a satisfying naturalistic explanation of the qualitative nature of the *colours*, there will be no mysterious qualitative residue left in *experience*. (2006: 223-4)

The reason that this conclusion seems so natural to the authors above is that they focus on properties such as colour and shape. The colours and shapes we experience are experienced as the properties of objects in the world and, crucially, when we reflect on our experience we do not find our experience itself to be blue or shaped – if that idea

even makes sense – nor do we find any objects other than the objects in the world to attach such properties to. Thus, in so far as we are only aware of such colour and shape properties, we are not aware of intrinsic, non-representational properties of experience. Consequently, the negative transparency claim seems to rule out introspective knowledge of any intrinsic, non-representational properties. As Byrne puts it, "As to the intrinsic nature of the experience, we are completely in the dark."[71]

The case of time shows that this is a mistake. For in the case of experience, we precisely *do* think of our experience itself as having a temporal structure of its own, as having intrinsic temporal properties that we can come to be aware of (i.e., know about). This is quite compatible with such properties being transparent in the above sense. The temporal properties of experience are intrinsic, non-representational properties of experience and yet also transparent.

Nonetheless, the fundamental concern of the theorists above is to establish that representational properties *fully determine* phenomenal character. That might still be so in the temporal case. However, even here, Strong Transparency does not establish this metaphysical thesis. As I suggested above, to say that a property is transparent is non-committal with regard to the metaphysical issue concerning the determination of phenomenal character by representational content or vice-versa. In the case of time the key issue can be stated as follows. Are the temporal properties of experience determined by the representational content of experience, or is it in virtue of having a certain temporal structure that experience represents the world as being temporally structured? Put another way: is the temporal structure of experience determined by the temporal structure of the apparently presented scene, or is the temporal structure that the apparent scene is presented as having determined by the temporal structure of experience?

In arguing that there are temporal constraints on the representation of time, I have not taken a stance on this issue. Our judgements (and plausibly knowledge) as to what that

---

[71] This discussion should not be assumed to build into it the idea that we do find representational properties as occasionally occurs, for example in Shoemaker's remark that "qualia, if there are such, are diaphanous; if one tries to attend to them, all one finds is the representative content of the experience" (1990: 101). Here, as with Tye and Harman on occasion, the positive claim that one does find representational content is built in to the transparency thesis. This is misleading. Rather, as Martin (2002) argues, a strong transparency thesis presents the theorist with an explanatory burden which the representationalist *but also the naïve realist* is prima facie in a position to discharge.

temporal structure *is* are dependent upon, indeed always go via judgements about the apparent temporal ordering of the objects presented. This is quite compatible with thinking that experience represents the world as being temporally structured in a certain way in virtue of having a certain temporal structure of its own.

6.2 Transparency and the Principle of Presentational Concurrence

In discussing the views of Brentano and Husserl, Miller introduces the following principle which he calls the *Principle of Presentational Concurrence* or PPC.

> [T]he duration of a *content* being presented is *concurrent* with the duration of the *act* of presenting it. That is, the time interval occupied by a content which is before the mind is the very same time interval which is occupied by the act of presenting that very content before the mind. (1984: 107)

Dainton (2000: 134) has recently endorsed this principle.[72] As noted above, Tye singles out PPC as a clear case where the properties of the represented are being confused with the properties of the representation, no different to assuming that, if a red light flashes before a green light, then the verbal representation of said events must be structured so that the representation of the red light comes before the representation of the green light, ruling out a representation such as, 'The green light flashed after the red light.'

The discussion of Temporal Transparency suggests that those who endorse PPC are not confused. Although they do not make explicit appeal to introspection, the view that the properties of experience must match the properties of what is apparently experienced in the temporal case does have introspective support. Thus, they cannot be charged with making a simple act/object confusion.

Having said that, a gap remains between PPC and the claim that the temporal properties of experience are transparent. Just because our only rational strategy for attributing

---

[72] In this Dainton follows Foster who holds "we have to take experience to extend over a period of real time in a way which exactly matches the phenomenal period it presents" (1991: 249). Dainton ultimately denies the act-object distinction upon which PPC is premised. However, as noted above, he declares that "even if we draw an awareness-content distinction it makes no sense to suppose that an act of awareness can apprehend a content of greater temporal duration than itself" (2000: 180).

temporal structure to experience commits us to judgements which accord with PPC, we have not yet provided grounds for thinking that our rational judgements are always correct. We may be misled about the nature of our inner lives. The burden of the next chapter is to close this gap, and thereby defend PPC.

Relatedly, and as already noted, I have yet to discuss the most striking arguments against PPC, namely, those found in Dennett's work. Dennett agrees that PPC is an overwhelmingly natural view to hold. Indeed, having argued against it at length, he considers how we might react.

> You may still want to object that all the arguments in this chapter are powerless to overturn *the obvious truth that our experiences of events occur in the very same order as we experience them to occur*. … [This is] a thesis that is true in general, and does indeed seem unexceptioned so long as we restrict our attention to psychological phenomena of "ordinary" macroscopic duration. (1991a: 168, my emphasis)

However, on the basis of short-term phenomena, Dennett notoriously rejects this picture. Quite how to understand Dennett's view is a complex issue which I take up at some length in Chapter Five. However, two points should be noted. Firstly, Dennett's view is not obviously one on which we are misled about the ordering of conscious experience when we engage in introspective reflection. Dennett is arguably better understood as endorsing an anti-realist view, according to which there is no more to the temporal ordering of conscious experience than the way we represent that order in short-term autobiographical memory. If that is right, Dennett may be in a position to endorse Temporal Transparency under an anti-realist construal. In other words, Dennett is arguably best understood as attacking Realism not Temporal Transparency.

Secondly, insofar as critics of PPC like Tye rely on Dennettian arguments, as opposed to dubious analogies, they certainly do not endorse other central tenets of Dennett's view. Tye (1993), for example, offers a robust defence of what he calls Phenomenal Realism against Dennett's attack. As noted above, one might well wonder whether Tye can really have his cake and eat it in this regard.

Chapter Three:

Self-Intimation

It must be possible for the 'I think' to accompany all my representations; for otherwise something would be represented in me which could not be thought at all, and that is equivalent to saying that the representation would be impossible, or at least would be nothing to me.[73]

1. Introduction

What is the relation between consciousness, and self-consciousness or self-awareness? Why is answering this question crucial to issues concerning experience and time?

We have already seen one reason. Chapters One and Two argued that experience has an experiential temporal structure of its own and that that structure is *transparent* in that it is rational to judge that one's experience is temporally determined in a given way only by taking its temporal structure to mirror the apparent temporal structure of the world experienced. However, rational judgements need not be true, less still need rational beliefs count as knowledge.[74] As a result, nothing that has been said so far rules out the sceptical possibility that our introspective judgements in this (or any other) domain are systematically mistaken. Nothing rules out the thought that, although we may be rational to judge that our experience is structured in one way or another, such judgements about the temporal structure of our experience regularly, even always, fail to be true.

---

[73] Kant 2003: B131-2.

[74] It is rational to trust news reports, teachers, encyclopaedia and maps (for example) even though sometimes the resulting beliefs are false, or even if true, do not count as knowledge for reasons beyond our ken. Certain maps contain 'copyright traps,' fictional post offices, streets, or even towns designed to catch plagiarism. Likewise, notoriously, the *New Grove Dictionary of Music and Musicians* (Sadie 1980) contains an entry concerning an entirely fictitious Italian composer. More pedestrian and common-place are the straightforward slips we all (journalists, teachers and editors included) make, mistakes which are occasionally accidentally right.

Perhaps our experience seems some way to us and yet is in fact some other way entirely. At the extreme, perhaps our apparent awareness of our experience as a continuous stream of overlapping and successive processes and/or events is only apparent. Perhaps it only seems that way – perhaps our experience is in fact a series of discrete and independent, instantaneous events.[75] Perhaps this supposed fact is something that we are constitutionally incapable of achieving awareness of. One aim of this chapter is to cast doubt on such scepticism.

Consideration of this sort of scepticism also anticipates issues which occupy centre stage when it comes to considering Dennett's influential views in Chapter Five. Dennett's arguments concerning time and experience are often dismissed as verificationist.[76] Yet as Dennett writes,

> Some thinkers have their faces set so hard against "verificationism" and "operationalism" that they want to deny it even in the one arena where it makes manifest good sense: the realm of subjectivity. (1991a: 132)[77]

This chapter aims to explore what does make manifest good sense when it comes to the realm of subjectivity. As such, the considerations here underpin much of the discussion in Chapter Five.

My strategy is not to present a decisive case in favour of one particular way of understanding the relation between consciousness and self-awareness. Rather I motivate and present a natural way of thinking about the relation and argue that extant reasons for rejecting it are unsuccessful. Thus, in the next two sections I sketch in broad terms two influential lines of thought regarding self-awareness and consciousness both of which drive us to the same conclusion, viz., a claim I call Self-Intimation. In subsequent

---

[75] A possibility suggested by Austen Clark (1989: 291), for example.

[76] For example, Tye comments, "I am inclined to reply [to Dennett] by modifying a remark of Mr. McCawber: "Verificationism! Foul play, sir! Take a drop more grog and you'll get over the weakness of believing in verificationism." (1993: 894)

[77] Cf. Hurley 1992: 212, "To claim that a difference in the content of conscious experience must be a consciously experienceable difference is not to subscribe to verificationism, of which some may be tempted to accuse D & K." Likewise, Searle (1995), who echoing Dennett declares, " … where the existence of conscious states is concerned, you can't make the distinction between appearance and reality, because the existence of the appearance is the reality in question. If it consciously seems to me that I am conscious, then I am conscious".

sections I turn to standard objections to the picture adumbrated. Responding to these objections helps clarify and flesh-out the proposed view.

2. Scepticism and Consciousness

It is commonly held that knowledge of our inner, conscious lives is somehow sceptic proof. Such knowledge brings the Cartesian sceptic's remorseless undermining of knowledge claims to a halt. We can begin to think about what exactly this amounts to by considering what could explain such a special epistemic position with respect to our conscious lives.

One might think that some especially tight and reliable connection between second-order beliefs and conscious experience could do the trick. To this idea Chalmers objects as follows.

> The knowledge that a reliability theory grants us seems too weak to count as the kind of knowledge that we have of our conscious experience. … The trouble is that if our beliefs about consciousness were justified *only* by a reliable connection, then we could not be *certain* that we are conscious. The mere existence of a reliable connection cannot deliver certainty, for we have no way to rule out the possibility that the reliable connection is absent and that there is no consciousness at the other end. (1996: 194)

Though I entirely agree that the reliabilist move is unsatisfying, it is not obvious that *certainty* is the right place to object to the reliabilist proposal. Certainty is not enough to explain the special status of self-knowledge of conscious episodes. After all, Chalmers invokes a specific principle about certainty in the above passage viz., that one cannot be certain in virtue of a reliable connection unless one can rule out the possible absence of the reliable connection. The reliabilist has two options at this point. She can simply reject the principle. Alternatively, she can claim that in the good case where the reliable connection does hold, a subject *can* come to know the relevant fact and in so knowing rule out the absence of the reliable connection. Either way, it seems that the reliabilist can make room for knowledge and certainty in the good case.

Nonetheless, the key Cartesian intuition noted at the start does conflict with the reliabilist idea. Indeed, Chalmers puts his finger on the crucial point as follows.

> In the case of perceptual knowledge, for example, one can construct a case … where the subject is a brain in the vat, say – and everything will still seem the same to the subject. … But in the case of consciousness, one cannot construct these sceptical hypotheses. … There is no situation in which everything seems just the same to us but in which we are not conscious … (ibid.: 195)

In other words, it is not that the reliabilist cannot account for knowledge in the good case. Rather it is that they must allow for bad cases. But there cannot be bad cases in this realm.[78]

This is not a problem just for the reliabilist. As Chalmers notes the same point affects a causal account of knowledge. We can be more general still – the problem is with any account of introspective knowledge which models it on perceptual knowledge. For any account of perceptual knowledge, however direct or immediate such knowledge is claimed to be in good cases, must allow for bad cases – cases of illusion or hallucination where things are not as they seem.[79] Yet that is precisely what we cannot make sense of in the inner arena.

Shoemaker spells out the feature of perceptual knowledge which must not be translated to the introspective case as follows.

> [I]n perception we have access to things or states of affairs that exist independently of their being perceived and independently of there being any means of perceiving them. (1994: Lecture I, 252)

---

[78] In the same passage Chalmers says that, "Our core epistemic situation already *includes* our conscious experience" (195). That is right. But note that the naïve realist can claim that our core epistemic situation includes objects of awareness in the external world in the perceptual case. Thus, it is the denial of bad cases which is crucial and not such claims about the good case. On this point see Martin 2006: 388.

[79] As Martin puts it, "all that need be denied is that we can make sense of a subject's situation being this way: describing how things seem or are from a subject's point of view characterizes her phenomenal consciousness one way; attending to how things really are, requires that we describe it another way" (2006: 389). Such a possibility is evidently present in the case of perceptual belief.

The core mistake of perceptually modelled accounts of introspection, according to Shoemaker then, is to attribute this independence to self-conscious awareness.

> [T]he [broadly perceptual] view [of introspection is] that the existence of mental entities and mental facts is, logically speaking, as independent of our knowing about them introspectively, and of there being whatever means or mechanisms enable us to know about them introspectively, as the existence of physical entities and physical facts is of our knowing about them perceptually, or of there being the means or mechanisms that enable us to have perceptual knowledge of them. (1994: Lecture II, 271)

In distilling the central failing of perceptual models, we need not commit ourselves or the perceptual theorist (as Shoemaker implicitly does) to the existence of *mechanisms* that make knowledge of consciousness possible. Strictly, all we need is Shoemaker's idea that, on the perceptual view, the existence of perceptual experiences is (a) independent of their being known and (b) independent of their subject's being in a position to know them. Moreover, since being in a position to know is strictly weaker than knowing, we can drop the first conjunct of the condition.[80] Thus, in rejecting this condition, we finally reach the idea that what is fundamentally correct about the Cartesian intuition is that a subject's undergoing experience of some experiential kind is not independent of her being in a position to know that she is undergoing experience of that kind, is not independent of the availability of that experience to self-consciousness. Or again: there is no situation in which a subject can be truly said to be undergoing some course of perceptual experience of a given experiential kind and yet that subject not be in a position to know that they are enjoying experience of that type.[81]

This echoes Chalmers' insistence on the "fact that there is no way to construct a sceptical scenario in which I am in a qualitatively equivalent epistemic position, but in which my experience experiences are *radically* different" (196; my emphasis). Yet here note Chalmers' caution revealed by the word 'radical'. Such caution is hard to justify once the

---

[80] It is not clear why Shoemaker includes it in any case since he happily acknowledges that people may have an experience of a certain kind and fail to notice it, remarking that he is concerned not with not-noticing but with the impossibility of noticing. See his 1994: Lecture II, 276 n.1.

[81] As Shoemaker puts it, "it is of the essence of various kind of mental states that they are introspectively accessible" (Lecture III, 291); "there is a conceptual, constitutive connection between the existence of certain sorts of mental entities [amongst which are perceptual experiences] and their introspective accessibility" (Lecture II, 272). See further Shoemaker's argument against 'self-blindness' (276-8).

central point is granted that consciousness does not provide room for a gap between how things seem from a subject's point of view and how they really are. Given that, it is very hard to see what could motivate anything less than the stronger claim that there is no way to construct a sceptical scenario in which I am in a qualitatively equivalent epistemic position but in which my experience differs *at all* (radically or not). If, as Chalmers holds, "there is something intrinsically epistemic about experience" (ibid.), how could there be some experiential aspects which were not intrinsically epistemic? Or, in Shoemaker's terms, how could we be self-blind to *any* aspects of our conscious lives? How could there be *any* aspect which was independent of our epistemic position?

The discussion so far puts us in a position to state the central claim to be defended in this chapter.

> **Self-Intimation** If a subject is undergoing perceptual experience of a certain experiential kind, then that subject is in a position to know that they are undergoing perceptual experience of that experiential kind simply in virtue of so undergoing.

Self-Intimation tells us that there cannot be conscious, experiential properties which are simply beyond the ken of their subject. Note that the asymmetric dependency adverted to by stating the thesis in terms of what subjects are said to be in a position to know *in virtue of* their experience makes the formulation of Self-Intimation quite explicitly realist. Experience is assumed to be explanatory of a subject's epistemic position. In Chapter Five I consider views of phenomenology which reject this picture and endorse some form of anti-realism about experience. For now the explicit realism is embraced.

Chalmers raises a prima facie worry about this picture as follows.

> This raises a question. What is it that justifies our beliefs about our experiences, if it is not a causal link to those experiences, and if it is not the mechanisms by which the beliefs are formed? (ibid.: 196)

He also provides what I take to be just the right answer: "it is *having* the experiences that justifies the beliefs" (196).[82] Dretske (2003) has recently shown what happens if you reject this natural answer. Dretske assumes that if we know something, there must be some *way* we know it and then worries that in the absence of an inner sense model of some kind that we have no *way* of knowing about our own conscious lives.[83] His remarkable conclusion is that perhaps we do not in fact know that we are conscious, that for all we know we are zombies! As he puts it,

> We are left, then, with our original question: how do you know you are not a zombie? … To insist that we know it despite there being no identifiable *way* we know it is not very helpful. We can't do epistemology by stamping our feet. Sceptical suspicions are, I think, rightly aroused by this result. Maybe our conviction that we know, in a direct and authoritative way, that we are conscious is simply a confusion of what we are aware of with our awareness of it. (2003: 9)

The right response to this is to insist that my way of finding out that I am undergoing experience of some kind can on occasion simply be to undergo that experience – as Self-Intimation insists, the having of that experience is not independent from our being in a position to know that fact. Perhaps Dretske tacitly subscribes to a principle held by traditional sense-datum theorists, viz., the constraint on knowledge that its objects must be logically independent of it.[84] Read in a certain way that might support Dretske's view. But in the absence of a defence of some such reading of that eminently questionable principle, there is no reason to deny that conscious experience is *itself* the epistemological ground and source of our awareness of it.

---

[82] Some writers on self-knowledge (e.g. Wright 1998) encourage the idea that because questions such as 'How do you know?' are supposedly "always inappropriate" in response to claims such as 'I am having an auditory experience as of a piercing shriek,' that there is nothing on which such claims are based. We should insist that the avowal is grounded in the experience itself and that the correct reply to the question is simply: 'Because I'm undergoing an experience of a piercing shriek'.

[83] Dretske notes, "The only way in which I am aware of myself seeing an ant is in the sense of being aware that I see an ant, but this, the awareness of the fact that I see an ant, is not my way of finding out I see an ant. It is a restatement (using the words 'aware that…') of the fact that I know I see an ant. If one fails to distinguish, in this way, the facts we are aware of from the objects (events) we are aware of, one will mistakenly suppose that our undisputed (…) awareness that we experience things is an answer to the question about how we know we experience them. This [would be] double dipping under the cloak of an equivocation: citing an awareness that *p* as one's way of becoming aware that *p*" (2003: 8).

[84] For example, in a discussion of Kemp Smith on Berkeley, Prichard writes, "For we all, including Professor Kemp Smith, think that the existence of what we know is independent of our knowledge of it, on the ground that otherwise our knowledge of it would not be knowledge" (1938: 204).

In the next section I consider a distinct line of thought which leads us to the same conclusion. Once again, I do not claim that the considerations are decisive. They simply highlight a very natural way of thinking about consciousness and self-consciousness.

3. Report, Judgement and Consciousness

In his paper 'Explaining Consciousness' Rosenthal claims that

> Whatever else we may discover about consciousness, it's clear that, if one is totally unaware of some mental state, that state is not a conscious state. (2002: 407)

Of course, Rosenthal's claim strikes many people as too strong and, I hazard, strikes most as far from clear. Cases of absent-minded truck drivers wholly unaware that they are consciously perceiving and responding to passing traffic do not seem obviously incoherent. However, what does seem incredibly difficult to conceive is a mental state or episode that is conscious and yet which is *in no way available* to self-awareness. We find it hard to credit the driver with consciousness of passing traffic if he is simply incapable of becoming aware of that fact.[85] Indeed, if one understands consciousness in terms of there being something it is like, subjectively, for you to be in a state/undergo an episode, it seems very hard to understand how there could be something it is like, subjectively, for you to be in that state/undergo that episode and yet that state/episode be irrelevant to how your mental life might even potentially seem from your own perspective. How would that be contributing to what it was like *for you*? What would the *for-you-ness* consist in?

Just as Rosenthal's statement of what is uncontroversial in this area over-reaches, so too does his motivation for his position. However, by considering what conclusion his motivating argument does support, one can motivate the weaker but more compelling view sketched above.

---

[85] Compare the case of bodily action. It is controversial whether unwittingly jiggling your leg up and down counts as a voluntary, intentional action. But it seems to be uncontroversial that if one has no direct control at all over one's leg and cannot stop or even modulate its movement, then it does not so count.

Rosenthal (1990) defends the claim that a mental state or event is conscious just if it is the intentional object of a (roughly) simultaneous higher-order thought formed neither by inference nor observation and having the content that one is undergoing just such an experience/in just such a state.[86] His basic argument for this view he calls 'The Argument from Reporting and Expressing' (1990: 736-7). It takes the form of an inference to the best explanation.

(i) "Reporting one's mental states is distinct from verbally expressing those states" (747). If I believe The Magpies will win the Grand Final I can *report* my belief by saying 'I believe the 'pies will win the Grand Final' or I can *express* that belief by saying simply 'The 'pies will win the Grand Final'.

(ii) "Reporting … has a special connection to consciousness. If we restrict attention to creatures with the relevant linguistic ability, a mental state is conscious if, and only if, one can report being in that state." (ibid.)

(iii) If I sincerely, and with understanding, report that *p*, I express a propositional mental state of mine with assertoric content that *p*.

(iv) Thus, "When I report being in a mental state … my speech act expresses my [higher-order thought] that I am, myself, in that very mental state."

(v) Consequently, a mental state of mine is conscious just in case I can express the relevant higher-order thought that I am in such a state.

(vi) Finally, "The best explanation of our ability to express these [higher-order thoughts] about all our conscious states is that our conscious states are actually accompanied by such [higher-order thoughts]. Similarly, we can best explain our inability to report mental states when they are not conscious by assuming we lack the relevant [higher-order thoughts]."

Note the following about this argument. Firstly, there are in fact two purported explananda: (a) that we can always report conscious states/episodes and (b) that we can never report non-conscious states/episodes. If one is sceptical that consciousness is a necessary condition for reportability, one might still wish to hold onto the thought that (in the linguistically able) it was sufficient. That fact would still require explanation. This is fortunate since it seems highly doubtful whether consciousness is necessary for

[86] As Rosenthal puts it, "a mental state's being conscious consists simply in its being accompanied by such a higher-order thought" (2005: 56).

reportability. I often know what I believe or desire without there being anything it is like for me to have the relevant belief or desire. I know that I believe that Picasso is a much greater artist than Pollock but the belief reported by the statement that I believe Picasso is a much greater artist than Pollock doesn't obviously need to be a conscious belief. Is there always something it is *like* to have that belief? Similarly, is there always something it is like, subjectively, to want one's steak *saignant* rather than *bien cuit*? If not, does that prevent me informing the waiter?

Let us focus then on the more plausible direction of the argument. As it stands it is patently too strong. For though Rosenthal claims that it "is unclear how one could have the ability to express some particular thought without actually having that thought" (1993: 204), as Byrne points out, "it isn't at all unclear."

> For it may simply be that in virtue of being in [some conscious mental state] *S* I have these two abilities: the ability to acquire a higher-order thought about *S*, and the ability to express it, when I have acquired it. One might sum up these facts by saying that in virtue of being in *S* I have the ability to express the higher-order thought that I am in S. And plainly I can have this composite ability without having the higher-order thought whenever I am in *S*. (1997: 109)

However, assuming that we accept the point about reportability – viz., premise (ii) of the argument as read right to left – it does seem that Byrne's weaker revised claim *is* well motivated. In other words: in virtue of being in *S* I have at least *the ability* to acquire (and express) the higher-order thought that I am in S.

All this of course assumes that I *am* linguistically able. One clear situation in which I might be linguistically disabled would be if I was temporally or permanently mute. Clearly we do not think this circumstance would impinge on my conscious life, let alone render me unconscious! There is a difficult issue here regarding whether any ability to externally manifest one's judgement is required. However, prima facie this seems implausible.[87]

---

[87] Of interest here are supposed cases of total locked-in syndrome (TLIS) or *coma vigilante*. Locked-in syndrome (LIS) is a condition caused by local and severe brain stem damage and clinically defined as quadriplegia (complete neck down paralysis) and anarthria (loss of speech-enabling motor-ability) *with preservation of consciousness*. In so-called 'Classic' LIS, patients often retain the ability to shift their gaze and move their eyelids. That is they retain the ability to publicly report on their condition. However, in TLIS patients lack even these limited means of communication, and yet there is good reason to think that they

What does seem plausible is that in situations where I cannot express the state I am in to others, if I am conceptually competent, I remain perfectly capable of forming the higher order belief that I am in the relevant experiential state and of *manifesting* that knowledge in a judgement *in foro interno*. Given this, we should replace Rosenthal's linguistic qualification with a conceptual one.

'The Argument from Reporting and Expressing' is badly named. It is really an argument from *manifesting*. And what the argument supports is not a higher-order thought model but rather Self-Intimation. Undergoing a conscious episode, I am in a certain epistemic position, a position which places me to form a judgement about said state which will manifest knowledge. In other words, undergoing a certain conscious episode puts me in a position to know about said episode.

Two common intuitions (one stemming from the rejection of a certain kind of scepticism in the conscious realm, the other stemming from a connection between consciousness and reportability) lead us to Self-Intimation. However, many will be dissatisfied. Such dissatisfaction arises, I suggest, either because of the alleged existence of decisive counter-examples to Self-Intimation, or because of other commitments that many wish to hold which are supposed incompatible with Self-Intimation. In the following sections I clarify and defend Self-Intimation showing (a) that it is compatible with significant error about our mental lives (§4), (b) that it is not in conflict with relational or naïve realist theories of perceptual experience (§5), and (c) that it can be understood so as to reconcile it with the alleged non-transitivity of indiscriminability (§6). Following on from that discussion, I argue that Self-Intimation grounds a certain principle concerning how our experience can seem to us, a principle I call Seems → Is (§7). Seems → Is rules out systematic, ineliminable error about the positive phenomenal characteristics of our own conscious lives. I conclude by summarizing the picture of experience in time set-out in the first three chapters of the thesis (§8).

---

are perfectly capable of thought and other mental acts and that they are well aware of things being said or done around them. The existence of such cases is evidenced by (a) the occurrence of transient or improving forms of TLIS, (b) the implausibility of a sudden cut off given the apparently very high degree of cognitive and mental capability in 'classic' LIS patients, and lastly (c) brain imaging of patients with such conditions. For a brief clinical overview, see Smith and Delargy 2005. For discussion of the relevant philosophical implications see Kurthen et al. 1991. On the evidence from neuro-imaging see Kotchoubeya et al. 2003 which concedes that such data can only strictly indicate what the authors describe as "the intactness of mechanisms necessary but not sufficient for conscious information processing". Also see Block 2007: 483-4 with which I agree that we should focus on cognitive accessibility not reportability.

4. Self-Intimation and Error

Self-Intimation, as it stands, does not mention the terms 'self-consciousness' or 'self-awareness' (terms I use interchangeably) but it can easily be rephrased in those terms. Self-Intimation is the claim that conscious experience is necessarily available to self-awareness, i.e., to self-conscious reflection. I talk about knowledge in part to distance myself from any commitment to introspective faculties conceived as a form of inner perception taking experience as its object.[88] I assume that all self-awareness is awareness *that*. It is a form of propositional knowledge, a knowledge of truths about our mental lives.

Self-Intimation is quite compatible with our being profoundly ignorant of many aspects of our minds.[89] Self-Intimation is a claim about *perceptual experience* in its *conscious aspect*. This sets aside two obvious concerns often raised concerning self-knowledge. Firstly, Self-Intimation has no immediate bearing on our knowledge of other, non-perceptual aspects of mind. This is clearest in cases where the aspect of mind does not have intrinsic phenomenal aspects, for example the standing attitudes of belief and desire, and plausibly also knowledge.[90] It may be that Self-Intimation also holds of other conscious episodes, sensations such as pain being prime candidates. However, there are clear reasons to resist hasty generalisations in this regard. Perceptual experiences are, very plausibly, individuated by their phenomenology, i.e., by their contribution to what it is like to undergo them.[91] Given this, Self-Intimation commits us to the availability from a subject's point of view of the *kind* or *type* of experience that they are undergoing.

---

[88] My impression is that such a view is criticised far more commonly than defended. See Shoemaker 1994: esp. Lecture I for some scepticism concerning the attribution of an 'Object Perception Model' of introspection to Locke and Armstrong.

[89] Insofar as I endorse the Kantian thought at the head of this chapter, 'representations' must be taken to have a narrower extension than the contemporary use of the term does.

[90] Not that I regard knowledge of our mental states as independent of phenomenology. It seems likely, for example, that an epistemology of our standing attitudes will need to appeal to their conscious manifestations in at least some cases. This is very plausible when it comes to emotional states and moods such as love, fear and sadness. But it may well be true that our epistemology of self-knowledge will need to appeal to the manifestation of standing attitudes, such as the manifestation of belief in conscious *judgement*.

[91] As Soteriou puts it, "We should accept that two conscious experiences are of the same fundamental kind if and only if what it is like for the subject to be having the one experience is the same as what it is like for the subject to be having the other" (2005: 194). Note that such a claim implicitly assumes that certain ways of classifying experiences are the most privileged. To paraphrase Martin (2006: 361), that for all such events there is a most specific answer to the question, 'What is it?'.

Occurrent thoughts also have phenomenal properties but are, very plausibly, *not* individuated by those properties. Thus, the considerations in favour of Self-Intimation will not straightforwardly entail that subjects are always in a position to know what they are thinking.

Self-Intimation is also compatible with profound ignorance of many of the properties of our own perceptual experiences. Mere unaided reflection on my own experience cannot reveal to me that my experience occurs on the same day that Britney Spears was born on, or that it is the last taste I will ever have of an *Old Fashioned*. Moreover, if experiences are identical to physical events, then they have intrinsic physical properties. If so, there will be intrinsic properties of experience which unaided reflection cannot reveal to me, viz., these physical properties of my experience. Self-Intimation is silent as regard properties which do not themselves 'contribute to what, subjectively, it is like for the person who enjoys them'. Its scope is only a subject's experience's, present, phenomenal properties. In talking about properties of experience, it is these properties I am concerned with.

To illustrate this, consider a series of objections that Paul Snowdon raises to a principle closely related to Self-Intimation, a principle he calls the Negative Seems Principle (NSP) (2006: 194).[92]

NSP $(\forall e)((e \text{ is an experience} \wedge \text{not Seems } (F(e))) \rightarrow \text{not } (F(e)))$

Snowdon offers four objections to NSP.

(a) "A proponent of NSP must claim that each episode of consciousness seems to fulfil NSP. I suggest that this is obviously false" (2006: 194).
(b) "My pain can be improved by taking paracetamol. It does not seem so" (195).
(c) "[Arguably] perceptual consciousness depends on how the subject is neurally. But does it seem so?" (195)
(d) "The episode of phantom limb pain does not take place in the limb (because there is no limb) but elsewhere. It does not seem so" (195).

[92] Snowdon finds NSP and PSP (see the discussion of naïve realism below) at work in Honderich's argument for 'Radical Externalism' but principles like them abound in philosophical discussion.

These objections can be rephrased as objections to Self-Intimation for in each case the relevant property of experience is arguably not available to self-consciousness simply in virtue of undergoing the experience. However, Self-Intimation's restriction to conscious, experiential properties disposes of these objections. Consider (a)-(c) first: although an experience of mine may play a certain epistemological role, be alterable by drugs, and be neurally dependent, these facts are not conscious, experiential aspects of my experience.[93]

In this way, Snowdon's objections (a) through (c) to NSP can be avoided by the amended, NSP*.

NSP* (∀e)(∀F)((e is an experience ^ not Seems (F(e))) → not ((F(e) ^ Experiential property(F))))

Handling case (d) needs a little more care. Snowdon's objection relies on the assumption that because the episode of phantom limb pain does not take place in a limb (for there is none) it must take place "elsewhere". A defender of NSP* might deny this. That is, they might deny that the pain takes place anywhere else, having only an apparent location. It would then be true that the experience had the property, 'not taking place in a genuine limb,' but there would be no reason to think that this was an experiential property. On the other hand, it might be true that the pain takes place elsewhere (say in the brain, as is commonly asserted that all pains do). However, here too the defender of NSP* can deny that this location property is experiential.

A significant number of other candidate counter-examples can be found in the literature. Let me mention just one, paradoxical thermal sensations.[94] There are a wide variety of cases falling under this head, ranging from simple ones akin to everyday cases where cold

---

[93] Precisely what status the epistemological property has is a difficult question. One might conceivably argue that conforming to a philosophical thesis *was* a conscious property of my experience. However, the only grounds for this claim seem to me to rely on the fact this knowledge is available to introspective reflection alone. If so, no counter-example is forthcoming.

[94] David Chalmers pressed me to mention this case and two others. One other candidate counter-example involves cases of consciousness outside attention where, one might hold, we are not positioned to know what phenomenal features we are aware of without attending, which, the thought would be, would be changing them. I see no reason to embrace this sceptical line of thought. Of course, *foveating* some location can dramatically alter our experience of that location. However, our ordinary conception of attention is quite compatible with features being attended to just as they were when unattended. A second source of difficulty too large to discuss here is the sceptical idea that the visual world might be a grand illusion. For discussion see the papers in Noë 2002. There is much to be said about the data put forward in support of such claims but, again, I am sceptical of the scepticism as it relates to claims in this chapter.

water is initially felt as hot (see the discussion in Susser et al. 1999 and references therein), to more complex cases, such as the Thermal Grill Illusion where "the simultaneous application of adjacent cutaneous warm and cold stimuli, whose temperatures are well below the heat and cold pain thresholds, is capable of producing a paradoxical painful sensation with a burning quality in a large majority of subjects" (Bouhassira et al. 2005: 164).

This is not the place to delve into the literature on the extremely complex systems that lie behind our perception of temperature (and the closely related perception of pain). Fortunately, we need not do this to address the supposed counter-example since there is no obvious description of the phenomena which cannot be handled by a proponent of principles such as those defended in this chapter. Some paradoxical thermal sensations are simply illusions, cases where a cold stimulus elicits a sensation as of heat. As Defrin et al. (2002: 508) note, the usual hypothesis is that these sensations are due to "the 'inappropriate' activation of thermal channels by cold and warm stimuli". No doubt naïve subjects judge incorrectly concerning the stimulus temperature in such cases but there is no reason to think that we have here a mistaken judgement about the sensation itself.

Perhaps there are cases where subjects do struggle to judge correctly concerning the sensation itself. This would be unsurprising if the sensation at some place were 'paradoxical' in involving both a warm and a cold sensation simultaneously. However, there is no conceptual reason to rule out cases of simultaneous hot and cold sensation at some single location – and indeed, an empirical story which involved two or more thermal channels (and perhaps also the integration of thermal information with nociceptor activation) might even predict them. Whilst such cases may often elicit confused judgements and demand careful attention in order to be correctly classified, they do not flout manifestation requirements. If simultaneous, co-located hot and cold sensations are possible, then a subject who judges that they *must* be having just one or the other is not rationally responding to their experience.

Even granting the qualifications above, a more general and common objection to theses concerning self-knowledge is that we can be grossly mistaken about our perceptual experience *even in its conscious aspect*. Consequently, it is vital to appreciate that Self-

Intimation is perfectly compatible with making serious mistakes about one's own conscious life. *Subjects may fail to exploit their evidential positions.* Just because a subject is undergoing an experience of a certain kind and consequently is, according to Self-Intimation, in a position to know that he is, this does not mean that he will exploit his evidential position. Indeed, he may completely fail to exploit his evidential position for various reasons. A lack of attention or interest may result in one's judgements being poorly constrained by the available evidence.[95] Temporary or permanent psychosis may render judgements wholly unconstrained by grounds.[96] This is quite compatible with these grounds nonetheless being available to one at that time.

One of the major problems with Dennett's discussion of self-awareness is his tendency to run together two quite distinct claims in this regard. In particular, Dennett fails to distinguish between the following claims.

(a) Necessarily, a subject cannot be conscious of a stimulus if they do not believe that they are conscious of that stimulus.
(b) Necessarily, a subject cannot be conscious of a stimulus if their consciousness of that stimulus is beyond their ken (i.e., Self-Intimation).

Dennett tells us that his own brand of "*first-person operationalism* … brusquely denies the possibility in principle of consciousness of a stimulus in the absence of a subject's belief in that consciousness" (1991a: 132). This seems like a commitment to (a). However, in the same breath he supports this claim with the contention that "[p]utative facts about our consciousness that swim out of reach of both 'outside' and 'inside' observers are strange facts indeed" (ibid.: 133). This only justifies a commitment to (b), or equivalently to Self-Intimation. At least with respect to perceptual consciousness, Self-Intimation insists that there cannot be facts about our consciousness that are out of our reach.

---

[95] Cf. Martin's discussion of inattentive John (2004: §9 and 2006: 380f.). Horgan and Kriegel enjoin us to "Imagine that you are engrossed in a heated argument with your conversational partner, and therefore do not notice the refrigerator's hum. You nonetheless have a dim auditory experience of the hum, but are simply not in a position to form true beliefs about your auditory experience" (2007: 131). They consider this case "plausibly effective" (ibid.) against any thesis of self-intimation. I take it our thesis counts as such. Yet it is very far from clear what grounds they have for claiming that an engrossed party really is not in a position *at all* to form beliefs as opposed to just not likely to.

[96] Martin makes this move in responding to A.D. Smith's (2002) objections to disjunctivism. In cases of "hallucinogenic medication, schizophrenic delusion, or simply hypnotic suggestion", Martin insists, "[a] subject … need not be rationally responding to how things sensorily seem to him," i.e., to their actually sensory experience. See his 2006: 389.

However, to be out of reach or beyond our ken contrasts being *within* reach, or *within* our ken; it does not contrast with already being reached, or already known or believed. This though is what Dennett's belief requirement, (a), assumes. It assumes that facts about our conscious lives not already believed would be strange facts. I see no reason to claim that.

Dennett also supports his operationalism by insisting that we should reject any view that "creates the bizarre category of the objectively subjective – the way things actually, objectively seem to you even if they don't seem to seem that way to you!" (ibid.) If one thinks that how things seem should be identified with the actual judgements that one makes, or actual beliefs one has, then endorsement of this intuition will commit one to (a), i.e., the brusque denial of the possibility of consciousness in the absence of belief in that consciousness. But Dennett gives no reason to think that we have to identify how things seem with *actual* judgements. If we identify how things seem with one's current evidence, that is with what one is in a position to know (in virtue of undergoing one's current experience), we can take the above intuition as an expression of the truth of Self-Intimation. Running (a) and (b) together has led to an exaggerated sense of the importance of counter-examples to (a). They have been thought to establish the complete independence of consciousness and self-awareness. Focusing on (b) reveals that they do no such thing.

Focusing on the evidential state of experiencers (what experiencers are in a position to know) may also allow us to avoid the objection that non-human animals and human infants present a clear counter-example to Self-Intimation. Non-human animals and young infants plausibly lack the concept EXPERIENCE, and thus straightforwardly lack abilities necessary to know or judge that they are enjoying experience with some property. Nonetheless it is implausible simply to deny that animals and infants lack conscious experience.[97]

One can, as is commonly done, simply modify Self-Intimation so that it applies only to creatures with conceptual capacities. However, this is not wholly satisfactory. If Self-Intimation captures something fundamental about perceptual consciousness, then it ought to speak to consciousness in all its myriad guises, including the kind of

[97] Though Carruthers seeks to make this option less unappealing, see for example his (2005b).

consciousness enjoyed by infants and animals, as well as the constitutionally inattentive and irrational. Exploiting a robust conception of evidential state or *availability* suggests an alternative response.

In his defence of disjunctivism, Martin argues at length that we need to distinguish between two kinds of discriminability: impersonal and individual indiscriminability.[98] Martin explicates the distinction between the two kinds of indiscriminability as follows.

> In general, where we ascribe an incapacity to someone, we indicate not only that they have failed to do something, but also that there is some ground in virtue of which they so fail. When we think of an individual's incapacity in relation to the specific ground for his or her incapacity – a ground which explains not only why they do not do *F*, but would not even do *F* in similar circumstances – we can still recognise that this impossibility or incapacity is quite consistent with the possibility that others do do *F* or at least could do it. On the other hand, when we talk of an incapacity or inability without indicating a subject lacking the capacity or incapacity, then we need not think in terms of a ground relative to an agent in virtue of which the act cannot be carried out. To say that something is invisible is not to indicate some specific lack in certain viewers, but rather to indicate something about it, that it cannot be seen. In parallel, when we talk of two things being indiscriminable, we need not mean that there is something about a given agent in virtue of which they cannot be told apart [individual indiscriminability], but simply that it is not possible to know that they are distinct [impersonal indiscriminability]. (2004: 75)

In the present context, we need to consider *availability* and not discriminability. But the same point holds: to say that something is visible is not to indicate some specific capacities in some particular group of viewers, but rather to indicate something about it, that it can be seen, that it is available *to vision*.

Similarly, to say that some state or fact (here that a subject is enjoying experience of a certain kind) is *available* to self-consciousness is not to indicate some specific capacities in the subject in virtue of which that fact is available to them to come to know (attentive and introspective mechanisms, say, or possession of relevant conceptual capacities). All we need mean is that the fact can be known about simply in virtue of undergoing the

[98] Here see Martin 2004: 75f. and also Martin 2006: 379f..

experience, i.e., that the fact is available *to self-consciousness*. If we understand the locution 'in a position to know' in this sense, we can declare that bats and babies are in a position to know about their experience despite their lack of conceptual sophistication. They are enjoying experience that is available to self-consciousness. As such, they are positioned to know about that experience. This claim is quite consistent with the idea that samenesses and differences are beyond the ken of the particular creature or child, for here we do indicate a specific deficit to explain the unavailability.[99]

Relatedly, note that focusing on what we are in a position to know also has a strategic value. Discussions of self-consciousness and consciousness typically begin by focusing on the self-consciousness side of the relation and then engaging in what we might call the epistemology of introspective judgement, seeking to understand the privileged status that self-conscious or introspective judgements apparently enjoy. However, any attempt to proceed in this way with introspective epistemology must, at some point, specify which judgements are to count as introspective. After all, a view which claims a special status for some class of judgements will lack content unless it provides some way of gaining a grasp on that class.

The literature contains two broad approaches here. On the one hand, one might seek to provide a positive description of what it is to be an introspective judgement. On the other hand, one might try and define introspective judgements in a negative way, as judgements *not* formed on the basis of observation or inference. Either way we face severe difficulties. In the latter case, it is difficult to know how to exclude judgements formed on no basis at all from those granted a special status. Yet clearly ungrounded judgements are not epistemically privileged in any way. It is simply not true that every judgement that has our inner lives as its subject matter is epistemically privileged.

In the former case, the obvious approach is to treat introspective judgements as judgements that result from the operation of some kind of mechanism or faculty. However, if introspection is some kind of mechanism by which we track our inner lives, then (assuming that that claim has genuine content) it must make sense for the mechanism in question to misfire, and so for our judgements to go wrong. Yet, as we

[99] Those sceptical should toe the standard line, treating Self-Intimation and related claims as restricted to the conceptually sophisticated.

saw above (§2), intuitions driving Self-Intimation suggest that such failure is impossible – for failure would introduce a kind of sceptical scenario concerning our inner lives, which we are unwilling to allow.[100] Consequently, picking out introspective judgements as those that result from the operation of some faculty or mechanism (however weakly conceived) is something a friend of Self-Intimation does well to avoid. On pain of begging the question against themselves, they should focus on what consciousness puts us in a position to know.

5. Self-Intimation and Naïve Realism

A commonly held view is that Self-Intimation and its kind are a vestige of a disreputable Cartesian internalism concerning the mind. In particular, it is very often suggested that a defender of Self-Intimation will have to reject relationalist, disjunctivist or naïve realist approaches to perception. Depending on one's prior philosophical commitments this may or may not seem like a serious bullet to bite. Regardless, I think it is a mistake to think that Self-Intimation is inconsistent with disjunctivism. In this section, I explain why.

The argument from hallucination can be thought of as proceeding in roughly the following manner. Consider a subject *S*, and a normal, veridical perceptual experience that they enjoy of a phenomenal kind *K*. For any such experience, we must concede the possibility of a situation in which *S* is enjoying a hallucinatory experience that they cannot discriminate from experience of kind *K*. Furthermore, if mind-independent objects literally constitute aspects of the phenomenal character of experience (as naïve realists contend), and assuming that experiences are individuated by their phenomenal characters, then such hallucinations must differ in phenomenal character and so kind from veridical experiences.

It certainly can appear to follow from this that there are differences in phenomenal character that are inaccessible to their subjects. After all, don't hallucinations seem just the same as veridical perceptions and yet differ in phenomenal character? That in turn

[100] As Martin puts it, "Introspective judgement cannot result from the correct operation of a specific mechanism of introspection without the possibility of one's phenomenal consciousness merely seeming some way to one" (2006: 393). For extensive discussion see his 2006: §7.

appears inconsistent with Self-Intimation. In what follows, I show why this is mistaken. That said, it must be acknowledged that, in one specific regard, the disjunctivist will have to allow that experience can seem some way and yet not be that way. I discuss this point shortly (§7).

Let us then assume a naïve realist account of perceptual experience according to which perception is conceived of as a relation to mind-independent objects, which (in part) literally constitute the phenomenal character of such experience. Consider in particular a perceptual experience as of a yellow mango as such. According to the disjunctivist such cases divide into two distinct kinds. For simplicity, call them the good case and the bad case. (Of course there are intermediate cases where the mango is misperceived, say, as green when it is in fact yellow. I set these cases aside for present purposes.) We can now consider the application of Self-Intimation in each case.

5.1 The Good Case

In the good case, I perceive the yellow mango in front of me as such. One of the properties of my experience according to the disjunctivist is 'being an encounter with a yellow mango as such'. In accordance with the above proposal, in virtue of undergoing the experience, I am in a position to know that my experience is an encounter with a yellow mango. As a consequence, I am in a position to know that there is a mango in front of me purely in virtue of undergoing the experience.

One might challenge this as follows. What about hallucinations? Do you not concede that there is a distinct experiential state (viz., the state I would be in were my neural states to remain precisely as they currently are and yet the mango be removed from view) in which, if I were in that state, I would not be in a position to tell it apart from a genuine perception of a yellow mango despite making the very same judgement made above? (Reply: *Yes*.) Does this not undermine the claim that I can know that I am in the good case when I am in it?

Reply: *No.* For it to undermine the claim to know in the good case, one would have to assume that, just because the bad case is indiscriminable from the good case, that the good case is indiscriminable from the bad.[101] The disjunctivist can and should deny this. Instead, they should hold that, if you are lucky and in the good case, then you are in a position to know that you are – and hence in a position to know that you are not in the bad case. In the good case, you can discriminate your experience (activate knowledge) that you are not in the bad case. This is far from *ad hoc*. If one thinks, (a) that one's experience literally involves objects in the good case but not in the bad, and (b) that there is an intimate connection between one's experience and one's epistemic position, then it will hardly be surprising that one's epistemic position is not neutral across the two cases.

Someone might press this: What if the possibility of hallucination is highly salient in your context? Perhaps – as Martin imagines (2001: 444) – you are inside a highly reliable hallucination generator, designed to cause you to hallucinate a yellow mango. A one-in-a-million electrical fault occurs and the machine does not work so, by chance, one happens to see (quite veridically) the yellow mango that has fortuitously been placed in front of you. Given the salience of the hallucination case, many would argue that you are not in a position to know that your experience is a genuine presentation of a yellow mango.

We can envisage two ways of telling this story. In the first I am fully aware that I am in a hallucination machine and have no idea that it has misfired. In this case we need to distinguish what one is in a position to know from what it is reasonable to judge. In the scenario just described, I claim that one *is* in a position to know that one's experience is a presentation of a yellow mango. However, it would be quite *unreasonable* for one to judge this, knowing what one knows. One has the evidence but it is not reasonable to exploit it. As Martin puts it,

> Someone who is alive to a genuine possibility of illusion or hallucination will allow doubts about whether they really are perceiving to constrain the judgements he or she makes. Even if a subject is in the privileged situation, the presence of doubts can put him or her in a position where he or she cannot responsibly exploit the privilege that

[101] Here compare Pryor's objection that 'phenomenally indiscriminable' is a symmetric relation noted by Martin (2006: 364, fn.16) and also Ruben's assumption of just this symmetry in his discussion of disjunctivism (2008: 228).

> they have. And it is in this way that collateral information can defeat the justificatory role that otherwise perception would play. (2001: 447)

In the second telling of the story, one has no idea that one is in such a machine. One has been tricked into entering it and believes that all is entirely normal. Does the nearness of a world where one is in fact hallucinating deprive one of knowledge? It is not at all obvious that it does. If one thinks that there is an intrinsic epistemic connection between experience and evidential position, then one *is* in a certain epistemic position if one is having the experience. Facts about nearby possibilities are irrelevant. I remain in a position to know; I am simply stupendously lucky.[102]

5.2 The Bad Case

In the bad case, I am undergoing a hallucination as of a yellow mango as such. My experience lacks the property of being a presentation of such a fruit though it seems to have it. I am not in a position to know that my experience has such a property because I cannot know what is false. But nor am I in a position to know that my experience is not a presentation of the world. Things seem the same but my evidence is not the same. My experience misleads me not only about the world but about itself. Does the bad case present a counter-example to Self-Intimation? Recall that Self-Intimation was the claim that for any *conscious property* of experience, one is in a position to know that one's experience has that property purely in virtue of undergoing it. Which conscious property am I in the dark regarding here? The disjunctivist interested in defending Self-Intimation must retort: *None*. Certainly, I overshoot in my judgements insofar as I claim that my experience is a presentation of a yellow mango. However, I am wholly aware of all the conscious aspects of the hallucination.

How can that be? How can my hallucination differ in phenomenal character from a veridical perception and yet lack any conscious properties which distinguish it from a veridical perception? Put this way the answer is obvious. There is no logical difficulty in claiming that every conscious property of a hallucination is had by corresponding

[102] Of course this is far from the only story that we might want to tell. Another popular option would be to think of one being in a position to know *ceteris paribus*, i.e., in the absence of defeat. Taking such a route would weaken Self-Intimation but would far from deprive it of force.

veridical perceptions, whilst denying the converse, i.e., insisting that veridical perceptions have aspects of conscious character which distinguish them from hallucinations. This is a feature of Martin's approach to hallucinations. (See Martin 2004, 2006.) According to Martin, hallucinations are exhaustively and fundamentally characterizable in epistemic terms, as not knowably not a certain kind of veridical perception. Any corresponding veridical perception, not being discriminable from itself, will have this property (a property manifest to subjects), but will also have more determinate aspects, like actually being a presentation of a yellow mango.

Hallucination is a form of experience in which subjects are deceived about the nature of their own experience. They will rationally judge that their experience has properties which it does not have. A cost of disjunctivism is being forced to allow that there are cases where one's experience seems to one some way and yet is not that way (§7). However, that consequence does not involve a rejection of Self-Intimation.

6. Self-Intimation and Non-Transitivity

The supposed non-transitivity of indiscriminability poses perhaps the most serious challenge to Self-Intimation for the following reason.[103] Let us say that A-kind experience is indiscriminable (in the sense relevant for present purposes) from B-kind experience just if it one is not positioned to know that A-kind experience is not B-kind experience simply in virtue of undergoing it. Let us also assume what Self-Intimation maintains, namely, that we are in a position to know of all phenomenal aspects of our experience simply in virtue of undergoing said experience. We can now informally derive the following contradiction with assumptions which represent the standard gloss on the non-transitivity of indiscriminability.

(1) Assumption: A-kind experience is indiscriminable from B-kind experience.
(2) There can be no phenomenal aspect of our A-experience which is incompatible with being B-experience otherwise (given Self-Intimation) we would be in a

[103] For a flavour of the large literature on this topic see Dummett 1975, Wright 1975, Fara 2001, Mills 2002, Hellie 2005 and discussion in Chapter Seven. For the style of objection here considered see Williamson 1996 and 2000.

position to know that our A-experience was not of the B-experience kind (contra 1).

(3) Assumption: B-kind experience is indiscriminable from C-kind experience.

(4) There can be no phenomenal aspect of our B-experience which is incompatible with being C-experience otherwise (given Self-Intimation) we would be in a position to know that our B-experience was not of the C-experience kind (contra 3).

(5) Assumption: A-kind experience is discriminable from C-kind experience.

(6) So there *is* some phenomenal aspect of A-experience which is incompatible with it being C-experience placing us to activate knowledge that our A-experience is not C-experience.

However:

(7) There can be no phenomenal aspect of our A-experience which is incompatible with being C-experience otherwise (from Self-Intimation) we would be in a position to know that our experience was not of the C-experience kind and (given 4) not of the B-experience kind (contra 3).

Contradiction.

Why think there are three kinds of experience which conform to assumptions (1), (3) and (5)? That is, why think that we can enjoy three kinds of experience: A-experience, B-experience and C-experience, such that A-experience is indiscriminable from B-experience, and C-experience is indiscriminable from B-experience, but discriminable from A-experience? The reason is that this seems entailed by the existence of situations, $\Sigma_{1\text{-}3}$, such that a subject fails to distinguish the scene in Σ_1 from that presented in Σ_2, likewise fails to distinguish the scene in Σ_2 from that presented in Σ_3, but nonetheless succeeds in distinguishing the scene in Σ_1 from that presented in Σ_3. Consider, for example, three scenes each involving a different pair of straight-lines from the set $\{L_1, L_2, L_3\}$. A subject may fail to distinguish L_1 from L_2, and likewise L_2 from L_3, but nonetheless discriminate L_1 from L_3.

Now one option here is to deny that such cases genuinely exist. Fara (2001) can be understood as pursuing this strategy. Chapter Seven discusses aspects of Fara's view and raises some doubts about the plausibility of that approach, committed as it seems to be to a gap between what we experience and what we are in a position to notice (i.e., to

unnoticeable appearances). In any case, it seems to me that a different tack is needed in general if we are to secure Self-Intimation as a general thesis against the above line of argument. The idea can be outlined as follows.[104]

That a subject on occasion fails to distinguish the scene in situation Σ_1 from that presented in Σ_2, likewise fails to distinguish the scene in Σ_2 from that presented in Σ_3, but nonetheless succeeds in distinguishing the scene in Σ_1 from that presented in Σ_3, does not immediately show that we have three kinds of experience in play here which conform to the assumptions in the above argument. Rather it may be that a given situation can give rise to more than one kind of (non-illusory) experience. If that were so, we might imagine that Σ_2 was a situation in which subjects were liable to enjoy two distinct kinds of experience as follows.

$\Sigma_1 \rightarrow E_1$
$\Sigma_2 \rightarrow E_1$ *or* E_2
$\Sigma_3 \rightarrow E_2$

In this light, consider the following margin for error principle invoked by Williamson:

> [C]ases in which one is in a position to know *p* must not be too close to cases in which *p* is false, otherwise one's belief in *p* in the former cases would lack a sufficiently reliable basis to constitute knowledge. (2000: 17)

Now, take situation Σ_2. Here one's situation is such that one might very well have enjoyed (non-illusory) experience of one of two distinct kinds. Imagine one enjoys, as it happens, experience of kind E_1. Consequently, one is not in a position to discriminate one's situation from Σ_1; after all, one is enjoying just the kind of experience that one would be having if one were in Σ_1. But, because of the margin for error required for knowledge, the fact that in the very same situation (Σ_2) one might as easily have enjoyed experience of kind E_2 means that one is too close to E_2 experience to know that one is enjoying experience of kind E_1 as opposed to E_2. As a result one is not in a position to

---

[104] The discussion here is heavily indebted to Martin 2004: 76-9. He there points out that non-transitivity worries are generated by speckled-hen type cases as well as the cases of supposed phenomenal continua Fara discusses. That might seem a reason to prefer the account in the text. In fact, it seems that Fara would likely make parallel moves in the speckled-hen cases.

discriminate one's situation from Σ_3 either, since in Σ_3 one is enjoying precisely the experience that one is not in a position to discriminate from one's current experience.

Nonetheless, for all that, experience of kind E_1, is *of a kind* that is distinguishable from E_2-type experience; put another way, it is impersonally indiscriminable from E_2-kind experience. This is because E_1-type experience *can* be had in a situation which does not afford the risk of having E_2-type experience, viz., Σ_1. E_1-experience in Σ_1 is safe. Although it is true that E_1-type experience can be enjoyed in a situation where it would not be discriminable *by its then subject* from E_2-type experience, this does not mean that it is not impersonally discriminable from E_2-type experience.

How does Self-Intimation fare in this light? Recall,

> **Self-Intimation** If a subject is undergoing perceptual experience of a certain experiential kind, then that subject is in a position to know that they are undergoing perceptual experience of that experiential kind simply in virtue of so undergoing.

Non-transitivity shows why, in certain cases, this is not quite correct. Enjoying a perceptual experience in certain situations (such as Σ_2) above need not place a subject in a position to know whether they are enjoying E_1- or E_2-type experience. Nonetheless, according to the above account, a subject enjoying experience of that kind *is* undergoing a kind of experience such that knowledge that one is undergoing experience with its phenomenal properties *is to be had* solely in virtue of undergoing said kind of experience. The point is that knowledge can sometimes be blocked in a given situation by the closeness of other experiential kinds that the situation in question affords.

Now, of course, given the nature of cases in question, it may well be that almost any similar situation (for example, any case where a subject is presented with a line of some length, or a colour swatch of some shade) in fact affords multiple distinct kinds of experience. In other words, the vast majority of situations here in view may be of the form: $\Sigma_i \rightarrow E_i$ *or* E_j … *or* E_n But this is not obviously a problem. What it shows is that it

may be that *no* situation affords knowledge of *all* properties of experience. But this does not entail the existence of any property which *cannot* be known about.[105]

Thus, strictly speaking, Self-Intimation should be the denial that there are (determinate) types of experience (i.e., phenomenal kinds) such that knowledge that experience is of that (determinate) type is simply unavailable solely in virtue of undergoing said experience. That does not rule out cases where margin for error principles prevent the occasion of having the experience from being one where the subject is in a position to know. It does rule out cases where some (determinate) aspect of experience is never such that a subject is in a position to know about it simply through undergoing the experience. In other words, if a perceptual experience has some (determinate) experiential property, then a subject of that kind of experience *can* be positioned to know that it has said property solely in virtue of undergoing it, even if they happen not to be so positioned on a given occasion.

This emendation, though crucially important in this context, has no significant impact on the argumentation that follows. As a result, I retain the original simpler formulation of Self-Intimation, simply noting the interaction at the relevant points.

7. Self-Intimation and Seems → Is

Paul Snowdon uses hallucination as an objection to a principle he calls the Positive Seems Principle (PSP).[106]

PSP $(\forall e)((e \text{ is an experience } \wedge \text{ Seems } (F(e))) \rightarrow F(e))$

In now familiar terms his objection runs as follows.

> If when I see a large pink elephant the consciousness seems to me to consist of a pink elephant then when I am having a large rat hallucination it will equally seem to me to consist of a large pink rat. Clearly, though the hallucinatory episode cannot consist of

[105] Here compare Goodman's (1977) criteria for the identity of qualia.
[106] Snowdon 2006: 192.

> that, there being no pink rat to do the constituting. I conclude that we cannot rely on PSP (2006: 194)

Snowdon's objection must be granted. However, I do not think that this means we should simply abandon all hope of establishing a claim in the vicinity of PSP. In this section, I suggest that, hallucination notwithstanding, a commitment to Self-Intimation does in fact involve a commitment to a close cousin of PSP.

Self-Intimation tells us that the experiential properties of certain experiences are within their subject's ken. Consider again the example of experiencing a yellow mango as such. Self-Intimation tell us that the subject of such experience is in a position to know that they are undergoing experience of something yellow. Given this, could it *seem* to them that they were instead having an experience of something green? That depends on how one understands *seeming*. Someone paying little attention and perhaps harbouring certain expectations might *judge* that they were seeing something green. Clearly such inattentive judgements are not always correct. Nor are irrational judgements. As Martin puts it, "[a] subject who is deluded into supposing that he now experiences angels talking to him need not be rationally responding to how things sensorily seem to him" (2006: 389). Thus, if we identify how things seem with what subjects do in fact judge, then experience can certainly seem other than it is in these ways.

However, it is hard to make sense of a fully rational and attentive subject judging that they are enjoying experience of something green as such, despite being in a position to know that they are enjoying experience of something yellow as such. In order to make sense of this kind of situation, we need to find a barrier or explanation of the failure to exploit the available knowledge.[107] In this light, the following principle might seem acceptable.

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

---

[107] In addition to inattention and delusion, believing that such a judgement would be unreasonable in the circumstances is such an explanation. See above.

Here the term 'systematically' rules out experience which is, in its very nature, such as to seem a way that it is not. As such it handles complications just discussed relating to non-transitivity. Unfortunately, the disjunctivist will be forced to allow that hallucinations *are* systematically misleading in one fundamental respect; they appear to involve mind-independent constituents (yellow mangoes, pink rats, and the like) but they do not.

Nonetheless, there is a simple amendment to PSP* that the disjunctivist can make which renders it acceptable. That amendment depends upon accepting a negative, epistemic characterization of hallucination.[108] This is the claim that hallucinations are most fundamentally characterized as episodes that are not knowably not veridical perceptions of a certain kind. In other words, the *only* positive mental characteristics of hallucinations are negative epistemological ones as opposed to positive 'hallucinatory' character properties. As Martin puts it, hallucinations of this class are characterizable "in nothing but epistemological terms, in terms of their unknowable difference from cases of veridical perception" (2004: 38).

Accepting this characterization, the following substantive principle does still hold.

> **Seems → Is** If some course of experience *has* a phenomenal property of a certain kind, then that course of experience cannot *systematically* seem to *differ* phenomenally in that regard.

This principle is supported by the discussion of the yellow mango above. If your experience does present or apparently present you with a *yellow* mango as such, then that experience cannot systematically seem to present you with a *green* mango as such. Things cannot be like this: a phenomenal aspect of your current experience is its being as of a yellow mango as such, yet it systematically seems to be an experience as of a green mango as such.

The new principle also avoids the objection from hallucination which refutes PSP*. If hallucinations are characterized epistemically in terms of their indiscriminability from

---

[108] As already noted, such a characterization is defended at length by Martin (2004, 2006). Strictly speaking Martin only makes claims concerning causally matching hallucinations, i.e., hallucinations brought about by replication of the proximate causes of veridical perceptual experiences. Martin discusses the first steps of extending the account to other cases in his 2004: 80-1. See also Siegel 2004.

ordinary veridical perceptions, then hallucinations do not have phenomenal properties with regard to which they seem to differ. They seem to be indiscriminable from veridical perceptions and they are. It is true that they seem to be presentations of mind-independent objects though they are not. But they do not have presentational phenomenological properties which differ from those they seem to have, only properties of being epistemically indiscriminable from experience with presentational phenomenology.

Of course, the complexities here only arise if one is committed to naïve realism or disjunctivism. For all that has been said here, a common factor theorist can rest content with the simpler formulation, PSP*. In any case, the weaker principle, Seems → Is, has all the strength needed to serve as premise in what follows and I only appeal to that claim henceforth.

8. A Naïve View of Temporal Awareness

A natural, if one likes *naïve,* picture of perceptual experience is the following.[109] Perceptual experience is the taking in of world by mind. In such experience we are directly acquainted with the structure of the world; the world presents itself to us in experience. In the light of the above principles, a natural working out of this thought in the temporal case is the following. The temporal structure of experience is quite simply directly inherited from the temporal structure of the world; mind to the external world is fitted. On this view, the ordering and relative duration of experiential episodes exactly mirrors the ordering and relative duration of the events experienced. If one event lasts twice as long as another and you directly experience the whole of both events, then your experience of the first will last twice as long as your experience of the second. If a series of six tones is played to you in a certain order, your experience of them will be ordered in just the same way. This picture is too naïve of course – temporal illusion abounds.[110] However, the spirit of the picture can survive the lesson of illusion if it claims instead

[109] Naïve in the sense of being a philosophical articulation of what "we all pre-theoretically accept concerning the nature of our sense experience" (Martin 2006: 404).

[110] For temporal order illusions and debate about their interpretation see: Warren et al. 1969, Bregman and Campbell 1971, Sussman et al. 1999, Näätänen et al 2001, Macken et al. 2003, and Micheyl et al. 2005. For temporal duration illusions and similar debate see: Brown 1995, Kanai et al 2006 and brief discussion in Chapter Five, below.

that the temporal structure of experience matches the *apparent* temporal structure of the world. Then, at least in cases of veridical perception, mind to the external world is indeed fitted.

In the last three chapters I argued for a series of claims. They can be combined so as to articulate the naïve view of temporal awareness just sketched. Here are the claims.

> **Realism** Experience itself in its experiential aspect consists of events and/or processes of which it is true that they persist through time and occur before and after one another.
>
> **Temporal Transparency** For any particular temporal determination of experience, it is never the case that when I consider how things are with me experientially, I am rationally positioned, through reflection on my experience alone, to judge my experience to be so determined other than on the basis of a judgement concerning (and typically perceptual attention to) the temporal structure of the apparent objects of perception.
>
> **Self-Intimation** If a subject is undergoing perceptual experience of a certain experiential kind, then that subject is in a position to know that they are undergoing perceptual experience of that experiential kind simply in virtue of so undergoing.
>
> **Seems → Is** If some course of experience *has* a phenomenal property of a certain kind, then that course of experience cannot *systematically* seem to *differ* phenomenally in that regard.

Realism tells us that experience has its own *experiential* temporal structure. Self-Intimation tells us that we are always in a position to know experience's own temporal structure. Temporal Transparency tells us how; our grip on that structure can only come via taking the temporal structure of experience to match the temporal structure of the apparent objects of experience. Putting these claims together we can conclude that we are always in a position to take the temporal structure of experience to match the temporal structure of the apparent objects of experience. Moreover, Seems → Is insists that (since, after all,

we are always in a position to know) we cannot be systematically in error when we make judgements about the temporal character of our experience itself. Thus, the temporal structure of experience cannot systematically come apart from the temporal structure of the apparent objects of experience.

We have arrived at a picture strikingly close to the naïve view sketched above. In what follows, I refer to the set of claims above in those terms, as our natural or naïve view of temporal awareness. Natural as it is, the view of experience in time just articulated has come under attack from many quarters. The rest of this thesis considers the most important of these challenges and strives to answer them.

PART II – EXPERIENCE OF TIME

Chapter Four:

Anhomoeomery and Experience

Time should be imaged with a paint-brush instead of a scythe; he knows how to wield the former even better than the latter.[111]

1. Preview

The rest of this thesis is devoted to the treatment of a number of puzzles that arise in relation to our experience of time. Although a great deal of work has been generated by these puzzles, they are rarely connected, nor are they framed against a clear view of how we naturally conceive of experience *in* time. I have two central aims in discussing these puzzles. First, to vouchsafe the naïve picture of temporal experience developed in Part One. For a common reaction to at least two of the puzzles is to reject some element of that picture. Second, to argue that at the root of three of the seemingly distinct puzzles is the same mistaken way of thinking about experience. The mistake is to think that experience is *homoeomerous* down to very short durations or instants. The core thought here is the idea that we can analyse experience down to a series of independent short-chunks or slices, and explain the nature and features of the stream of consciousness in terms of those temporal units. In contrast, I argue that experience is significantly *anhomoeomerous*. 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

[111] Sherwood (1897: 134) *An Epistle to Posterity*.

not vice-versa. That is, the stream is structured such that over short periods the explanatory direction runs from temporal whole to temporal parts.

If we reject the assumption of homoeomery, the puzzles dissolve. Indeed, I argue that a commitment to the picture of temporal experience developed above *commits* us to rejecting homoeomery. Seeing this allows us to appreciate the difficulties philosophers have encountered when trying to account for temporal awareness whilst wedded to the naïve picture of experience in time and also to a commitment to homoeomery. Before introducing the puzzles, I introduce the general concepts of homoeomery and anhomoeomery since they serve as the backdrop to subsequent discussion. I begin with the spatial case.

2. Spatial Anhomoeomery

We naturally distinguish between countable *things* and amassable *stuffs*. Cars are a kind of thing and one prima facie mark of being a thing like a car is lacking proper parts which are themselves cars; cars are not made of cars. In contrast, gold is a kind of stuff and one prima facie mark of being a stuff like gold is that *any* volume contained within a lump of gold also counts as a lump of gold. Gold is thus homogeneous or, better, homoeomerous, literally "like-parted"; it is gold all the way down.

Of course, this is all too picturesque. For one, we can imagine constructing a car out of cars such that the former car literally has car parts – cars as parts. For another, once we drop to volumes of gold measuring only a few cubic picometres, we will no longer find ourselves with lumps of gold for we are at the level of atomic nuclei, electron clouds and (mostly) 'empty' space.[112] Thus, even if van Inwagen is right to think that "[i]t was possible for an ancient Greek to believe that matter was continuous and that all stuffs were 'homoeomerous' " (1993: 27), it seems that chemists and physicists have now demonstrated that none in fact are.

[112] See, for example, van Inwagen on bronze, who concludes, "Every stuff that is visible to the naked eye arises from an arrangement of atoms that are not themselves made of that stuff." (1993: 27) Also Taylor 1985: 69.

One does not need the aid of microscopy to find examples of stuffs where spatial homoeomery breaks down. Barry Taylor's nice example is fruitcake, division of which may or may not leave us with fruitcake as remainder. After all we may simply find ourselves with a sultana and "a mere sultana does not in itself constitute a lump of fruitcake" (Taylor 1985: 70). Many stuffs are like fruitcake in this respect – salad, moussaka, rope, chain-mail, granular sedimentary rock (e.g., coquina), and cloth (e.g., hessian).[113] Whereas gold is homoeomerous down to the molecular scale; these stuffs seemingly cease to be homoeomerous at macroscopic, indeed, perfectly visible scales.

Despite the realisation that many stuffs are not homoeomerous down to arbitrary scales, we are inclined to think that a bowl can be *filled* with fruitcake, a table *covered* with cloth, and so on. Someone might object to this, 'But, look, here is a sultana in the middle of the fruitcake; sultanas are not fruitcake; so the bowl isn't filled with fruitcake because this volume of the bowl isn't!' In response, we need to recognize that although Taylor is very plausibly right to say that a sultana, in itself, is not a lump of fruitcake, a sultana-filled volume of the bowl can count as a fruitcake-filled volume *in virtue of what fills a larger surrounding volume*. Whether a small volume is filled with fruitcake is thus essentially dependent on what is true of the volume's surroundings. We explain the nature of the small volume in terms of facts about the larger, surrounding volume. Put another way, with anhomoeomerous stuffs, like fruitcake, the direction of explanation at a certain scale is whole-to-part not part-to-whole. This has important consequences.

Imagine asking whether any fruitcake occupies a finite volume picked out by a set of Cartesian co-ordinates. This seems to be a perfectly good question. And, at first blush, it seems that to answer it we only need to consider what fills the volume specified. But what has just been said shows that we need to be very cautious. We need to distinguish between two questions that we might be asking. Firstly, we might be asking whether there is any of the stuff, *fruitcake*, filling the relevant volume. What the above discussion suggests is that this question cannot be answered solely on the basis of what is contained within the volume, considered in independence from what fills surrounding volumes. If the volume contains half a sultana, then it may or may not be a volume which is fruitcake filled. That depends on whether the half-sultana is part of a fruitcake or not. In other

[113] Note that Aristotle's examples are natural substances like wood and bone, flesh and metal. There is some plausibility to the idea that such stuffs *appear* homoeomerous. See Aristotle 1922: 314a.

words, the volume only counts as a fruitcake filled volume in virtue of what fills a larger volume encompassing it. Its own filling is insufficient to establish an answer.

On the other hand, we might be asking whether the volume *by itself* counts as a lump of fruitcake. What the above discussion shows is that this question is quite independent of whether the volume is filled with fruitcake. Even if it is, it may be that only half a sultana fills the volume in question, and so the volume in and of itself will not count as a lump of fruitcake. Taylor seems to me right that even if the half-sultana is a part of a fruitcake, half a sultana is not in and of itself a lump of fruitcake.[114] One might deny this and insist that the half sultana *does* count as a lump of fruitcake in virtue of its being embedded as it is in the fruitcake. Although I find this way of speaking counter-intuitive, I do not want to take issue with it. What is important is the thought shared by both views, that the nature of the volume in question – whether it is *filled* with fruitcake and/or constitutes a *lump* of fruitcake – is essentially dependent on the nature of the volume's surrounds. Either way, we explain the nature of the small volume in terms of facts about the larger volume.

For someone who thinks that a volume can count as a lump of something in virtue of being embedded in a lump, there is a sense in which stuffs turn out to be 'like-parted' after all. Any volume contained within a lump of the stuff will count as a lump of the stuff. [115] However, there remains good reason to think of such stuff as anhomoeomerous because small volumes considered in independence from their embedding will not count as lumps of the stuff. In what follows, I shall speak about anhomoeomerous stuffs in this broad sense. Anhomoeomerous stuffs have parts the nature of which constitutively depends on the nature of the wholes in which they are located. Gold was not traditionally conceived to be like this. Fruitcake visibly is.

If, on the other hand we think, with Taylor, that a volume may only be stuff-filled and not necessarily a lump of stuff in virtue of being embedded within a lump of stuff, we face the question of how to differentiate stuffs from things. We cannot of course do this

---

[114] Just as Inwagen seems to be right to say that the atoms that make up visible stuffs are not in general themselves made of that stuff.

[115] Or at least any genuinely constitutive volume. No-one will want to claim that *anything* embedded in a fruitcake might count as a lump of fruitcake. A ring that slips off one's finger and into the mixture whilst making the cake needn't so count.

by saying that any volume of a lump of stuff is also a lump of that stuff. Taylor provides an alternative way of fleshing out the intuitive idea that stuffs *fill*, where substances *delimit* the spaces that they occupy as follows. If S is a stuff then S fills the spaces it occupies either homoeomerously or anhomoeomerously. If S is a homoeomerous stuff, then the above condition holds – any volume of a lump of the stuff is also a lump of that stuff. If S is an anhomoeomerous stuff, on the other hand, there will be parts of the stuff which do not count as lumps (chunks, hunks etc.) of the stuff (as sultanas do not count as lumps of fruitcake). Nonetheless, *any* volume *containing* a lump of stuff will so count. Add a sultana to a lump of fruitcake and you still have fruitcake. Add another gold atom to a lump of gold and you still have a lump of gold. In contrast, a thing like a car may have a car as a part. However, it is not true that any volume containing that car part will count as a car. Thus imagine this super-car of cars has a car as a roof. The car-as-roof may be a car but the car-as-roof plus windscreen and wipers is not.[116]

As I say, I am inclined to follow Taylor here and deny that sultanas, however embedded, count as lumps of fruitcake. However, what is important in what follows is a thesis about explanation which holds independently of this claim. Thus, to repeat, for our purposes, a stuff is said to be anhomoeomerous just in case a lump of that stuff has a sub-part whose nature (be it as a lump of stuff, or simply as a volume filled with stuff) depends constitutively on facts about the stuff beyond the volume of the particular sub-part in question. I also talk about stuffs being homoeomerous *down to* points as well as to tiny areas. Thus, a stuff is homoeomerous down to areas of a cubic centimetre in size just if, for any area A of at least a cubic centimetre which falls wholly within a lump of the stuff in question, A is also a lump of that stuff in its own right and independent of any facts about surrounding areas. Exploiting this flexibility, we need not worry if all stuffs turn out to be anhomoeomerous – some may still be homoeomerous down to any finite volume, for example.

Issues concerning anhomoeomery do not arise simply when we consider individual very small volumes. Consider, for example, a fairly large volume with precise boundaries. We might naïvely think that questions such as, 'Does this volume contain any fruitcake?' or 'How much fruitcake does this volume contain?' can be answered by establishing what is

---

[116] This offers a rough and ready summary of the picture detailed by Taylor 1985: 72. I am inclined to drop his condition (ii)(a) given the possibility of an infinite universe filled with luminiferous ether.

within the volume, in independence from what is outside of it. However, imagine that the original volume overlaps one protruding half-sultana which is part of a fruitcake. In that case, the original volume overlaps part of a fruitcake-filled volume, and so does contain fruitcake. Yet the above method would not have reached that conclusion. Rather, ignoring the fact that the half-sultana was part of a fruitcake, it would have concluded that there was merely a half-sultana in the volume.

3. Temporal Anhomoeomery

Mourelatos (1978) proposes that the thing/stuff distinction mirrors the event/process distinction.

> Events … occupy relatively to other situations [Mourelatos' generic name for the referents of verb predications] a position analogous to the one objects or things or substances occupy relatively to stuffs and properties or qualities. (1978: 430)

Thus, mirroring the above discussion, consider the prima facie plausible claim that events are like cars in lacking proper parts which are of the same type. Thus, "the capsizing of a boat is not made up of boat-capsizings" (Mourelatos 1978: 430); the second world war is not made up of second world wars; and a performance of Shostakovich's Second Piano Concerto is not made up of performances of Shostakovich's Second Piano Concerto. In contrast, processes are prima facie analogous to stuffs like gold in being temporally homoeomerous. For example, prima facie, it seems that if the glider is gliding through the sky from noon until one o'clock, then any interval between noon and one will be an interval during which the glider has glided. Again, prima facie, if a spinning top is spinning on the table for twenty seconds, then any period of time within that twenty seconds will be a period of time during which the top has spun.

Once again this is all too picturesque. For one, we can imagine bringing about certain events which have events of the same kind as sub-parts. For example, a large town might put on a large fireworks display which has as parts sequential smaller displays organized by individual neighbourhoods: a firework display with fireworks displays as parts. More importantly, there are also examples of processes where homoeomery breaks down. Given the importance of this breakdown in what follows, I consider several examples.

3.1 Walking and running

Consider first the two processes (which are also activities), walking and running, and the following passage taken from a textbook on human motion.

> Human gait involves alternating sequences in which the body is supported first by one limb, which contacts the ground, and then by the other limb. Human gait has two modes, walking and running. One distinction between these two modes lies in the percentage of each cycle during which the body is supported by foot contact with the ground. When we walk, there is always at least one foot on the ground; and for a brief period of each cycle, both feet are on the ground. Accordingly, walking can be characterized as an alternating sequence of single and double support. In contrast, running involves alternating sequences of support and nonsupport, with the proportion of the cycle spent in support varying with speed. For both walking and running, however, each limb experiences a sequence of support and nonsupport during a single cycle. The period of support is referred to as the stance phase, and nonsupport is known as the swing phase … one complete cycle … is defined as a stride. (Enoka 2002: 179)

What this passage reveals is that, regardless of whether someone is running or walking, a certain sub-period of that process will include a single support stance phase. During this period not enough is going on for what is going on, just on its own and independently of what is going on during surrounding periods, to count as running or walking. If enough was going on, we would face an impossible trilemma: Is enough going on to count as walking, running or both? Yet as the above passage suggests not enough of a cycle has gone on to determine the answer to this question.[117]

This suggests that the process of running, whilst prima facie homoeomerous, is in fact only homoeomerous down to small intervals.[118] The same goes for walking, and plausibly also skiing, marching, dancing, gambolling, strutting, ambling, striding, staggering and

---

[117] Note that this is not an epistemological worry but a metaphysical one. Not enough has gone on to ground running or walking.

[118] *Contra* Vendler "running and its kind go on in time in a homogeneous [i.e., homoeomerous] way; any part of the process is of the same nature as the whole" (1957: 146). Cf. Rothstein: "a cumulative predicate such as *run*, although intuitively homogeneous [i.e., homoeomerous], has non-homogenous minimal parts: there are parts of running events which are just too small to count as events of running." (2004: 11)

moseying. Indeed, the more activities we list, the more powerful the above consideration becomes. For consider a very brief period of human body movement. It is highly implausible to hold that for all activities such a period might be a sub-interval of, that enough always goes on during such a period to determine which it is a sub-interval of and which not. This would be like saying that although sultanas are used in baking dozens of different cakes, there is always enough in any individual sultana to determine which kind of cake it is and isn't a part of! This is clearly nonsense.

Such activities (and so processes) are not homoeomerous down to very small intervals. Arguably, in many contexts, they are not homoeomerous down to quite significant durations. On certain occasions, one can truly say, 'I have been walking all day,' even if from time to time one has paused to catch one's breath or take in the view. Indeed, on some occasions, one can truly claim to have been walking all day despite having sat-down for a pint of *Jennings* and a ploughman's lunch at *The High Cross Inn*. Consider (a) the appropriateness of the claim, 'We walked all day, taking lunch at the inn en route,' and (b) the inappropriateness in certain contexts of the reply, 'No you didn't [walk all day], I saw you sat in the pub garden at lunchtime,' and, finally, (c) the jokey ineptitude of the following exchange at the bar: 'What are you up to this fine day?' 'We're walking the Haystacks.' 'Doesn't much look like it!'.

Of course, there are high standards cases, i.e., contexts to be found in which it would be false to make the claims above. One can imagine an army training exercise which demanded that those involved walk all day around a mountain, where that instruction meant, *without stopping at all even to catch your breath*. But this does not mean that ordinary fell-walking should be held to be a discontinuous process as opposed to a continuous but anhomoeomerous process.[119]

Just as we ordinarily think that a fruitcake can fill a bowl, we consider that someone, say Paula Radcliffe, can be running or walking *throughout* a certain period of time. Someone

[119] It is an interesting question how we should think of these different cases. If I am watching the army exercise from a vantage point all day I might truly say, 'I saw a dozen men walking all day around the mountain'. I suggest that this would be true even if the officer in charge of the exercise caught some or all of the men resting briefly and so quite rightly criticised them for not walking all day as they were ordered to do. This suggests a contextualist as opposed to a subject/event-sensitive account. In other words, in some cases, whether a given process counts as continuous or discontinuous will depend on the concerns of the context that you as speaker/thinker are in.

might object, 'But take that brief interval during a period you claim that Paula is running throughout; Paula only accomplishes a single support phase during that interval; a single support phase isn't running (or walking for that matter); so Paula isn't running throughout the period.' As in the spatial case, we should respond by noting that a period *and its sub-periods* may be *filled* with Paula's running (just as a tin can be filled with fruitcake) in virtue of the fact that Paula ran over the period. We explain the nature of the sub-interval (its being an interval during which Paula was running) by reference to a larger whole (over which Paula ran). Thus, again, explanation is essentially whole-to-part and not part-to-whole. Such explanations cannot refer solely to the past either. Imagine Paula has just stepped out the door. With her single-footfall she has neither run nor walked. Nonetheless, she may well have started running or walking. If she has, that can only be in virtue of what is yet to come.

3.2 Talking

Mourelatos gives the examples, "thundering, giggling, or talking … as paradigms of anhomoeomerous process" (1978: 430).[120] Why might Mourelatos think that talking (and also: speaking, chatting, discussing, shouting, rambling on etc.) is such a process? One plausible reason is that to count as having talked one has at least to utter a phoneme/distinctive feature (i.e., a genuine phonological unit), otherwise one would simply be making noise. Analogous arguments to the running/walking case here apply; not enough has been achieved in a very brief period considered in independence to ground talking as opposed to mere noise-making, or indeed singing, stuttering, grunting and so forth. Intentions are of no help here since the person intending to talk might well be struck dumb (or struck by lightning) before succeeding in having talked.

Another, equally plausible reason is that two people can truly be said to be talking for hours on end despite there being silences between words, between remarks, and indeed lulls in the conversation where nothing is said by either participant in the exchange for

[120] Cf. Taylor: "… an example like 'chuckles' provides a case more naturally conceived on the analogy of a heterogeneous [i.e., anhomoeomerous] stuff, since any sounds emitted in a microsecond during a period of chuckling (at the normal rate) hardly constitute chuckling themselves, but rather appear to stand to chuckling as a sultana might stand to fruit-cake, viz. as at best falling within some period of chuckling though themselves occupying a time too short to constitute such a period" (1977: 212). Taylor also mentions: 'giggles', 'talks', 'walks', and 'strokes (the dog)' (ibid.).

some time. There seems no reason to think that from the perspective of many, perhaps all contexts, some silences during the maximal period of talking (i.e., some periods which on their own are not periods where anyone has talked) are compatible with a continuous stretch of talking.[121] Such periods will count as periods during which talking is going on in virtue of being located within durations over which the participants talk.

3.3 Other Examples

Other examples of anhomoeomerous processes include:

Watching (and also: listening, tasting, feeling etc.) – for one can truly claim to have been watching television all evening despite popping to the bathroom or running to the door to pick up the pizza delivery, let alone blinking.

Swimming (and also: playing, exercising, sparing etc.) – for one can truly claim to have been swimming for thirty minutes straight even if one has put one's feet down at the shallow end every other length.

Growing (and also: rising, falling, shrinking, dying, swelling etc.) – for one can truly claim to be growing all year despite doing so only at night, or despite not gaining any height or even shrinking for a brief period, just as the FTSE can have been gaining all year despite the occasional drop as judged minute by minute.

Likewise, consider: nudging, downloading, interrogating, writing, painting, dressing, and snoring. In each case, it is quite implausible to think that all sub-intervals of periods during which these processes are occurring will count on their own as periods during which the relevant process is going on, that is, on their own make it true that someone has nudged etc. over that period. In most of the above cases, this holds even from the perspective of the strictest of contexts. In sum: many processes are analogous to fruitcake and not to gold.[122]

---

[121] To say, 'We were talking for hours,' is entirely felicitous despite clear silences etc.

[122] Of course, there may be reasons to think that just as gold is not homoeomerous, no process is actually homoeomerous for reasons to do with fundamental physics. For example, it may be that Planck time (5.39×10^{-44} seconds) is the in principle minimum time period for which any process can go on for. My point in

4. Consequences

These everyday facts about walking, talking and so forth have important consequences. Imagine asking whether walking is going on during some interval. We might be asking two distinct questions. Firstly, we might be asking whether the interval is an interval throughout which walking is going on. What the above discussion shows is that this question cannot be answered just on the basis of what is going on during the interval in question considered in independence from what is going on in surrounding periods. If, during the interval, only a single support stance phase is unfolding, then the interval may or may not count as one during which walking is going on when considered independently – whether it does so depends on what is going on in the surrounding interval. In independence, everything that is going on in the interval is compatible with running and not walking; whatever is going on during the interval considered in independence is insufficient to establish an answer to the walking/not walking question.

On the other hand, we might be asking whether the interval *by itself* counts as a period over which the individual in question has walked. What the above discussion shows is that this question is independent of whether the interval is one in which walking is going on. Even if walking is going on throughout the interval, it may be that only a single support stance phase fills the interval in question, and so the interval *by itself* will not count as one during which the individual has walked. A single support stance phase is not enough for walking in this sense – though walking can be said to be going on during an interval in virtue of its being part of a period during which the walker has walked.

Again the issue is not just about individual small periods. Imagine a long stretch of time, at the end of which someone completes a single support phase. Has that person run at all during the long stretch? We cannot say without taking into account what occurs in the following period. Perhaps the person *is* running during the period in virtue of that last support phase being a part of a period during which they do run. Or again, perhaps they have just begun walking.

---

the text has been that many processes are not homoeomerous even at quite gross timescales, timescales which will be relevant to our experiential lives.

The metaphysical moral is this: if we try and ask very precise questions about the lengths of time certain processes go on for, their precise starting points, or their characteristics, we must consider long enough periods of time. Failure to do so may blind us to the presence or nature of a process unfolding during the period of our primary concern. For instance, if we probe whether a given process, say a computer's humming, is going on over a given time period, we must consider how things stand with respect to a period, long enough to support the process in question (which is not to say long enough for the process to be going on within it, but long enough *on its own* and in independence of surrounding periods to count as a period during which the computer has hummed). When the process is spinning or falling, the relevant length of time may simply be some finite duration. However, when the process is walking or talking, the length of time may be considerably longer. And indeed, in the case of some processes (in certain contexts), the relevant timescales may be surprisingly long.

An important consequence of this is that we should be extremely wary of assuming that, just because x has not Φ-ed over some period δt, a process of x Φ-ing is not going on *throughout* some period Δt where $\delta t \subseteq \Delta t$. Conversely, we cannot assume that, just because a process is going on throughout some period Δt, x will have Φ-ed over any period δt such that $\delta t \subseteq \Delta t$. In the case of some processes, it may be that questions such as, 'Is Φ going on at some instant *t* or over some short duration δt?' are problematic in just the same way as questions like, 'Is there fruitcake at some point *p* or in some tiny volume (δx, δy, δz)?' These questions may only be answerable if we know whether a process is going on throughout Δt where $\delta t \subseteq \Delta t$, just as the analogous question concerning fruitcake requires us to know whether there is fruitcake filling a volume (Δx, Δy, Δz) where $(\delta x, \delta y, \delta z) \subseteq (\Delta x, \Delta y, \Delta z)$.

To say that a process is anhomoeomerous then is to say that it has temporal sub-parts whose nature depends constitutively on facts about the nature of the process before and/or after the particular temporal part in question. In other words, the process is structured such that over certain short periods the explanatory direction runs from temporal whole to temporal parts and not vice-versa. A key contention of subsequent chapters is that experience is significantly anhomoeomerous, just like walking and talking.

This is not obvious and will require extensive argument. However, the above considerations demonstrate two things. First, that if experience is a process, it would not be especially unusual if it were an anhomoeomerous process. Second, that if experience is anhomoeomerous, then it would be seriously mistaken to think that we could always appreciate whether a given experiential process was occurring at very short timescales (or with limitless precision) in independence of a subject's experience over longer timescales. It is possible that at such brief timescales considered in independence of surrounding periods we will simply not appreciate how things are experientially, since how things are experientially depends on what is unfolding over longer, encompassing periods.

Just as there was a question above as to whether a sultana counted as a lump of fruitcake or merely as a fruitcake filled volume in virtue of its embedding, there is a question about how best to think of parts of anhomoeomerous processes. If Paula has only accomplished a single support phase she may certainly be running throughout that interval in virtue of what follows; but has she *run* in virtue of what follows? I am strongly inclined to think not.[123] But again the crucial point is one about explanatory direction and so for our purposes we can treat any process as anhomoeomerous just if the process has temporal parts whose nature (be it as periods filled with S's Φ-ing, or periods over which S has Φ-ed) depends constitutively on facts about the process beyond the period of the temporal part in question.

Similarly, just as I talked about stuffs being homoeomerous down to points and small areas, so I talk about processes being homoeomerous down to instants and small intervals or durations. Thus, a process, P, of x Φ-ing is homoeomerous down to periods of a second in duration just if, for any period Q of at least a second which falls wholly within the course of P in question, Q is period during which x has Φ-ed in its own right and independent of any facts about surrounding areas. Exploiting this flexibility, we need not worry if all processes turn out to be anhomoeomerous – some may still be homoeomerous down to any finite duration, for example. Above I proposed gliding and spinning as plausible examples of temporally homoeomerous processes at least down to

[123] In this I take my lead from Rothstein and Soteriou who deny that running is homogeneous, defining homogeneity in the following ways. Rothstein: "If a predicate is homogeneous then *x P-ed for y time* ENTAILS that at any time during y, *x P-ed* was true" (2004: 14). Soteriou: "[w]here Φ-ing is an activity, if it is true that S was Φ-ing from $t_{1\ \text{to}\ t10}$ (e.g. walking), then at every point during that interval it is true that S Φ-ed" (2007: 552).

any finite duration. Another is falling; all finite sub-durations of a period during which an object is falling are plausibly durations during which the object has fallen.[124] My principle interest in this thesis is to argue that experience is *significantly* temporally anhomoeomerous, i.e., not homoeomerous below temporal intervals of *significant* duration (i.e., of the order of several hundreds of milliseconds).

Finally, just as we can distinguish stuffs from things without claiming that all stuffs are homoeomerous; we can distinguish events and processes even if we deny that, if S Φs (e.g., runs) throughout a period, then it is true that S has Φ-ed (i.e., run) during every sub-period.[125] A process fills time either homoeomerously or anhomoeomerously. If the former, the condition just given holds. If the latter, then any sub-period of the process which contains a period during which S has Φ-ed will itself be one during which S has Φ-ed. If Paula has run during a one second period, then a one and a half second period comprised of that period together with a further half second during which she is running will be one over which she has run. Contrast events. An event such as a war may be made of up separate smaller and sequential events of the same kind. For example, the Hundred Years' War comprised a number of other wars, some distinguished by locality such as the Breton War of Succession and the Castilian Civil War; others by their dates, such as the Caroline War, 1369-1389, and the Lancastrian War, 1415-1429. The Hundred Years' War is nonetheless a war (event) and not simply *warring* (process) because it is not true that any sub-period of the Hundred Years' War which contains a war will itself count as a war. Warring during the period 1369-1417 is not itself a war even though the Caroline war occurred during a sub-period of that period. More generally, if war W is made up of war, w_1, followed by w_2, followed by w_3, then w_1 together with the opening salvos of w_2, need not be a war.

---

[124] Cf. Taylor: "there are cases where the analogy with homogeneous stuffs is appropriate: even a microsecond within a period of falling is plausibly reckoned as itself genuinely a period of falling, even though it can be told as such by means of normal empirical criteria only indirectly, via the knowledge that it does indeed come within some wider period long enough for those criteria to be applied." (1977: 212). As already mentioned, this claims only holds with respect to certain contexts.

[125] For his detailed treatment of temporal homoeomery see Taylor 1977: 213 and 1985: Ch.3.

5. Experience as process

The discussion above assumes that experience is a process. It might be objected that philosophers of mind often treat experience as a state or event. However, philosophers of mind rarely display sufficient sensitivity to ontological categories.[126] The linguistic evidence clearly demonstrates that experience is best thought of as a process. For example, the verb 'to experience' takes the progressive form. It is rarely used in such a form since there are usually more specific things to say. But consider Mary released from her black and white laboratory.

> I am experiencing colour for the first time!

And similarly the following examples (garnered from a quick Google search).

> I am experiencing a nuisance odour/problems with noise in my neighbourhood.
> I am experiencing palpitations/sharp pains in my lower back.
> I am experiencing auditory, tactile and visual hallucinations.
> I am experiencing after-images on certain objects/a pleasant sensation.

Vendler claims that only process predications have progressive forms.[127] Certainly, they contrast with state predications such as the following.

> * I am knowing what it is like to see colour for the first time!

---

[126] Three representative examples: "Visual experience is a state characterised by its typical causal role" (Lewis 1980: 274); "the most philosophically convenient way of describing experience is an event one" (Sprigge 2002: 235); "*belief and perception form a single class of mental phenomena marked by their representational nature and any mental difference between any two members of the class must be a difference in what they represent*" (Thau 2002: 14-5). Confusion here is compounded by the availability of both the mass noun 'experience' (process) and the count noun 'an experience' (event). More on that distinction below. But even when the concern is specifically with experiences, they are often misclassified as states as in, "A visual experience is a state whose phenomenological properties enable me to identify it without external observation" (Nagel 1986: 46).

[127] "The question, 'What are you doing?' might be answered by 'I am running (or writing, working, and so on),' but not by 'I am knowing (or loving, recognizing, and so on).' On the other hand, the appropriate question and answer, 'Do you know . . .?' 'Yes, I do,' have no counterparts like 'Do you run?' 'Yes, I do.' This difference suggests that running, writing, and the like are processes going on in time, i.e., roughly, that they consist of successive phases following one another in time" (1957: 144).

As mentioned, Mourelatos improves upon Vendler's typology. He also improves on Vendler's test for being a process predication by introducing the idea of a nominalization transcription (the idea of which will be clear from the example below). According to Mourelatos, all and only process predications have mass quantified nominalization transcriptions.[128] Thus, a nominalization transcription of, "I am experiencing colour for the first time!" is, "There is, for the first time, (some) experiencing of colour by me!" This is mass quantified; hence experience is a process. Contrast the event predication, "I saw colour for the first time!" which transcribes as, "There was a (one) seeing/sighting of colour for the first time by me!" This is count quantified and so an event predication.

This linguistic evidence supports the intuitive classification made vivid in James' metaphor of the stream of consciousness. Our experience flows or unfolds in time (flowing and unfolding being paradigm processes); it consists of successive phases which occur before and after each other.

Thinking of experience as a process is not in competition here with our count quantified talk of *experiences* which are best thought of as events. Particular experiences result when we conceptually divide up experience up in certain ways. Sometimes the division will fall-out naturally where the object of experience is an event with its own temporal shape as, for example, when you recall your experience of a lightening strike. Sometimes the division will be arbitrary as, for example, when you contemplate your experience from 10:52 a.m. until 11:24 a.m. In either case, we have what is distinctive of events, namely, a final temporal part – the last moment of experiencing before 11:24 a.m., or the last moment of experiencing the lightening. Experience *qua* process consists of successive phases which occur before and after each other, but does not require a final temporal part, it does not climax as say, "Hillary and Tenzing conquered Everest" does with the moment that they reach the summit.

There are reasons for taking experience *qua* process as primary. In particular thinking of the stream of consciousness as divisible into experiences but denying that it is built up from separate experiences allows us to dissolve certain worries about the diachronic unity

[128] See Mourelatos 1978 and Steward 1997.

of consciousness.[129] However, the considerations below do not turn on this point for events are clearly anhomoeomerous. If I wrote this thesis during 2008, it certainly does not follow that I wrote this thesis during March 1st 2008. I may well have been *writing* it then, of course. Thus, the central contention of this chapter would stand regardless of whether we treated experience as a process or as an event. However, I suspect that the fact that processes have traditionally been assumed to be homoeomerous in part lies behind the assumption that experience is homoeomerous. Given that, the most sympathetic way to understand the debate is to treat experience as a process. In this way, when we see why the assumption of experiential homoeomery should be rejected, it will be for deep as opposed to shallow reasons. I now turn to those considerations.

[129] A thought which goes back at least to Bergson 1914: 3. Much the same idea plays a central role in Tye 2003. I see no reason to make the further denial – as Tye does – that there are such things as experiences (events) in the sense noted above.

Chapter Five:

Phenomenal Anti-Realism

Tell us, too, the way
Time, in its fullness, fill us
As it flows: tell us the beauty of succession
That Breton denied: the day goes
Down, but there is time before it goes[130]

1. 'Time Seemed to Slow Down': a Warm-Up Case

In the discussion of cinematic representation above (Ch.3, §3.4) I discussed a view which denied that there was any more to the depiction of temporal properties than the representation of temporal features in the narrative. I suggested that such a view might provide a starting point for thinking about anti-realism in relation to the temporal structure of experience. In this section, I want to consider an objection to the naïve picture developed in Part One. Although I do not think the objection is, in itself, one that seriously troubles that picture, answering it opens up the prospect of a more extreme, anti-realist view. Approaching anti-realism in this way helps prepare us for this chapter's main concern, the serious challenge to the naïve view posed by Dennett's arguments.

The basic objection to the framework developed above is that it is ill-equipped to handle the phenomenon of 'time slowing down'. Consider the following narrative, which relates the experience of a 6.7 magnitude earthquake in Northridge, California.

> I had just gotten up to go to the bathroom when the earthquake came. It was unbelievable. It threw me off my feet. My roommate and I ran to get underneath the dining room table. As I was running I saw my computer flying through the air. Once we

[130] Tomlinson (1997) 'In the Fullness of Time' from *Selected Poems*.

> got under the table we held on for dear life! It felt like it lasted forever but really it was only a few seconds. Once it stopped we went outside and saw the mess. Pipes had broken and everything was flooding, there were little sand volcanoes in the front yard. There was lots of damage to peoples houses and chimneys but everyone on our block was uninjured. It was the most frightening experience of my life.[131]

Statements such as, 'It felt like it lasted forever but really it was only a few seconds,' are very common expressions of people's experience in extreme circumstances, such as earthquakes, falls and car crashes. Of course, the claim here is an exaggeration. But this need not detain us since our problem is generated by any claim that some event, like an earthquake, felt like it lasted a long time, when in reality it only lasted a few seconds. It is also generated by claims to the effect that time seemed to slow down which are found in similar contexts. How should we account for such judgements?

Given our framework, one thing that it is *not* open to us to say is that such statements express judgements about how long some event seemed to last *in contrast* to the temporal extent of the experience of the event, as if our experience shifted in the way illustrated below.

ORDINARY CASE EARTHQUAKE CASE

Object Time

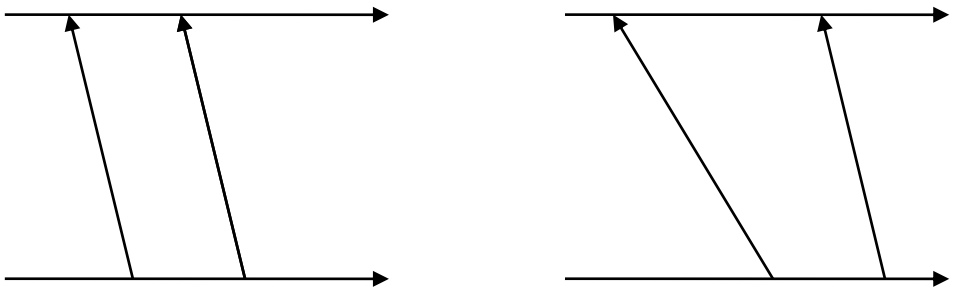

Act Time

[131] This is one of dozens of narratives collected as part of a project funded by the U.S. Geological Survey Earthquake Hazards Program. See http://earthquake.usgs.gov/learning/teachers/Mag_vs_Int_Pkg.pdf.

Divertingly, a recent paper by Stetson et al. (2007) entitled, 'Does Time Really Slow Down during a Frightening Event?' takes this kind of model seriously and attempts to put it to the test. Their idea is that during a stressful or frightening event there should be, according to the above picture, an increase in temporal resolution. To test this, subjects were dropped fifty metres from a platform into a net, freefalling for about two and a half seconds. A chronometer was strapped to their forearms, set to show random numbers alternating just too fast to be detected in normal circumstances. Subjects were asked to attempt to read the digits as they fell. As the authors surmise,

> If higher temporal resolution were experienced during the freefall, the alternation rate should appear slowed, thus allowing for the accurate reporting of numbers that would otherwise be unreadable. (2007: 2)

Unsurprisingly, the authors found no evidence to support this hypothesis.[132]

James, in discussing related phenomena suggests an important idea, namely, that "tracts of time … shorten *in passing* whenever we are so fully occupied with their content as not to note the actual time itself" (1890: 626). Evidently, this is precisely the wrong explanation of the earthquake case. Here an occupation with content leads time to *'lengthen*' in passing. However, there are two ideas here that can be considered in abstraction from the particular hypothesis in play. First is James' basic thought that our occupation with the *content* of experience is at the root of our talk of time slowing down or speeding up. Second is James' suggestion that this engagement of attention means that we do not notice the actual time-span itself. Employing the first of these ideas, we can develop the following account of what is going on in the cases in question by way of an objection to the naïve picture articulated above.

The objector might put things this way. Our judgements of duration and rate of passage are not based on the temporal content of experience (as Temporal Transparency would

---

[132] Not that their experiment should be thought of as decisive. Someone might claim that stress slows down the apparent alteration rate but also lowers resolution at the same time. Stetson et al. go on to suggest that duration judgements are due to misremembering. What follows suggests that this is at best half-right. Certainly, duration judgements will be grounded in propositional memory. However, if metrical properties are not amongst the perceptible features of events and so not amongst the experiential properties of experience, we should not think of this as misremembering what we once perceived correctly, but rather as part of the autobiographical narrative that we construct around our episodic memories.

have it) but rather on non-temporal aspects of the content of experience. The content of our experience captures our attention in various ways but, roughly speaking, the currency of attention is *interest*. Thus, a moving, varied or unusual stimulus can be expected to engage attention to a greater extent than a static, monotonous or mundane stimulus. Judgements of absolute duration are based on the engagement of attention. This, in essence, is the explanation of the phenomena above. This shows, the objector continues, that the naïve theorist is wrong to think that temporal judgements are sourced in the temporality of experience itself; rather they are grounded in facts about how attention is engaged in undergoing experience, facts which determine what gets written into our retrospective, autobiographical representation of our stream of consciousness. Consequently, we have no reason to believe that experience itself has any temporal structure independent of our representation of it in autobiographical memory.

How should the naïve theorist react to this? The idea that it is the *interest* of what is taking place in the perceived world (and so the degree to which our attention is occupied, as opposed to temporal features presented in experience) which grounds our metrical durational judgements seems very plausible in the earthquake case. The description above made vivid how much there was to capture the individual's attention over the course of a few seconds – the scramble to get under the desk, the computer flying through the air, the noise and so forth. Consequently, the hypothesis that such dramatic and engaging stretches of experience are judged to last longer than more pedestrian ones, is compelling.[133] The naïve theorist should agree; what she should disagree with is the way that this claim is developed into an objection.

The judgements in earthquake cases concern *metrical* temporal properties. In Chapter One (§1.1) I argued that metrical temporal properties do not form part of the content of perceptual experience. Given that, it should be common ground that judgements concerning metrical durations will be not grounded in temporal aspects of experience. There are no such aspects. Events are not presented as taking, say, six seconds (or proceeding at a certain metrical rate). Given that a judgement that a stretch of experience lasted six seconds (or just 'a long time') cannot have its source in the (metrical) temporal

[133] The idea that interest drives these judgements is plausible but extremely monotonous experience can also seem to drag on for an age. We might account for this by noting that in such cases attention is being deployed in watching out for change of some kind. Thus, in both cases, it is plausible that the extent of attentional involvement connects with the judgements of duration.

content of experience, it must have some other source.[134] The hypothesis above provides an attractive account of that source, and so our thought and talk about time slowing down and events seeming to take 'forever'. However, the hypothesis does not provide any basis for denying that *non-metrical* temporal properties are genuinely experiential properties that ground our judgements about the temporal *structure* of experience. Since that is all the naïve picture claims, the objection fails. It offers no reason for rejecting the claim that the stream of consciousness is temporally structured in a way that contributes to what it is like, subjectively, to undergo it.

For all that, the dialectic above does reveal the possibility of a more radical objector. The more radical objector urges that, once it is granted that metrical temporal judgements are not sourced in the metrical temporal content of experience, we should take seriously the thought that *no* temporal judgements (including those concerning structural properties) are grounded in temporal content of experience. The objector proposes, for example, that we judge that one event was experienced after another, not because of the temporal structure of experience, but rather because taking experience to have been so structured is the best way that we have of making sense of the non-temporal content of our experience. Less metaphorically, the objector might propose that autobiographical memory constructs a narrative that represents the stream of consciousness as having a certain structure based on its non-temporal content.[135] The radical conclusion is that there is nothing more to the experiential temporal structure of experience than the

---

[134] Often people cite dreams as the paradigm case of transparency failure. Treanor (ms.), for example, writes, "dreams … seem to go on for ages though evidence suggests they were really quite brief. Here I am thinking of a dream in which, for instance, some loud banging precipitates an (apparently) extended course of events. Upon wakening, you find a friend is at the door and knocked just a few seconds ago." Cf. Walton 2008: 180. Dreams are not my focus. Nonetheless, points in the text are relevant here. If we assume that dreams are a vivid kind of imaginative episode and that we imagine events in dreams through imagining perceiving them, two points can be made. Firstly, given what I say above, imagining seeing will not be imagining an episode with experiential metrical structure. Thus, metrical judgements can enter in two ways. As above, they can enter through the *content* of the imagined experience. Alternatively, they can enter through associated *propositional* imagining. That is, we may simultaneously imagine some event and imagine that it is taking a certain amount of time. Since dream-contents are unconstrained by physics, dreams may well be about events that would normally take a long time. It would therefore be no surprise if both these mechanisms led to the kinds of judgements Treanor suggests. For some extremely interesting and perhaps surprising empirical findings concerning judgements about dream duration, see Noreika (forthcoming).

[135] Here we have a way of thinking about Kantian synthesis where the narrative structure is compelled to conform to the principles of the Analogies. Here we also have a way of thinking about Dennett and Kinsbourne's notion of content-sensitive settling (see their 1992a), where a far less constrictive set of heuristics are in play.

structure represented in autobiographical memory's narrative, sourced in non-temporal aspects of experience.

The objector has not provided any argument in favour of his extreme view. Distinguishing between metrical and non-metrical properties allows us to endorse the claim that autobiographical narrative driven by attentional occupation is the source of our *metrical* judgements, whilst resisting the view that *all* temporal aspects of experience should receive this account. Nonetheless, in the light of the account of duration judgements tentatively endorsed above, the radical objector's extreme anti-realism no longer appears the alien prospect that it might initially have seemed. As yet it is unmotivated. Nonetheless, where we do find motivation, such a view must be taken seriously as a prima facie viable competitor to the naïve view of Part One.

1.1 Empirical Connections

Clearly many of the claims made above are subject to empirical investigation and to try to put flesh on the skeleton explanation of duration judgements renders us hostage to empirical fortune. In that light, let me quote the following passage from a recent paper, which attempts to summarise the current state of play on the issue.

> Because time cannot be directly measured at a given moment, it has been proposed that the brain estimates time based on the number of changes in an event. Consistent with this idea, dynamic visual stimuli are known to lengthen perceived time. However, the kind of information that constitutes the basis for time perception remains unresolved. (Kanai et al. 2006: 1421; for references see original)

This passage deserves a number of comments. (1) In the first sentence, it is claimed that "time cannot be directly measured at a given moment". It is not clear what this means.[136] However, if I am right that metrical properties are not part of the content of our perceptual experience, then we can rephrase the point as the claim that metrical durations of events are not perceptible features. (2) Talk of the brain estimating time is perhaps better put in terms of the neural mechanisms under-pinning judgements of apparent

[136] Though compare discussion of PSA in Chapter Six.

duration. (3) The idea that moving stimuli "lengthen perceived time" should be understood as the claim that (when presented on their own) moving stimuli are *judged* to last longer than stationary ones; it should not be taken to be a claim about phenomenology. In general, the empirical literature uses the phrase 'time perception' to refer to judgements concerning stimulus durations rather than phenomenological durations. (4) The authors note that the presumption in the field is that the amount of change (in some sense) lies behind our duration judgements. This is very much in line with the discussion above. They also note that the precise basis is as yet unclear.

Kanai et al. go on to provide evidence "that the temporal frequency of a stimulus serves as the 'clock' for perceived duration. Other aspects of changes, such as speed or coherence, were found to be inconsequential" (1421). This suggests a rather simple basis for duration judgements in the simple paradigms considered. I take it that it is an open question whether such a simple mechanism can account for all the phenomena. It seems more likely that a variety of mechanisms come into play in more complex situations. Stetson et al. "speculate that the involvement of the amygdala in emotional memory may lead to dilated duration judgements" (2007: 3). Here the fact that "highly salient events may be erroneously interpreted to have spanned a greater period of time" (ibid.) is plausibly a consequence of emotional salience as opposed to frequency per se. And we might speculate that judgements are affected by how stressed or aroused the subject is.[137] All this remains entirely in line with the general idea in play above.

It is well-attested that, at time scales of 200-1000ms, moving stimuli are judged to persist longer than stationary stimuli even when presentation times are identical. In general, the effect increases with stimulus speed.[138] In the literature, such effects are often termed *illusions*. It is not clear that this is the most appropriate terminology. If it correlates with the rather misleading use of the phrase 'time perception' to refer to duration judgements (and if I am right about the absence of metrical properties from perceptual content), then these are, strictly speaking, *delusions* not illusions. That said, I am by no means ruling out cases of genuine temporal illusion. Perhaps one event may look to last longer than another despite their occurring during the very same interval of time; perhaps the 'rat' of

[137] It is also noteworthy in this context that (at least anecdotally) schizophrenics tend to over-estimate the lengths of time periods and depressives tend to under-estimate them.
[138] For a survey, see Brown 1995.

a postman's 'rat-tat' may sound as if it lasts longer than the 'tat' even though both last the same amount of time. These cases do not trouble Temporal Transparency for they provide no reason to deny that experience itself is structured in the same manner as its apparent objects.

2. Dennett on Time and Consciousness

The rest of this chapter is devoted to Daniel Dennett's writings on time and consciousness, wherein we find the most forceful arguments for anti-realism that have been propounded in recent times.[139] However, its aim is not primarily exegesis; rather, I focus on two related questions. Firstly, to paraphrase Wright (2002), what could anti-realism about phenomenology (and in particular *temporal* phenomenology) possibly be? Secondly, why would anyone feel driven to endorse such a conception of the stream of consciousness?

Dennett's arguments in *Consciousness Explained* and, in particular, those based on 'Time and the Observer' (co-authored with Marcel Kinsbourne) concerning masking and apparent motion phenomena, created a storm in the early nineties and have spawned a vast secondary literature. The storm evinces the widely held view that Dennett presents a powerful challenge to our intuitive ways of thinking about the stream of consciousness. There is little consensus within that literature concerning precisely how one should understand Dennett. Nonetheless, most commentators agree that he is arguing for some form of anti-realism about consciousness, where that term is a catch-all for fictionalist, eliminativist and irrealist views, as well as more traditional anti-realist positions. Thus, Searle (1995) claims, "Dennett denies the existence of consciousness"; Seager avers that for Dennett "there is *no such thing* as phenomenal consciousness!" (1999: 85); Anthony holds that "the most charitable reading of [Dennett and Kinsbourne's] critique commits them to the view that consciousness *does not exist* – to *eliminativism*" (1992: 201); McGinn insists that Dennett "has to regard conscious states as fictional" (1995: 246).[140]

---

[139] For convenience I often talk about Dennett's view, noting only here that Dennett's views on time and experience were largely developed in collaboration with Marcel Kinsbourne.

[140] See also Aronson et al. who conclude: "[D & K] seem to have thrown the baby out with the bath water: They seem to have rejected the essential phenomenology of consciousness in an attempt to avoid dualism and implausible neurophysiology" (1992: 202). Block (1992), Van Gulick (1992: 228) and Lloyd (1992: 216) also find anti-realism explicit in Dennett; see also Tye 1993 discussed below.

These characterisations certainly apply to Dennett's 1979 view, which he described in the following terms.

> I am left defending the view that such [phenomenal] judgements *exhaust* our immediate consciousness, that our individual stream of consciousness consists of nothing but such propositional episodes … My view, put bluntly, is that there is no phenomenological manifold in any such relation to our reports. There are the public reports we issue, and then there are the episodes of our propositional awareness, our judgements, and then there is – so far as introspection is concerned – darkness. (1979: 95)

However, it is not clear whether *Consciousness Explained* represents a change of view or rhetorical strategy.[141] Certainly, post-1991, Dennett has resisted such straightforward characterisations.[142] Thus, Dennett and Kinsbourne respond to their early critics as follows.

> [W]e consider our position to be unproblematically "realist" and materialist: conscious experiences are real events occurring in the real time and space of the brain, and hence they are clockable and locatable within the appropriate limits of precision for real phenomena of their type. (…) Certain sorts of questions one might think it appropriate to ask about them, however, have no answers, because these questions presuppose inappropriate – unmotivatable – temporal and spatial boundaries that are more fine-grained than the phenomena admit. (1992b: 205)[143]

Carmen offers the following explanation of the situation.

---

[141] For evidence of the latter see his aptly named, 'Caveat Emptor,' where Dennett replies to Mangan's criticism of *Consciousness Explained* as follows: "I take pleasure in confirming one of [Mangan's] charges. He accuses me of deliberately concealing my philosophical conclusions until late in the book, of creating a "presumptive mood," of relying on "rhetorical devices" rather than stating my "anti-realist" positions at the outset and arguing for them. Exactly! That was my strategy … Had I opened with a frank declaration of my final conclusions I would simply have provoked a chorus of ill-concealed outrage and *that* brouhaha would have postponed indefinitely any remotely even-handed exploration of the position I want to defend" (1993: 49). That said, there are also suggestions that Dennett wishes to distance himself from his earlier stance.

[142] Cf. Jackson (1993: 899) who comments, "Dennett has long resisted those who demand a definite answer to whether he is or is not a realist about mental states, and, correspondingly, to whether he is or is not an instrumentalist about the mind." See also Clark 1992.

[143] Note though the scare-quotes around "realist".

> Dennett is often criticized for denying the existence of consciousness, and he has often denied denying it. Disagreement of this sort can arise, it seems to me, only because his view equivocates between on the one hand a subtle and plausible challenge to a cluster of traditional assumptions about consciousness, and on the other hand a boldly counterintuitive conjecture that challenges not just expert opinion but also common sense, indeed manifest appearance (assuming there is such a thing). (2007: 100)

I think that there is something quite right about this. Dennett does indeed offer "a subtle and plausible challenge to a cluster of traditional assumptions about consciousness". To that extent, I aspire to be a fellow traveller. A central aim of this thesis is to challenge certain traditional assumptions about temporal consciousness, which prevent us from making sense of various phenomena such as the perception of succession or constant motion. That said, if Carmen's subtle Dennett is to be an ally in a defence of a sophisticated realism, we must at the very least tackle the arguments of the anti-realist Dennett that Carmen and so many other commentators have found. If we can show how a subtle realism can fully account for the problematic phenomena Dennett adduces without embracing anti-realism, we will have demonstrated that Dennett *should* have changed his 1979 view, whatever the truth of the matter.

The rest of the chapter is structured as follows. In §§3-6, I introduce the Dennettian considerations which have courted the anti-realist charge and show how they can be marshalled into an argument threatening a form of anti-realism about temporal phenomenology, drawing on resources from Crispin Wright's *Truth and Objectivity*. Whether or not the argument is Dennett's own is not of great importance; my aim is only to show that the argument can be made, and made forcefully. In doing so, I raise (without attempting to resolve) a number of subtler issues related to the formulation and tenability of anti-realism in the phenomenological domain.

In §§7-9, I argue that the reconstructed Dennettian argument only threatens an anti-realist conclusion if we make an assumption about the metaphysics of experience – viz., that experience is homoeomerous down to very small durations. If experience is significantly anhomoeomerous like many of its sibling processes, all such arguments collapse. Someone who finds Dennett's other assumptions plausible, but who finds the anti-realist conclusion unwelcome, should encounter welcome refuge in anhomoeomery. Indeed, if I am right, realism *requires* anhomoeomery. In the final section, I return briefly

to exegesis. I suggest that many of Dennett's explicit claims about time and experience are in fact quite compatible with the account defended; where disagreement may lie is rather with Dennett's quite independent views concerning phenomenal judgement (§10).

3. Visual Masking

Dennett offers a number of considerations designed to overturn our 'intuitive' or natural ways of thinking about experience in time.[144] Those which have, in my view rightly, garnered the most attention, involve appeals to psychological phenomena which arise at short timescales. Backward visual masking provides the simplest introduction.

Backward visual masking is typically described as the reduction and even elimination in visibility of a target object, brought about by the presentation of a second masking object a short time after the presentation of the target. Enns and Di Lollo introduce the phenomenon as follows.

> [A] target that is highly visible when presented briefly by itself can be rendered completely invisible by the subsequent presentation of a non-target object in the same (or nearby) spatial location. 'Backward masking' of this kind has its strongest influence not when target and mask objects are presented simultaneously, as intuition might suggest, but rather when a brief temporal gap is inserted between the presentation of the target and the mask. (2000: 345)

In metacontrast masking, a form of masking "which involves closely adjacent but non-overlapping contours" (ibid.), a plot of visibility versus stimulus onset asynchrony (SOA), i.e., the temporal gap between target and mask presentation times, gives a roughly U-shaped curve, with visibility at its lowest at an 'optimal' SOA of about 50-100 ms. A wide-range of parameters are involved in determining this function and certain effects can be found in some paradigms with SOAs as long as 200-300ms. However, in what follows, we can focus on the following, highly simplified case.[145]

---

[144] I do not discuss the other considerations Dennett brings to bear. Chalmers (1996: 360, fn.9 and references therein) offers powerful reasons for thinking that the realist should be unfazed by the arguments in Part Three of *Consciousness Explained*.

[145] Todd 2006 provides a useful introduction to masking as well as trenchant criticism of Dennett and Kinsbourne (1992a) for failing to consider the details and range of masking phenomena.

The case to be considered involves a brief (10ms) presentation to a subject of a disc, followed very shortly afterwards (SOA = 70ms) by a similarly brief presentation of a ring in the same location. The ring's inner circumference is just larger than the disc's circumference (as below) and both presentations are at constant luminance.

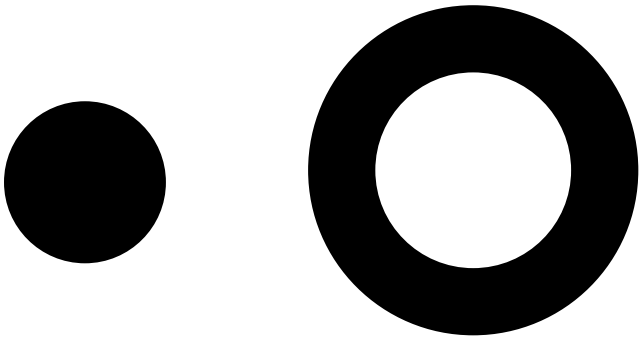

First stimulus Second, masking stimulus

As Dennett correctly notes, the standard understanding of backwards visual metacontrast masking in the literature is that "the second stimulus somehow *prevents conscious experience* of the first stimulus" (1991a: 142; his emphasis).[146] Certainly, this much seems true (and I assume that it holds in the set-up here in question), those subject to such an experiment consistently "swear that there was only one stimulus: the ring" (141).[147] In other words, if they are genuinely conscious of the ring, this is not something that they are in a position to attest to.

Dennett notoriously suggests two rival accounts of masking phenomena, a Stalinist account and an Orwellian account. The Stalinist account claims that only the ring is

---

[146] Alpern (1953) is typically credited with the first detailed examination of metacontrast masking. However, his results only indicate that the "brightness of a flash of light is *reduced* when it is followed by a second flash in an adjacent region of the field" (1953: 648). Dennett, although he cites Alpern, discusses a case of complete, first stimulus elimination masking. For such cases and for evidence that they are standardly interpreted in the way Dennett suggests, witness: "Some very efficient metacontrast masking conditions can be used even to completely block a visual test stimulus from accessing conscious perception: Metacontrast masking can be so strong that the test stimulus can no longer be discriminated and becomes invisible" (Ansorge and Heumann 2006: 62). See also Enns and Di Lollo 2000: 345 above and Todd 2006: 480 which talks of "the striking invisibility of the disc at optimal SOA".

[147] Neumann and Klotz (1994) describe masked stimuli as "non-reportable". Lachter et al. (2000) state, "With … backward masking observers claim *no* conscious experience of the first stimuli." Todd (2006: 481) avers that "at the optimal SOA the subject will report that [the target] is not visible".

perceived and posits a delay in consciousness to explain the missing first disc; the Orwellian account eschews delay and instead claims that both disc and ring are perceived but the disc rapidly forgotten. Dennett then argues that there is nothing that can decide between these two accounts.

Why might the seeming availability of these two accounts threaten anti-realism? A potentially helpful way to construct the argument is to exploit Wright's notion of Cognitive Command. According to Wright a discourse about some given domain exerts Cognitive Command just if:

> It is a priori that differences of opinion formulated within the discourse, unless excusable as a result of vagueness in a disputed statement, or in the standards of acceptability, or variation in personal evidence thresholds, so to speak, will involve something which may be regarded as a cognitive shortcoming. (1992: 144)

Very plausibly, realism concerning the domain of some discourse is committed to the discourse meeting this constraint. Thus, as Wright puts it, "show that a discourse lacks it and you will blow away with one stroke all conceivable forms of realist resistance" (148). In the case at hand, the question is whether a subject's own discourse about their phenomenal stream exhibits Cognitive Command. In particular, we are interested in whether there can be (and what we should say about) differing accounts a subject might consider concerning their own mental life.[148]

With the notion of Cognitive Command at hand, we can construct an anti-realist argument as follows.

i. Assumption: Cognitive shortcoming on the part of the subject of experience is ruled out a priori with respect to their own current experience.[149]

[148] Non-subjects will evidently exhibit cognitive short-coming given that they do not have the privilege of undergoing the stream in question.
[149] I consider what exactly 'cognitive shortcoming' amounts to shortly. I do not mean to suggest that we cannot make mistakes about our own, current stream of consciousness.

ii. Thus, assuming Cognitive Command, competing accounts of a subject's own phenomenology must also be ruled out a priori or excusable as a result of vagueness etc.
iii. But, competing accounts of a subject's own phenomenology can be found which are not so excusable (i.e., Orwellian/Stalinesque accounts of the masking case).
iv. Thus, the assumption of Cognitive Command is false; discourse about the phenomenal does not exhibit Cognitive Command.
v. Hence, realism cannot be true.

This reconstruction makes good sense of one standard response to Dennett. An exemplar is Tye (1993) who takes Dennett's target to be what he calls Phenomenal Realism, "the view that there really is such a thing as phenomenal experience or consciousness, conceived of as distinct from judgement or belief" (1993: 893). As Tye reconstructs him, Dennett proceeds as follows.

> Dennett holds that if Phenomenal Realism is true, then there are two possible alternative accounts (one Orwellian and one Stalinesque) of the relationship of experience to belief or judgement. Since these accounts ultimately do not differ, according to Dennett, the conclusion he draws is that Phenomenal Realism is false. (894)

To this Tye responds thus:

> [T]he defender of Phenomenal Realism will balk. Is there really no possible way for us to decide between the two accounts? And even if there were, would that establish that there was no genuine difference? After all, by hypothesis, the Orwellian account admits a very brief phenomenal experience of [the first disc] while the Stalinesque story does not. This looks like a pretty definite difference, whether or not we are ever smart enough to devise scientific experiments, the results of which favour (by standard scientific criteria) one account over the other. To deny that this is the case is to embrace old fashioned verificationism. And to this I am inclined to reply by modifying a remark of Mr. McCawber: "Verificationism! Foul play, sir! Take a drop more grog and you'll get over the weakness of believing in verificationism." (ibid.)

In other words, Tye rejects premise (i) and insists that cognitive shortcoming is perfectly possible with respect to one's own current experience.[150]

Now, as already noted, Dennett pleads guilty to verificationism as entirely appropriate in the 'inner' realm. Tye understands that special pleading in terms of Dennett's *identifying* phenomenal seeming with judging or believing. Given that assumption, Orwellian and Stalinist accounts will, of course, collapse since the judgements made are the same according to both accounts. However, Tye now complains that such special pleading is simply to assert the falsity of Phenomenal Realism.

> The first major problem here is that, with this reasoning in place, Dennett's main argument against Phenomenal Realism becomes redundant. For the defence of the second premise of that argument now rests on a claim which itself directly entails that Phenomenal Realism is false. The second obvious difficulty is that no advocate of the Orwellian or Stalinesque story would accept Dennett's identification of phenomenal seeming with judging or believing. (895)

I briefly return to Tye's interpretation of Dennett in §10 below. For now, I want to draw on resources developed in Chapter Three, to show that a version of premise (i) which avoids identifying judging and seeming can be formulated, and should be endorsed.

In Chapter Three, I suggested that it does not make sense to think of us as being 'self-blind' to a genuinely conscious aspect of our experience. It makes no sense, I claimed, to think of there as being conscious, experiential properties which are simply beyond the ken of their subject. What I labelled Self-Intimation holds: If a subject is undergoing perceptual experience with a certain experiential property, then that subject is in a position to know that they are undergoing perceptual experience with that property simply in virtue of so undergoing.

I acknowledged that subjects sometimes fail to exploit their evidential position (perhaps through inattention or irrationality). In such cases a subject's judgements may come apart from how things genuinely seem to them. However, it is not plausible to hold that

---

[150] See also Lycan (1992: 217) who complains, "Can't we finally call a halt to these creaking verificationist arguments?" and many others, for example, Block (1993: 190) and Lee (2007: 364).

subjects in masking cases are systematically failing to exploit an evidential position that they are nonetheless in. There is no positive reason to think that subjects occupy such an evidential position. Nor do they fail any necessary condition for report, which would plausibly explain their lack of report despite the presence of conscious experience. The only plausible thing to say is that the only stimulus they are in a position to know about is the second stimulus, the ring. Given this, we can ground the standard presumption in the literature, viz., that subjects have no conscious awareness of the first stimulus disc, in an implicit commitment to Self-Intimation.[151] This kind of reasoning holds quite generally. Thus, Self-Intimation commits us to a version of premise (i) on which *systematic* cognitive shortcoming on the part of the subject of experience is ruled out a priori with respect to their own current experience.[152]

Tye then cannot simply dismiss the anti-realist argument on the basis that it identifies judging and phenomenal seeming. If he wishes to reject premise (i) he must reject Self-Intimation and cling on to his extremely strong Phenomenal Realism, which insists on the total mutual independence of self-consciousness and consciousness, i.e., allows that a subject can undergo phenomenal experience of some kind without being in any position to know about it. As Chapter Three made clear, that is a serious bullet for the realist to bite.

---

[151] To this extent we will want to explain better than chance performance in forced-choice guessing as to whether one stimulus was presented or two by appealing to unconscious processing of the disc. See Dennett 1991a: 142. There is nothing especially problematic about this.

[152] As discussed in Chapter Three, strictly we should allow for occasions of having a certain experience where one is not in a position to know the nature of one's experience despite the experience being of a *kind* such that knowledge of its nature is to be had solely in virtue of undergoing it. This raises the following challenge. Why not think that, in the masking case, subjects do enjoy disc-experience even though they aren't in a position to know it and, moreover, think of the non-masked case precisely as an occasion where subjects enjoy experience of the same type (i.e., of a disc) but *are* in a position to know it? Two significant disanalogies with the cases discussed in Chapter Three show why this line of thought is not persuasive. In the cases in Chapter Three, margin for error principles were invoked on the basis that a single situation plausibly afforded two distinct kinds of experience, entailing that in such a situation one was always too close to having had another kind of experience to obtain knowledge. Here, in contrast, the suggestion is that one actually always experiences both disc and ring but is, for some unspecified reason, too close to an experience of a ring on its own to know what one experiences. Note firstly how, in the masking set-up, one actively judges that one only saw a ring, one does not simply withhold judgement on the basis that one cannot tell whether one's experience is a disc-then-ring experience or a ring-alone experience. Secondly, note that no explanation of closeness seems available which would allow margin for error principles to get a grip; it is precisely not the case that the same situation affords both ring-alone experience as well as disc-then-ring experience.

Block (2007) argues forcefully for this kind of strong Phenomenal Realist picture and in doing so defends an Orwellian line. Block's paper has received much criticism (as well as praise) so let me confine myself to just three very brief points.

(i) Block's general view is that neurological considerations are weighty when it comes to reaching psychological conclusions. Be this as it may, his conclusion that we should think of consciousness as present whenever a given, local part of the visual system is activated, quite independently of the activation of cognitive systems, must seem perplexing to anyone who thinks of conscious experience as essentially a way of finding out about our environment – as a *cognitive achievement* of some kind. The idea of sufficient neural conditions, in principle independent of cognitive capacities, is tempting only insofar as this conception of conscious experience is rejected.[153]

(ii) Sperling's experiments (Sperling 1960) have long formed the basis of an objection to Self-Intimation type-theses – in addition to Block, see Dretske (2006) and Tye (2006).[154] However, in their original guise at least, they are subject to precisely the kinds of worries that I argue masking and apparent motion experiments face below (see §7f.).[155] To anticipate discussion to come, Sperling assumes that one's visual experience of the grid will be the same irrespective of which tone is played immediately after the grid is displayed. If we reject homoeomery, this cannot be assumed, and so the partial reports cannot be summed as Sperling suggests. It may be that the events 'grid-followed-by-high-tone' and 'grid-followed-by-low-tone' are different events with different visual (as well as auditory) appearances.

(iii) The more recent developments of Sperling's paradigm that Block discusses may avoid this concern to some extent, but they also introduce other

---

[153] For Block's conception see in particular his 1995.

[154] Sperling's paradigm involves showing subjects a number of letters, say twelve, arranged in a rectangular grid. Asked to fill-in whole blank grids after the experiment, subjects consistently report approximately four letters in their correct positions (a finding which holds regardless of stimulus complexity). Sperling's question was whether this "limit on the number of letters that can be correctly reported" was also a limit of visual representation. To test this he played a brief tone very shortly after the grid had ceased to be shown to participants. The pitch of the tone indicated to the subject which row was to be reported on. These partial reports were invariably accurate (assuming the row contained only three or four letters). Sperling's conclusion was that this evidenced a visual representation of (almost) *all* the letters – after all had a different tone been played after the grid was removed, subjects would have reported equally well.

[155] Dennett and Kinsbourne (1992a) very briefly allude to this.

> complications.[156] As Byrne et al. (2007) note, their use by Block depends on closing a crucial gap between visible and informational persistence. However, Block only succeeds in closing the gap by appeal to a tendentious claim regarding what subjects report (and so see), viz., that, in Sperling-type experiments, for each and every determinate character in the grid, subjects are phenomenally conscious of it *as such*. As is often pointed out, it may well be that subjects are merely conscious of a twelve letter grid and only a few individual determinate letters as such.[157]

If the anti-realist (quite correctly in my view) insists on Self-Intimation, another consequence immediately arises, one that the realist might seize upon. For, although Dennett argues that there is nothing that can decide between Orwellian and Stalinesque accounts, in one respect the above discussion already has decided between the two accounts. For only the Stalinist account concurs that the first disc is not seen. In contrast, the Orwellian claims that the first disc is experienced but rapidly forgotten. In a sense this is right. I do not think that we need any "crucial experimental result" (Dennett 1991a: 142) to settle the dispute. The court of introspection is quite sufficient to settle matters. However, this need not mean advocating Stalinism. So far we are only committed to one aspect of the Stalinist account, viz., that the first disc is not experienced in the masking scenario. In fact, I think that both accounts are mistaken; they represent a false dichotomy.

Nonetheless, if the disc is not seen, this leaves us with a puzzle as to why Dennett holds that the realist is committed to the existence of competing accounts as opposed to simply endorsing Stalinism. The anti-realist can hardly appeal to Block-style considerations (persuasive or not), for they amount to a rejection of the picture developed in Chapter Three and in particular, of Self-Intimation. Since the anti-realist argument above relies on that picture,[158] Block-style considerations are quite out of place in pressing anti-realism. The best way to understand the dialectical situation, I suggest, is rather in terms of an Orwellian counter-argument, which attempts to show that the first disc *must*, after all, be perceived. We can then understand Stalinism as an attempt to preserve the invisibility of

---

[156] I have in mind Landman et al. 2003 and Sligte et al. 2006, 2008.
[157] See, for example, Papineau 2007 on scene phenomenology.
[158] Or, rather, that picture stripped of its implicit realist bias, see below (§5).

the disc in the masking trial *whilst acknowledging the central plank of the Orwellian challenge*. A stand-off arises because the costs of Stalinism parallel the costs of Orwellianism. As I have said, my reply to all of this is that both positions are premised on a false assumption.

4. The Orwellian Challenge

The Orwellian challenge can be developed as follows. The Orwellian first notes something to which everyone agrees, namely, that, if no second stimulus had been presented to the subject in the above masking trial, then they would have seen (and been in a position to report) the first disc. However, given this, it seems that we face a dilemma. Either the first disc is seen in both cases or consciousness of the first disc is delayed by at least 70ms quite generally – allowing time for the second stimulus to 'block' the conscious presentation of the first.[159]

The apparent dilemma is troubling. If the disc *is* seen, then the considerations relating to self-awareness above must be mistaken. On the other hand, a delay in consciousness of 70ms (which is the full-blown Stalinist proposal) is implausible since, as Dennett puts it, "There is abundant evidence that responses under conscious control … occur with close to the minimum latencies (delays) that are physically possible" (122).[160] The Stalinist must, it seems, quarrel with this. He can do so by arguing that, what retrospectively seem to be conscious responses within such a temporal window, are in fact unconscious. But now, the attractiveness of the Stalinist's initial move, viz., an endorsement of a picture of the mind which seems to be as it is, has been abandoned.

In other words, the Orwellian dilemma seems to force us to reject a background picture of self-conscious awareness (according to which conscious properties are self-intimating and on which we cannot be systematically deluded about our own conscious lives). On

---

[159] The precise time is not important for present purposes but it is clear that the delay would in fact need to be much longer. Optimal masking occurs with SOAs of up to 100ms. However, reduced visibility masking occurs with SOAs of up to 300ms and no less demands explanation. Indeed, as Todd (2006: 480) notes these effects are of "most theoretical interest to researchers". Hence, the Orwellian argument can appeal to these longer times. What is important is that the delay posited is longer than seemingly consciously controlled response times (plus processing, preparation and pulse travel times) – see below.

[160] That this evidence is compelling is a large assumption, especially at this timescale. However, I propose to grant the implausibility to Dennett since I think we can vouchsafe realism even so granting.

the first horn, we must accept that the disc is seen, despite this being something that we cannot be aware of. On the second horn, we commit to a delay in conscious awareness, a commitment which turns out to involve a commitment to there being an aspect of our conscious lives about which we are systematically deluded. We consider rapid – but not reflex – responses to stimuli to be under our conscious control; we believe our actions to be grounded in conscious awareness of their instigating stimuli. Yet, although our inner lives appear like this, they are not in fact like this. We are systematically deceived that a certain range of our actions are responses grounded in conscious awareness of the world, despite this being a pure fiction.

If we are wedded to the picture of self-consciousness developed in Chapter Three, we appear to be in an inescapable bind. For both the Stalinist and Orwellian accounts involve abandoning a core commitment of our understanding of our conscious lives. Having elaborated this predicament, Dennett (as here construed) in effect suggests that the culprit is Cognitive Command. It cannot, as premise (i) insists, be that subjects exhibit cognitive shortcoming in their inability to decide between the Orwellian and Stalinesque accounts; nor can the difference be put down to vagueness. Consequently, it seems that discourse about the phenomenal does not exhibit Cognitive Command and realism must be rejected.

Put another way, our difficulties arise from a realist commitment to there being a matter of fact as to which account, Stalinesque or Orwellian, correctly describes our stream of consciousness, and further to the principle of Self-Intimation which declares that such a difference is always available to subjects. We can avoid contradiction if we deny that there is really any matter of fact as to which account is correct. I explore how this idea should be developed in the next section.

5. Exploring Phenomenal Anti-Realism

A passage often cited by critics looking for a confession of anti-realism in Dennett is the following.

> There is no such phenomenon as really seeming – over and above the phenomenon of judging in one way or another that something is the case. (1991a: 364)

We have already seen here that Tye takes this to amount to an "*identification* of phenomenal seeming with judging or believing" (1993: 895; my emphasis). Similarly, Carmen tell us that "Dennett's eliminativist theory of consciousness rests on an implausible reduction of sensory seeming to cognitive judgement" (Carmen 2007: 99). We have also seen, however, that constitutive connections between phenomenal experiences and their representation (or how they seem to us) can be made out which allow for judgement to come apart from how things seem. In particular, Self-Intimation offers an understanding of how things seem in terms of a subject's epistemic position whilst allowing for failures to exploit that position in judgement. However, Self-Intimation is an unabashedly realist claim. It insists that subjects are positioned to know that they are undergoing a perceptual experience with certain properties *in virtue of so undergoing*. This asymmetric relation makes explicit the explanatory priority afforded to experience itself, conceived of as independent of and prior to a subject's epistemic position.

Martin gives voice to this way of thinking about the relation between consciousness and self-consciousness.

> It is not clear that we can quite conceive of the cognitive aspect [of our cognitive states of response to phenomenal consciousness] as other than awareness of one's mind, and hence a form of, or ground for, knowledge. In which case, one could not be this way cognitively (i.e., with this range of judgements, formed in this way) without being self-aware and self-knowing: phenomenal consciousness would have to be present. On this picture, although the *facts about phenomenal consciousness obtain independently and prior to any facts about our knowledge of it*, our introspective cognition of phenomenal consciousness need not be independent of that consciousness: seeming awareness of one's conscious mind will always be genuine self-awareness. (Martin 2006: 377-8; my emphasis)[161]

Our perceptual awareness can easily be conceived as merely apparent awareness of the world. The senses that provide us with perceptual knowledge can, on occasion, fail to be

[161] One should not assume that this applies to all cases of sensory awareness on Martin's picture. For in the case of causally matching hallucinations it is *not* clear that he thinks of facts about phenomenal consciousness as being prior to epistemological facts. Here rather we seem to have a no priority view; see for example his 2004: 84 and below.

a route to such. In contrast, Martin suggests, our introspective awareness cannot ever be considered as merely apparent awareness. Nonetheless, as Martin describes the constitutive relation between self-awareness and phenomenal consciousness, it is firmly realist: "facts about phenomenal consciousness obtain independently and prior to any facts about our knowledge of it". This, I suggest, is the core claim of phenomenal realism, and it is quite consistent with Self-Intimation. We might think of it like this: being in a conscious state brings with it the potential for knowledge that one is in that state. Thus, our epistemic situation is explained in terms of our experiential situation; experience is the *ground* of our epistemic position.

Anti-realism is sometimes *characterised* as the denial that truth in some domain outstrips what it is warranted to assert in that domain or, following Wright, what is *superassertibile* in that domain, where a statement is superassertibile just if "it is, or can be, warranted and some warrant for it would survive arbitrarily close scrutiny of its pedigree and arbitrarily extensive increments to or other forms of improvement of our information" (Wright 1992: 48). Now, riding rough shod over the finer details, the above discussion suggests that, even for the realist, truths about a subject's current phenomenal state never outstrip warranted assertibility.[162] Self-Intimation tells us, after all, that subjects are *always* in a position to know of, and so warrantedly assert concerning their conscious experience.

Consequently, contrasting assertibility or superassertibility with truth does not serve to distinguish realism and anti-realism in this domain.[163] In that light, and following Wright (1992: 79), we might think that the place to look for guidance on how to frame and prosecute the realism/anti-realist debate is to the famous contrast pointed out by Socrates in the *Euthyphro*. Are certain acts pious because they are loved by the Gods, or do the Gods love certain acts because they are pious? As Wright notes, all parties to that debate agree that an act is pious iff it is loved by the Gods. Debate remains because one

---

[162] One important complication arises from the discussion of disjunctivism above. According to the disjunctivist, it remains true that there can be no *positive* characterisation of our experience which outstrips warranted assertibility. However, for the disjunctivist, the attribution of presentational properties to experience in hallucinatory cases is warranted but false. The disjunctivist must therefore find some way of distinguishing between cases where the range of judgements is the same and yet their truth different. That can only be achieved by committing to a form of realism. This is perhaps not surprising. Indeed, below, I suggest a way in which our thinking about perception more generally does prejudice the debate in favour of realism. Nonetheless, for now, I grant that in the inner realm truth and warranted assertibility (or some such concept) line up. I do not want to argue from disjunctivism to realism. Even if that argument is perfectly sound, it is unlikely to have much suasive force in the current context.

[163] See Wright 1992: 78.

side insists that this biconditional holds because of a sensitivity that the Gods have to what is independently pious, whereas the other side insists that what the Gods think somehow *constitutes* piety.

The contrast in direction of explanation is a place we might hope to locate the realist/anti-realist contrast with respect to the phenomenal. If this is right, we need to find a biconditional agreed upon by both parties. We can then consider which direction of explanation (if any) should have priority. As already noted, the bi-conditional cannot be Self-Intimation given its obvious realist bias. The *Euthyphro* question cannot for example be: is a property experiential because a subject is in a position to know of its instantiation just in virtue of undergoing the relevant experience, or is it because the property is experiential that the subject is in such a position? This formulation begs the question in favour of realism. Read in what might seem to be the anti-realist's favoured direction, the claim is that experience's having a certain experiential property can be explained in terms of what a subject is in a position to know in virtue of undergoing that experience. But here the cognitive state which is supposed to be explanatorily prior turns out to be one which obtains only in virtue of undergoing a certain experience. Yet our only grip on *which* experience is in question comes via the experiential property in question which the experience is claimed to have. Thus, it transpires that the cognitive state is not explanatorily prior after all.[164]

The obvious remedy here is to include experiences themselves in the biconditional. Thus,

> A subject *S* is undergoing a course of experience with experiential property *P* just if *S* is in a position to know that they are undergoing experience with *P*.

This avoids the objection above, but a second worry remains because of the use of the verb 'to know' in the formulation. To invoke knowledge may appear to beg the question in favour of the realist since it implicitly brings with it the notion of justification. However precisely we answer the question, 'What is it to know that *p*?' it is hard to avoid the idea that it involves being *justified* in believing that *p*, even if that justification is not available to one, or is not characterizable independently of knowledge. Above I suggested

---

[164] Note how in the *Euthyphro* case all parties implicitly agree on a *piety-independent* mode of identifying acts. This is an implicit agreement we might well challenge.

that, what justified introspective judgements, was simply having the relevant experience. However, the anti-realist cannot say this since the subject's knowledge is supposed to be explanatorily prior to the experience. Yet without experience as the source of justification, it is hard to see where justification might come in. Whether or not that is a worry in itself (a point briefly discussed in the next section), it is certainly a worry if our formulation simply helps itself to the notion of knowledge.

The remedy is to remove the offending verb. Thus, the anti-realist will ultimately endorse a biconditional of the following form.

> A subject *S* is undergoing a course of experience with experiential property *P* just if *S represents* it as though they are undergoing experience with *P*.

A biconditional of this form can generate a *Euthyphro* debate since it admits of two understandings. The realist will claim it holds because subjects undergoing experience with the relevant phenomenal property are thereby in a position to judge that they are. The anti-realist will claim that the biconditional holds because the subject's representation of their mental life *constitutes* what it is to have experience of a given phenomenal kind.[165]

So construed, phenomenal anti-realism is a quite general doctrine about the experiential. However, the arguments I am concerned with, i.e., those of Dennett (1991a: Chs.5-6) and Dennett and Kinsbourne (1992a), arguably speak only to anti-realism about certain aspects of our experience, in particular, *temporal* features of experience. If successful, they only directly motivate an anti-realist reading of the biconditional with respect to these experiential properties. Thus, in the first instance, they serve to establish only an instance of a general view, which Dennett then proceeds to defend. This is not to say that no spreading argument can be provided. For example consider Anthony.

---

[165] Here compare Kant's view that "every advance of perception, no matter what the objects may be, whether appearances or pure intuitions, is nothing but an extension of the determination of inner sense" (2003: A210). As Martin (1998a: 119) notes, it is not obvious that we need to understand Kant here in terms of a perceptual model of inner sense (despite the connotations of that way of speaking). That is, of course, not to claim that we should not attribute an inner sense model to Kant; there is undoubtedly evidence in favour of such an attribution.

> [A first] interpretation [of Dennett and Kinsbourne] takes them as straight-forwardly denying that conscious experiences are temporally located, that they begin, persist for some interval, and end. Clearly this interpretation leads directly to eliminativism, given materialistic assumptions. For *every* physical state, event, process, and so on, is temporally located in this sense. If there exist conscious states and processes, therefore, they are as well. (1992: 201-2)

In the context of Dennett's materialism, Anthony sees no way of stopping at a mere temporal anti-realism. If this is right, beginning with temporal issues is a dialectically excellent choice of starting point for Dennett, if his larger ambition is the defence of a quite general phenomenal anti-realism.

Anthony's interpretation is questionable. Dennett holds that experiences are informational/neural states or events. Therefore, there is no question of their lacking temporal properties. There is the question, however, of whether those properties are relevant from a psychological point of view; in other words, of whether these properties of such events are to be thought of as *conscious*. In that light, temporal anti-realism can be conceived of as the view that there is nothing more to a particular neural event being *conscious* at t_1 as opposed to t_2 than when that event is *represented* as being conscious. The timing the event is represented as having is quite independent of the time of its neural tokening.[166] Consequently, temporal anti-realism is (at least, prima facie) quite compatible with other aspects of the stream of consciousness obtaining independently and prior to their recording as such.

A spreading argument can be found if one thinks that experiences, if they exist at all, are *essentially* dynamic events in the stream of consciousness. If that is right, then an insistence that the dynamic nature of such events must be construed in anti-realist fashion entails that an anti-realist account has to be given of experience itself. If there is nothing more to experiential time-structure than the structure recorded in working memory, that rules out there being experiences as understood by the realist, since the realist's experiences must essentially have a time-structure conceived of as distinct from a mere record of such.

---

[166] Cf. "…content-fixations in the brain are precisely locatable in both space and time, but their onsets do *not* mark the onset of consciousness of their content" (Dennett 1991a: 113).

Dennett and Kinsbourne clearly do intend a general conclusion, noting that, though they concentrate first of all on temporal features, their 'Multiple Drafts' model applies to all features of subjective experience. However, it is quite possible to defend a restricted form of anti-realism about certain features of experience. In particular, one might defend *temporal* anti-realism, a view that we have seen anticipated in the discussion of film and again in the discussion of §1 of this chapter.

In the next section, I suggest that even if we think that full-blown phenomenal anti-realism is a deeply unattractive view, restricted forms of anti-realism need taking far more seriously. Quite what that restriction should be is unclear. Dennett presents his arguments as concerned with temporal features of the stream of consciousness, considering them to provide the *best* case for the anti-realist. However, the cases he considers also, arguably, testify in favour of anti-realism about related aspects of spatial and colour phenomenology (e.g. the *presence* of a masked stimulus or colour change). I focus just on *temporal* features for two reasons. Firstly, such a focus provides a relatively simple and well-motivated restriction with which to appreciate the theoretical lie of the land, even if the restriction is, in the end, less clear cut. Secondly, the basic phenomenon in all the cases that Dennett adduces does seem to be a broadly temporal one. Even where the fact to be given an anti-realist interpretation concerns the presence of, say, a spatial feature, as opposed to the temporal ordering of certain events, what motivates the anti-realist interpretation ultimately comes down to a question of 'settling contents' into temporal structures. Thus, my interest in the next section is to further motivate *temporal* anti-realism and so show why the realist urgently needs to resist Dennett's arguments as I do in §§7-9. That said, the resistance mounted in §§7-9 in no way depends on accepting this line of thought.

6. Pursuing the Debate

Given a neutral bi-conditional of the above form, how might the debate be moved forward? This section tentatively sketches a number of approaches, noting how the temporal anti-realist has much more to say in response to standard realist objections than the full-blown phenomenal anti-realist.

One idea floated by Wright is to consider whether "the decidability of the statements in question [is] assured only as a matter of contingency or – closely relatedly – whether the coincidence of truth with superassertibility within the discourse [is] something that could be known only a posteriori" (1992: 81). This is a bad idea in the current context. The biconditional (if it holds) is arguably something that we can know a priori insofar as that test is a clear one in this context. But it is hardly plausible to think that this provokes a crisis for the realist. As Wright is well aware,

> Even in cases where decidability is guaranteed a priori – for instance, statements describing one's own sensations … – there ought to be space for a contrast between the thought that truth-conferring states of affairs are necessarily *detectable* by an appropriate subject, and the thought that they are, rather conceptually grounded in the judgements of such a subject. (1992: 82)[167]

Thus, as Wright puts it, for the realist the appropriate

> line of attack would [in fact: must] be to show that, even though extensional divergence [of truth from superassertibility] is excluded a priori, it is proper to think of the truth of a statement of the discourse in question as the *explanatory ground* of its superassertibility (1992: 143)

In pursuit of this strategy, realists are likely to make three broad objections to anti-realism. First, the realist might insist that, pre-theoretically, we find it quite *obvious* that our phenomenal judgements are grounded in, or rationalised by, the subjective character of our experience. 'Why did you judge that you were having an experience as of a tiger?' 'Because I *was* having such an experience – and I know what that's like!' 'Why did you judge that your bone felt broken?' 'Well, because it *did* feel broken.' In contrast to this natural picture, the full-blown anti-realist must claim that we simply find ourselves inclined to judge in certain ways without being able to offer any reasoned explanation why we so judged. Not so the temporal anti-realist. The reason is that the temporal anti-

[167] Later Wright remarks that the Euthyphronist is committed to the necessitation of the biconditional. In contrast, if detectivism [realism] is right, "it seems that it ought to be a possibility that the causal order be so constituted that opinions formed under the conditions which, as things are, ensure that they track the facts, might not have done so" (1992: 112). Thus, "at least a potentially sufficient condition for the propriety of the detectivist view" is that the biconditional is "at best contingently true". Merely sufficient, however, given that this possibility is not one we will wish to countenance in the inner realm (see above).

realist can appeal to non-temporal features of experience in order to ground the phenomenal judgements in question. This was the approach introduced in §1 of this chapter in relation to duration judgements. Moreover, in light of that discussion, the temporal anti-realist can cite duration judgements as demonstrative of the coherence of such an approach. That is not to disagree with the realist's thought that realism is our natural, default view. It is, however, to counter the impression that the temporal anti-realist has nothing to say regarding the basis for our temporal phenomenal judgements.

Second, the realist might argue that the anti-realist has no obvious way of explaining why only some properties and events figure in our experience. Hearing a cuckoo or feeling a lover's touch are experiential events, digesting a tuna sandwich or repairing damaged skin cells (typically) are not. Whence the difference? According to the realist, the answer is that the former are events with phenomenal properties. It is hard to see what the anti-realist can say other than once more, and rather unsatisfactorily, that we simply find ourselves making judgements with some contents and not others. The temporal anti-realist is able to provide a more convincing reply. The reply appeals to the existence of non-temporal phenomenal properties (construed in a realist manner) in combination with the unobjectionable notion of autobiographical memory. According to the anti-realist account, temporal properties are introduced into the representation of the stream of consciousness to make sense of the experiential contents that we find ourselves with.[168] For example, in apparent motion cases, we find ourselves with an experience of a disc in place A, in addition to one of a similar disc in place B. Our autobiographical memory then constructs the fiction that we experienced one and the same disc moving from A to B (perhaps changing as it does) to account for these multiple contents. In addition to properties conceived of in a realist manner, properties introduced in this manner are the subject matter of phenomenal judgements. However, only a very restricted range of properties are introduced in this manner, providing the required demarcation principle.

A third and final objection to phenomenal anti-realism questions how the anti-realist can handle the misjudgements of the insane or chronically inattentive, or simply judgements

---

[168] For the Dennettian, presumably the contents are those in working memory and the structuring is a matter of what gets written into short-term autobiographical memory. For a Kantian, presumably the structuring is a matter of synthesis over all experiences. This vastly more complex task is possible only because of the kind of principles established in the Analogies.

made for absolutely no reason at all. The realist thinks of such cases as resulting from failures to exploit the evidential position that experience puts one in, independently of particular judgements. In contrast, the full-blown anti-realist must introduce some kind of optimality condition to rule out 'ungrounded' inclinations to judge from being constitutive of 'genuine' phenomenal episodes. However, it is far from clear that optimality conditions are available, which succeed in ruling out bad cases and ruling in good cases without trivialising the anti-realist's thesis (and without committing the conditional fallacy). Unless some story is provided here, full-blown phenomenal anti-realism remains a promissory note.[169] However, though this is a serious concern for the full-blown anti-realist, the temporal anti-realist may freely appeal to non-temporal experiential features to explain misjudgements. Indeed, the Dennettian argument above claims to find failures of Cognitive Command, which are inexplicable on the realist account, but explained easily by the temporal anti-realist in terms of multiple narratives based on the same non-temporal contents. Thus, not only does the temporal anti-realist have an answer here; they also have a powerful rejoinder. Whatever our final verdict on the rejoinder, temporal anti-realism once again appears at least superficially coherent and defensible when put under pressure by the realist.

Aside from raising concerns about the coherence of the anti-realist's view, can the realist do anything positive to press his position? Wright's most appetising offering for the realist is the following.

> Where it is possible, without mention of human judgement or the conditions under which, in the case in question, such judgement would be best, either fully to analyse, or at least to draw attention to the general characteristics of the truth-conferring states of affairs in such a way that it is a *consequence* that there is an a priori guarantee that best opinion will be on track, then it is appropriate to think in terms of infallibility. Where it is not, … then the notion of (partial) extension-determination has its proper place. (1992: 124)

---

[169] Note how a form of anti-realism or response-dependency concerning properties like that of being stylish or fashionable seems in better shape in this respect. In those cases scope for reasons might be argued to enter by way of the community. Thus, we might evidence the claim that kitten heels are *en vogue* by watching people on The King's Road, reading *Cosmopolitan*, or by consulting Anna Wintour. No such options exist with respect to our own experience. I cannot consider what other people's experience is like, read a text book, or consult an expert; I am the only expert.

Now we need not think quite in terms of *infallibility* for this test to be appropriate. The key point is that in a case where it is a priori that a neutral biconditional holds, the above test may provide traction on the *Euthyphro* debate. So, let us ask, where the truth conferring states in question are phenomenal episodes, can we "draw attention" to their general characteristics in such a way as to avoid mentioning judgement? Let us focus on perceptual experiences as we have done elsewhere.

Reaching out to the Nagelian idea that such episodes contribute to what, subjectively, it is like to be the subject may not appear helpful. That phrase seems closely connected to a subject's epistemic position. As such it may appear to favour the anti-realist side of the debate. However, as Martin (1998b) points out, when we characterise perceptual experiences we typically reach for *objective* qualities in order to identify what the episode is like. Thus, we describe how we are enjoying the blueness of the ocean or the creaminess of the butter. Taking this datum seriously has led many to claim that the subjective qualities of our experience (what it is like to undergo our experience) are, in part, *constituted* by objective features in the world (or perhaps the representation of such). Put another way, this tradition thinks of perceptual experience as a subjective perspective on an objective world.[170]

If this tradition is right, we can draw attention to general characteristics of the truth conferring states of affairs without invoking judgement. We achieve that by identifying such episodes as presentations of mind-independent features to subjects. The idea here is, prima facie, most congenial to the naïve realist, according to whom the character of such states is literally constituted by certain mind-independent features (certain scenes). If we accept realism about the external world and the idea that perceptual episodes are literally relations between subjects and that world, materials are in play to secure a grip on experience, independent of judgement. Perceptual experience can be given a basic, judgement-independent characterisation as a relation to mind-independent particulars and their aspects. The perspectival nature of that relation can then be held to explain why subjects are positioned to know that they stand in it when they do, hence the apriority of the biconditional. In this way, the realist can pass Wright's test.

---

[170] Cf. Soteriou 2005.

As already noted, there is an obvious worry here for the naïve realist. Hallucinations are not episodes that involve mind-independent entities as literal constituents. Rather, according to Martin, causally matching hallucinations at least must be characterised in purely epistemic terms, as states that are not knowably not veridical perceptions of a certain kind. Evidently, hallucinations so treated dramatically fail Wright's test. Thus, realism about hallucinations looks untenable. There are two options here. Either one can insist on naïve realism and accept the rejection of realism as applied to hallucinations. This need not necessarily involve claiming the opposite direction of explanatory priority; instead one might deny that any direction of explanation takes priority in such cases (an assumption Wright is rightly criticised for ignoring).[171] Alternatively, one might deny the naïve realist's literal constitution claim and insist that a representationalist construal of the perspectivity datum according to which the character of perceptual experience is determined by represented features will do duty here. Consideration of the merits and demerits of that choice is well beyond the scope on this discussion.

Let me sum up. Our naïve picture of mind involves a stream of consciousness conceived of in realist terms. Moreover, a widely held view in the philosophy of perception, namely, that experience is a perspective on a mind-independent world, provides a way of gaining a judgement-independent grip on perceptual experience, thereby passing Wright's proposed test for detectivism/realism in the *Euthyphro* debate. In this light, I suggest that realism is rightly regarded as our default view in the sense that no-one would be inclined to accept an anti-realist view unless they hadn't seen how to answer certain kinds of argument. Nonetheless, in the temporal case, the realist precisely appears to lack an answer to Dennett's challenge. Moreover, a restricted temporal anti-realism appears to be a coherent alternative view which explains what is going on in the cases Dennett directs us to. Should we abandon our default, naïve view?

7. A Different Diagnosis

I am a realist. But I also think we should insist on the idea that it makes no sense to think of there being conscious, experiential properties which are simply beyond the ken of their subject. Likewise, I think we should insist that subjects cannot be systematically in

[171] For that view see Martin 2004: 84.

error about their own conscious lives. How then can one avoid Dennett's argument? In defending his Multiple Drafts model Dennett suggests the following challenge, giving his preferred response.

> Should we insist that the disc was experienced because *if the ring hadn't intervened* the disc would have been reported? That would be to make the mistake of supposing that we could "freeze-frame" the film in the Cartesian Theatre and make sure that the disc frame really did make it into the Theatre before the memory of it was obliterated by later events. (1991a: 142)

I want to suggest that there is an insight here, which can be achieved outside of the Multiple Drafts model and without embracing any form of anti-realism.

It is not uncommon to conceive of our stream of consciousness on analogy with a series of photographic frames.[172] The real point of that analogy in this context is to draw attention to the supposed fact that experience can be analysed (metaphysically speaking) into a series of instantaneous (or near-instantaneous) 'frames' or 'slices' – that is presentations or representations at an instant, each logically independent of all others. Call this principle (following Miller 1984) the Strong Principle of Simultaneous Awareness (PSA).

> **Strong PSA** Irreducibly temporal facts have *no* place in explanations of sameness and difference across experiences. Present tense psychological facts are necessary and sufficient to explain any sameness or difference between conscious episodes.

The idea behind this principle is that we can say everything we want to say about experience in terms of a series of independent instantaneous representations or presentations that together constitute our stream of consciousness.

I re-introduce and discuss this principle, its origins and motivation more fully and in its more usual context in the next chapter. For now, note how Strong PSA is tantamount to the view that experience is homoeomerous down to instants (or brief durations). To say that a process is homoeomerous, recall, is to deny that any stretch has temporal sub-parts

---

[172] Dennett sometimes also uses the metaphor of railcars travelling along a track. Similar points apply.

whose nature depends constitutively on facts about the nature of the process before and/or after the particular temporal part in question. This is just to deny that irreducibly temporal facts have any place in explaining the nature of the stream. To work towards the thought in the terms favoured in Chapter Four, consider that if experience is homoeomerous, then, if a subject is experiencing over some duration, over any sub-duration the subject will have experienc*ed*, i.e. had an experience, where the nature of that experience is constitutively independent of facts about surrounding experience. If a subject has so experienced (had such an experience) over some tiny duration, we can enquire about the nature of that achievement in isolation from subsequent experience, accounting for its similarity to and differences from other experiences in terms of facts local to the time of the experience. This precisely what Strong PSA claims can be done (either as applied to instants or at least to very short durations).

If we accept this picture, then we will think it an appropriate question to ask of any instant in the flow of the stream of consciousness, 'What is represented at that moment?' where the answer to this question will not depend in any way on what is represented at surrounding instants – just as we might stop a reel of film and enquire what image is currently in the projector, a question we should be able to answer without concerning ourselves with recent or forthcoming images.

In particular, if we accept Strong PSA, we will think it is appropriate to ask of any instant during a visual masking experiment: 'What is represented at that moment?' where the answer to this question does not depend in any way on what is represented at surrounding moments. I now argue that if we do not accept a principle like Strong PSA, i.e., if we deny that experience is analysable down to instantaneous representations or very brief intervals (homoeomery), we can escape the predicament into which Dennett corners the realist.[173]

Consider two trials of the masking experiment. On the first trial the only stimulus presented is a single disc. The subject reports seeing the disc. On the second trial the

[173] I argue in Chapter Six that we must reject Strong PSA if we are to provide an account of temporal experience. However, that argument is not independent of the issue of realism since it assumes a transparency thesis premised on the acceptance of realism. To this extent one should see the concerns as interlocking; a mistaken commitment to Strong PSA is the root evil behind many seemingly unrelated philosophical problems concerning time and perception.

same stimulus disc is presented but on this occasion a masking ring is presented 70ms after the first disc. The subject reports seeing only the masking ring.

It seems extremely natural to graph the situation in the first trial as follows. The stimulus is presented at time t_0; the subject then sees the disc at some later instant, $t_0 + \delta t$. The naturalness of this picture prompts us to ask: 'What does the subject see at $t_0 + \delta t$ in the second trial?' The above discussion of Orwellianism versus Stalinism suggests two options. Either δt is a very short period (perhaps just enough for light to reach the eyes and for the requisite visual processing), in which case the subject sees the disc at $t_0 + \delta t$ but shortly forgets about it afterwards (Orwellianism). Alternatively, δt is a longer period (70ms or more) – enough time for the disc to be 'processed out' by subsequent processing of the ring. In this case, the subject does not see the disc (Stalinism).

However, this natural way of thinking implicitly assumes that we can legitimately ask what is true at some instant of experience – $t_0 + \delta t$ – without taking into account the nature of the subject's experience at any subsequent times. That would be legitimate if experience were analysable down to instants, i.e., if Strong PSA were true. What if instead Strong PSA is false and experience is not so analysable? What if experience is anhomoeomerous and we cannot assume that what is true at some moment holds true independently of what is true over surrounding intervals of experience?

If we reject Strong PSA, then in thinking about visual masking we have no reason to assume that what is perceived at $t_0 + \delta t$ must be the same across trials. If there are no truths concerning instants of experience which hold independently of what is true during the surrounding interval of experience, then at a minimum there is no reason to assume that the experiential presence of a mask subsequent to the first stimulus is irrelevant to answering the question whether the first stimulus is perceived.[174] So, whilst it is true that if the ring hadn't intervened, the disc would have been reported, this has no bearing on the case where the ring was present. 'Has one seen the disc at t?' (where t is a moment briefly after the presentation) is not a question that one can answer independently of

[174] The independence here is from surrounding *visual* experience but this restriction might be lifted. It seems plausible that the whole unified stream of consciousness is anhomoeomerous as evidenced by cross-modal phenomena.

one's experience during the surrounding period of time. To assume that it must be is to insist that experience is analysable down to instants or at least very short periods.

If we reject Strong PSA, we can escape the Orwellian/Stalinist stand-off. Note that we began that discussion by establishing that the disc does *not* register on consciousness in the second, masking trial. To this the Orwellian objected that such lack of awareness could only be explained by positing an implausible slack in consciousness. The Stalinist accepted this slack and its unwelcome consequences. If we reject Strong PSA, we can avoid Stalinism whilst maintaining that the disc does *not* register on consciousness in the second, masking trial. For note how the Orwellian objection to this proceeded. First, it was correctly asserted that a disc *would* have been seen had it not been for the mask. It was then asserted as a consequence that, either the disc was seen in the masking trial (Orwellianism) or, alternatively, that there was a slack in consciousness (Stalinism). But this exhaustive disjunction only follows if we assume that some present tense fact at $t_0 + \delta t$ must distinguish the masking trial from the unmasked, single stimulus case if the situation is not to be exactly the same.[175]

This is precisely to assume Strong PSA. A picture of experience analysable into independent instants *would* commit us to such a view. For then we could not appeal to the nature of subsequent experience (an irreducibly temporal fact) to distinguish the disc case (first trial) from the no disc case (second trial) unless there was sufficient delay in consciousness for unconscious registering of the second stimulus to come into play. However, if we have rejected this picture of experience we have no reason to endorse the disjunction. And so we have no reason to follow the argument which causes problems for the simple view that we do not see the disc. We do not, in particular, need to embroil ourselves in positing delays in consciousness, and unconscious acts which appear to us to be conscious.

---

[175] When I talk of 'present tense facts' I intend to contrast 'irreducibly temporal facts' – facts whose truth logically depends on states obtaining at times other than the present instant. Purely present tense facts merely causally depend (if at all) on states obtaining at times other than the present instant. Cf. Kripke 1978: Lecture V.

8. Apparent Motion

The committed anti-realist may remain unimpressed with this discussion of visual masking. Masking is not, after all, Dennett's central example. I focused on masking because the case is simple and serves to exhibit the key features of the realist defence. I now want to turn briefly to Dennett's chosen case. I show how the defence applies there in equal measure.[176] I then turn to objections to the account set-out. First, a number of complications need addressing.

8.1 The Varieties of Apparent Motion

Dennett introduces his central case, calling it 'phi' (φ) as follows:

> if two or more small spots separated by as much as 4 degrees of visual angle are briefly lit in rapid succession, a single spot will move back and forth (1991a: 114)

In fact, when one turns 'on' two white discs against a dark background in succession, a number of distinct phenomena can be observed. When the succession is relatively slow (above 200ms), there is no illusion. *Ceteris paribus*, one sees (veridically) two stationary discs turning on and off in succession. At the other extreme (less than 30ms) one sees what appear to be two simultaneously presented stationary discs. However, in between three possibilities need to be distinguished (see table below). In sum we have the following ordering: succession – partial apparent movement – optimal apparent movement (β) – pure apparent motion (φ) – simultaneity.[177]

---

[176] The considerations of this chapter also apply to a potentially large number of other cases – some mentioned by Dennett, for example the cutaneous 'rabbit' (Geldard and Sherrick 1972), some not, for example, repetition blindness (Kanwisher 1987). See also the brief discussion of Sperling's visual display experiments above.

[177] For an online demonstration see http://www2.psych.purdue.edu/Magniphi/.

Type of apparent motion	Inter-Stimulus Interval	Phenomenology
None	> c.200msec	Succession
Partial or dual apparent motion	c.60msec – 200msec	Two white discs, each apparently moving a short distance from its actual location.
Optimal apparent motion, Wertheimer's β	c.50 – 60msec	A single white disk moving back and forth between the actual locations of the two disc.
Pure apparent motion, Wertheimer's φ	c.30 – 50 msec	Two stationary flickering discs with a contourless ('shadow', 'cloud' or 'flag-like') region with the colour of the background moving between and slightly around them in counterphase to the turning on/off of the discs.
None	< c.30msec	Apparent simultaneity

What Wertheimer called 'φ' in 1912 is a striking phenomenon. As Kai von Fieandt's puts it, one observes

> [A] peculiar phenomenal motion … an objectless movement, or 'pure motion' as Wertheimer described it. Without seeing any moving objects of figures, there [is] a clear impression of motion from one place to the other. (1966: 263; quoted in Steinmann et al. 2000: 2259, fn.3)

However, this is not a description of Dennett's chosen case. In passing, Dennett suggests that "there is an intermediate range of intervals where the phenomenology is somewhat paradoxical: you see the spots as two stationary flashers and as one thing moving!" (1991a: 123). This is closer to φ. However, Dennett suggests this occurs at intervals between succession and optimal apparent movement and, moreover, that the phenomenology is of a spot of the same kind as the stationary spots moving. None of this fits with φ.[178]

[178] For a fascinating discussion of these issues and the sources of this long-standing confusion in the literature see Steinmann et al. 2000.

In any case, it is quite clear that Dennett's interest is in β and that shall be my concern for the rest of the discussion. However, the distinction is important. In particular, there may well be issues concerning how one should make sense of the notion of 'objectless movement' that require consideration in the literature.

8.2 Dennett's Chosen Case

Dennett considers the following case of β-apparent motion. Two discs, the first red, the second green are turned on for 150ms each in succession separated by a 50ms interval (i.e., SOA = 200ms), as below.[179]

t = 0 – 150ms

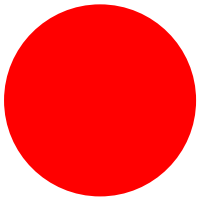

t = 200 – 350ms

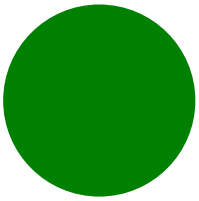

The results of such experiments, according to Dennett, are as follows:

> The first spot seem[s] to begin moving and then change colour abruptly in the middle of its illusory passage toward the second location (1991a: 114)

Assuming that subjects are not primed, and each disc is only turned on once (repeated cycles allow for explanations in terms of expectation), the reports of subjects raise a

[179] See Kolers and von Grünau 1976. This paper simply talks of apparent motion (as opposed to β).

problem. Put crudely, 'How does the brain know that the second disc will be green before it is too late?' (i.e., before the subject is already simply perceiving a stationary green disc). The problem is especially vivid in the coloured case, but it is not important that the discs are coloured. If one simply flashes a white disc on a black screen, there will be no perception of movement. However, if one flashes two discs in rapid enough succession, optimal apparent movement occurs. But the same question arises here: 'How does the brain know that the second disc will flash before it is too late?' (i.e., before it is already simply perceiving the second disc stationary, at its actual location).

8.3 Stalin and Orwell Revisited

Dennett now sets out the familiar Orwellian versus Stalinist dialectic. According to the Stalinist we do genuinely enjoy experience as of a disc moving across the screen, changing colour half way across. This is what subject's insist they see, and they should know.

Two familiar considerations then motivate the puzzle. Had there been no second, green disc, the subject would not have experienced any illusory motion at all. Consider then the moment in this single stimulus case when a subject ceases to see the first disc any more, $t = 150\text{ms} + \delta t$. Assuming that there are such things as truths about instants or brief intervals in the stream of consciousness which hold independently of truths about intervals of experience we have two options as to what to say about the analogous moment in the β-case. On the one hand, we can say that conscious experience is delayed by a time sufficient for our brains to have processed information about whether or not there is a second disc before we experience the apparent motion (or lack of it).[180] This is Stalinism. On the other hand, we can deny that we see apparent motion in either case and claim that it is merely a fiction of our memory that we do.[181] This is Orwellianism. In

[180] Dennett assumes that the relevant slack must be at least 200ms (see his 1991a: 121). However, his logic only holds if subjects report the disc as moving from the very beginning. Without this unwarranted assumption the lowest possible slack will be fixed by the inter-stimulus interval not the SOA, i.e., at 50-60ms. That said, dual apparent motion raises similar considerations with an inter-stimulus interval of up to 200ms so Dennett's figure may after all be correct. Of course, precise figures are an empirical matter. (Note that the related cutaneous 'rabbit' illusion occurs with inter-stimulus intervals of up to 200ms.)

[181] "You say and believe that you saw the illusory motion and colour change, but that is really a memory hallucination, not an accurate recollection of your original consciousness" (1991a: 121).

response to this line of thought, exactly parallel points can be made as were made in the masking case.

9. Naturalism

A reader sympathetic to Dennett might well object to the above accounts as follows: Even granting that experiential instants and brief intervals are not independent of each other, your account is incompatible with a plausible naturalistic account of our mental lives. After all, Dennett's argument rests only on denying backwards causation and not on metaphysical claims about the homoeomerous nature of experience or Strong PSA .

Certainly we should agree that experience is a part of the causal order. Experiential time cannot simply "float free" of processing time. Given this, the following constraint on any satisfactory account of masking phenomena may seem plausible.

> **Causal Constraint** For any experiential property whose instantiation by experience is the causal upshot of the visual (or auditory etc.) processing of certain information, experience cannot instantiate that property until the relevant information has begun to be received and processed.[182]

This constraint may seem problematic for the following reason. In the normal non-masking case, when the disc flashes up, the account above eschews (full-blown) Stalinesque delay. Thus, it is committed to claiming that we see the disc *before there is time for a possible ring to block awareness of the disc*. Yet, if this is so, the Causal Constraint appears to force us to attribute this same property to experience at the parallel moment in the masking case just before the mask appears. (Clearly, something we must avoid.) After all, attributing experience of no ring would be to attribute a property dependent on as yet unprocessed information regarding a future mask. Indeed, in both cases there appears to be an aspect of experience – disc absence or presence – which causally depends upon information being received or not in the future. That seems obviously incompatible with Causal Constraint.

---

[182] Cf. the truisms in Dennett 1995.

This objection reveals how hard it can be to appreciate the idea that experience might be anhomoeomerous. We can see why this is so with reference to a non-experiential example. Recall the discussion of walking above. That discussion concluded that whether one counts as walking depends on what occurs over some duration of time. We should therefore reject the following constraint on walking.

> **Walking Constraint** Walking is a part of the causal order. Thus, one cannot instantiate the property "… is walking," at *t* unless sufficient activity has gone on to ground walking at *t* regardless of what may come after *t*.

This is false. One may count as walking at *t* even if only a single support phase has so far occurred, a phase which in and of itself is not sufficient to ground walking (as opposed to running etc.). Another way of putting the point is by saying that one can be in the process of walking without yet having walked. One has not yet walked at *t*, but one can nonetheless be walking. What is true is that one cannot have *walked* unless sufficient activity has gone on to ground walking.

The same, I submit, is true of experience. One can be experiencing, i.e., in the midst of experiencing some event or process, without *yet having experienced* the event or process. In particular, one can be in the midst of experiencing a certain apparent event (viz., disc-followed-by-blank-screen, or blank-screen-followed-by-ring) without yet having experienced the event. We might put it like this: different events have different perceptual appearances for us. The perceptual appearances of events are constrained by the fact that our experiences themselves have a certain coarse-grained temporal structure. These constraints mean that we must think of the appearances of events as attaching first and foremost to durations above a certain minimum temporal threshold, in other words to temporally extended events rather than instantaneous or near instantaneous time-slices. In the masking case we should consider the events, (i) disc followed by mask, and (ii) disc followed by no mask, as different events with different appearances. What masking experiments show is that the appearances of sub-events of an event whole can contrast strikingly with the appearance of one of the sub-events presented on its own.

Thus, as with walking, one may be in the midst of experiencing in some way without yet having experienced in that way and so without yet having processed all the information

necessary to experience in that way.[183] Nonetheless, it is quite right to think that experiential time cannot simply "float free" of processing time for, as with walking, one cannot *have experienced* a particular masking episode unless mask presence (or absence) information has been received and processed. Once we abandon the idea that experience is homoeomerous, we can happily allow that our fundamental explanations (causal and otherwise) about experience will be explanations of temporal periods of a certain size. All we can say of instants/very brief portions of experience is that they are instants/brief periods during which subjects are in the midst of experiencing in some manner (e.g., a blank-screen-followed-by-ring) as determined by which stretch of the stream of consciousness they fall within.

This is not to say that there are not complications here. In particular, the picture here is one on which it is appropriate to attribute a property to a temporal point during some course of experience, where that point occurs before all the processing relevant for the attribution of that property has taken place. That is acceptable only because the truth of the attribution holds only in virtue of the occurrence of a longer duration of experience in which the moment is embedded and the insistence that the Causal Constraint applies only to this longer duration of experience. If, with Frege, we think of utterance or propositional truth as tenseless, no difficulty emerges here, for it will be true even just before the mask has been presented that the subject is in the midst of experiencing a masking event (i.e., blank screen followed by a ring). However, if we are convinced that the future is open in the strongest sense, one which demands tensed propositional truth, we will need to deny that at the moment just before the ring is presented and processed that there is a fact of the matter as to what the subject is experiencing. Relative to that context of assessment it will be indeterminate what the subject is experiencing. Subsequently, relative to a given future context, there will be a fact of the matter. [184]

---

[183] In the apparent motion case it is even more tempting to baulk at the idea that we might be experiencing motion without yet having experienced any yet. But this is simply because over ordinary timescales we treat experiential processes as homoeomerous. At the tiny timescales here involved that assumption fails.

[184] For details of this kind of approach see MacFarlane 2003 and forthcoming. This kind of context-relativity is not a particular issue for my approach to Dennett and it is certainly not the case that the approach here rests on the success of MacFarlane's programme.

10. Conclusions

Dennett tells us that our intuitions suggest "that our streams of consciousness consist of events occurring in sequence, and that at any instant every element in that sequence can be classified as either having occurred 'in consciousness' or as having not occurred 'there' yet" (1991a: 144). There are two immediate problems with this way of framing our intuitions. Firstly, the term 'classified' can be understood in two ways. It might simply refer to there being a fact of the matter as to whether some element is currently being experienced or not. On the other hand, it might refer to a subject's ability to determine whether some element is currently being experienced or not. Secondly, I have urged that the stream of consciousness be thought of as fundamentally processive. Though we can divide it up into events, it should not be thought of as the mereological sum of a sequence of events. In this light, a better statement of our intuitions is that our stream of consciousness is a temporally structured process, and that at any instant (i) there is a matter of fact concerning the nature of the stream at that instant, and (ii) the nature of the stream is available to the subject of the experience at that instant.

Is anything wrong with this conception? My own view is that there is, strictly speaking, nothing wrong with it. (I set aside the issue of the open future.) Rather what masking and apparent motion phenomena teach us is that we need to be extraordinarily careful about putting things this way. For one, what is true of any given course of experience at an instant is not independent of the course experience takes over longer timescales. So, although it is correct to hold that at any instant there is a fact of the matter concerning the stream of consciousness, such a claim is liable to mislead if we do not appreciate the anhomoeomery of the stream. Truths about instants only hold in virtue of truths concerning longer periods. It is even more misleading to talk about what a subject is able to determine at a given instant. A subject's epistemic position with respect to their experience is a state that they are in which constitutively depends on their stream of consciousness (i.e., of the processes and/or events that the stream is comprised of). Thus, the state that one is in at an instant will depend on what is occurring in the stream at that instant. Since that depends on what goes on over some longer period, however, so too the state that one is in at an instant obtains only in virtue of facts about one's mental

life over longer periods of time.[185] In other words, a claim parallel to experiential anhomoeomery applies to epistemic states too.

What, if anything, prevents Dennett from agreeing? One possibility is that the block is his conception of seeming. As we saw above, Tye assumes that we must either insist on the complete independence of consciousness and self-conscious judgement (a form of ultra-realism) or identify the two (anti-realism). This misses the possibility of the middle-position defended in Chapter Three, according to which the nature of conscious experience is not fixed independently of how it is available to self-consciousness, but nonetheless conceives of judgements as grounded (and on occasion ungrounded) in experience as opposed to constituting truths about it. It is possible that Dennett's notorious remark, "There is no such phenomenon as really seeming – over and above the phenomenon of judging in one way or another that something is the case" (1991a: 364), should be taken as revealing that Dennett shares Tye's conception of the landscape. If so, Dennett no doubt finds anti-realism compelling. But Tye's complaint that his argument begs the question against the ultra-realist will also stand.

A more attractive, but much more speculative diagnosis, focuses on the key idea behind the multiple drafts model, which Dennett tells us, "avoids the tempting mistake of supposing that there must be a single narrative … that is the *actual* stream of consciousness of the subject, whether or not the experimenter (or even the subject) can gain access to it" (1991a: 113). As I have said, on the realist account offered above, it makes perfect sense to insist that there is an actual, single, unified stream of consciousness despite acknowledging that this is not structured as it is in independence of the subject's ability to access it. Perhaps what Dennett does not quite see is the way in which a subject's access at an instant can be a matter of their being in an anhomoeomerous epistemic state as just discussed. If Dennett, were to think that there were only punctate judgements, i.e., events at instants, it would indeed be hard to combine an anhomoeomerous stream with a version of Self-Intimation framed in those terms. The nature of the stream at some point would be dependent on its course over some period; but the judgement about that point made at that point would not seem able so to depend. Later judgements would thus indeed be in conflict with earlier judgements

[185] Compare here Soteriou 2007. Thanks to Guy Longworth for pressing me on this point.

and we would find ourselves disposing of the idea of the realist's single narrative and thinking in terms of multiple drafts.

In any case, given the availability of a subtle form of realism, which emphasises the anhomoeomery of experience, I see no reason to adopt such a position. The experiments Dennett employs to attack our intuitive picture of conscious experience do indeed turn out to have serious implications for the way in which we think about temporal experience. If I am right, they show that the stream of consciousness is homoeomerous only down to timescales of perhaps several hundred milliseconds. This should not be such a surprising empirical discovery about how experience is for us – anhomoeomery is a common phenomenon amongst processes in general. Appreciating it in this context shows that empirical findings do not refute realism. In the next two chapters, I show how appreciating it dissolves two other puzzles about the experience of time. By the end of those chapters, a fully satisfying picture of temporal experience will have emerged.

Chapter Six:
Perceiving Temporal Properties

We normally take experience in larger chunks, and if we try to pulverize it by focusing attention on particles within … we usually find ourselves puzzled and uncertain.[186]

1. Puzzlement

This thesis began with a datum, the fact that myriad different temporal properties and relations can be made manifest in perceptual awareness. Consider, for example, the celebrated opening bars of Gershwin's *Rhapsody in Blue* during which the first B-flat clarinet, starting from a long low trill, crescendos flamboyantly through a smooth two-and-a-half octave glissando to arrive on a sustained minim concert B-flat (see below).

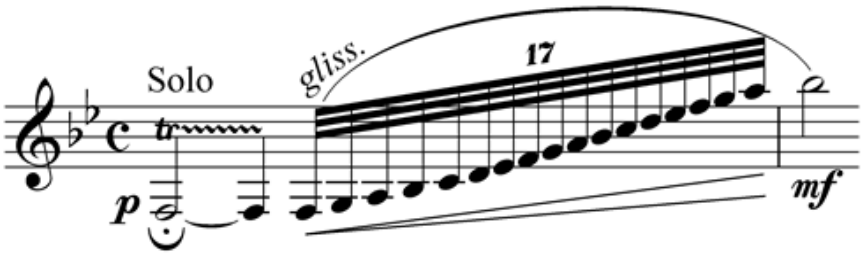

Listening to this passage, our experience seems to present us with various aspects: (a) the alternation of successive notes at the beginning of the phrase, first *piano* and steady, then louder and faster as the trill accelerates; (b) the rapid, long, continuous seventeen note ascent that emerges from it; and (c) the final sustained minim, lingering for a moment before relaxing into the next passage. In listening to the clarinet, we experience persistence, succession, and, with the final B-flat, simultaneity – assuming that the rest of

[186] Goodman 1977: 203.

the orchestra is brought in on time. We also experience properties logically connected to time such as change in volume, pitch, speed, tone etc.

Evidently, it is not just our enjoyment of music that depends upon the perception of such temporal properties and relations. The world buzzes and blooms around us and we are constantly experiencing the movement, change, persistence, succession and simultaneity of its denizen objects and events. Moreover, as noted in Chapter One, we seem to be no less directly acquainted with the temporal structure of the world around us than with its spatial structure. As John Foster puts it,

> duration and change through time seem to be presented to us with the same phenomenal immediacy as homogeneity and variation of colour through space. (1982: 255)

Barry Dainton agrees, calling the direct experience of change and persistence an "obvious truth" and terming it "the *phenomenological constraint*" (2000: 114–5).

Obvious as these facts are, they have been the source of profound and long-standing philosophical puzzlement. Indeed, Kant – arguably the greatest influence on thinking about issues relating time and experience – ultimately denies the "obvious truth" that we directly experience change and persistence. In particular, throughout the *Analytic of Principles*, he asserts that "time cannot be perceived by itself, and what precedes and what follows cannot, therefore, by relation to it, be empirically determined in the object" (2003: B233).[187] According to Guyer, Kant's claim here is "more intelligibly" rendered as the view "that particular temporal relations are not directly perceived" (1987: 167).

Kant is far from alone in this regard. The key father figure in Anglophone discussion of time-consciousness is perhaps Thomas Reid who writes as follows.

> It may here be observed, that if we speak strictly and philosophically, no kind of succession can be an object of either the sense or of consciousness; because the

[187] See further B219, B225, B257 and A183, references in Guyer 1987: 88 and also A33/B49–50.

> operations of both are confined to the present point of time and there can be no succession in a point of time. (1827: 169)[188]

Another 'denier' is Prichard (1950a) who claims that it is, strictly speaking, impossible to hear sounds. What could drive someone to this extreme view? More generally, why have philosophers struggled to make sense of our experience of succession, persistence and change?

In the first part of this chapter I suggest a principle that lies at the root of the puzzlement. I then show how the two leading, rival theories of temporal awareness – specious present theories and memory theories – should be seen as attempts to explain how temporal awareness is possible in the light of that principle. In the second and main part of the chapter I argue that, for the same fundamental reason, neither of these theories is satisfactory unless it rejects the very principle that drove its motivating puzzlement. Thus, the developments of both theories represent paths which take us back to where we started. However, what stands revealed at the end of our exploring is that the apparent rivalry between memory and specious present theories is illusory. Both theories must ultimately unite in rejecting the background assumption that forces them apart and makes them unworkable. Once this background assumption is rejected, the theories no longer appear to be rivals. What is more, we must reconsider what, if anything, is wrong with a very simple account of temporal awareness.

Put crudely, the problematic principle is that if one experiences succession or temporal structure at all, then one experiences it at a moment. Applying the claims defended in Part One, I demonstrate that this principle is false; in order to understand the perception of temporal properties we must look beyond the instant. Our experience of temporal phenomena cannot be understood if we attempt to break experience down into instantaneous slices. The attempt to do so drives the philosophical puzzlement. Of course this problematic principle is precisely the assumption already fingered above in

---

[188] Reid notes that "this observation seems to contradict the common sense and common language of mankind" (ibid.) but goes on to account for this by (a) providing a now traditional memory theory of temporal consciousness (see below) and (b) noting that the 'vulgar view' is explained by a failure to properly distinguish memory and perception. For discussion of Reid and his followers such as Dugald Stewart, Thomas Brown and William Hamilton see Anderson and Grush forthcoming. Note that a common objection to memory theories is that they are really denials that we actually perceive temporal properties. Reid simply admits this; cf. discussion of Brentano below.

resisting Dennettian anti-realist arguments, viz., the homoeomerous nature of experience. Puzzlement and denial is driven by the assumption that experience must be homoeomerous down to instants. Contraposing, I argue that experience can only be experience of temporal properties if it is anhomoeomerous.

2. Experiencing Succession: a Simple Example

It will help to have an example in front of us. Consider the following very simple case of auditory experience: hearing a *C* major, broken triad played *staccato* and *allegro* on a well-damped piano. Listening as the piano is played, one experiences each of the notes of the broken chord in turn. But one's experience also has an additional aspect which can be brought out by comparing *Case A* with *Case B* where one simply hears a *staccato G* played on the same piano.

Case A

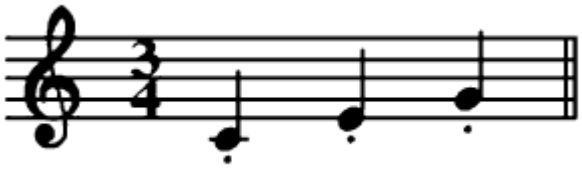

Case B

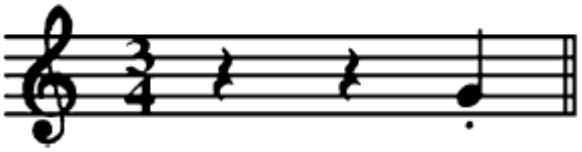

Cases A and *B* resemble one another in this respect: in both, a small time after the *G*-key has been struck, one hears a *G*. However, intuitively, there is also a difference in one's experience at this time, *t*. In *Case A*, one does not merely hear a *G*, one experiences *succession*. That is, one hears the *G following on* from the two previous notes of the triad and not in lonely isolation as in *Case B*. One would be failing to characterise fully how things were for you at *t* if one only mentioned the fact that one was then hearing a *G*.

At *t*, in *Case A*, one is in a position to attend to a series of note-soundings spanning an interval of time. This contrasts with a case in which the notes are spread out over a very long period of time. No doubt in such a case one might recall the past notes as one heard the *G*, but – assuming that the period is long enough – one would not be able to selectively attend to all the notes occupying that interval. This, I suggest, supports the thought that, in the slow case, one's experience of the auditory world at *t* would not require mention of successiveness (even if one's overall experience in some way did).[189]

More generally, the datum is this: there are cases (like *Case A*) in which one hears or perceives in such a way that one is able to attend to a structure of sounds, events or event parts which occupy a temporal interval. If what has just been said is right, we can ask: how should we account for such cases of perception, in particular for the additional aspect of experience in *Case A*? A simple-minded account might run as follows.

> The difference between the two experiential situations (*A* and *B*) at time *t* when the *G* is first heard is simply that in situation *A* the subject has just heard a *C* and an *E* in that order, whereas in situation *B* the subject has just heard nothing, merely two beats of silence before the *G*.

According to the simple-minded account, an irreducibly temporal difference is appealed to as a way of distinguishing the two experiences, namely, the past experience of the subject.

This kind of appeal is rejected by almost all participants in the debate. William James articulates the objection thus,

> *A succession of feelings, in and of itself, is not a feeling of succession.* (1890: 629)

Husserl makes the same point.

---

[189] Bill Brewer encouraged me to clarify the sense in which we have genuinely direct perceptual experience of temporal properties as opposed to a more general sense of the past's relation to present experience. His own example was of hearing a long, difficult string quartet by Morton Feldman – *String Quartet II* at over five hours being an extreme case in point – where notes heard an hour ago might still, in some sense, affect one's current experience – giving one a sense of finality, repetition or larger structure. Such cases are not the source of philosophical puzzlement at issue here.

> The duration of sensation and the sensation of duration are different. And it is the same with [succession].[190] The succession of sensations and the sensations of succession are not the same. (Husserl 1964: 31)

Careful to distinguish between acts and objects of acts, he goes on to make the equivalent point about acts.

> We must naturally raise precisely the same objection against those who would trace the idea of duration and succession back to the fact of the duration and succession of the psychical act. (ibid.)

James' and Husserl's thought is that merely having a series of experiences with differing objects, or one extended experience whose objects change over time, is compatible with the absence of temporal experience. This is, of course, true in many cases. When the triad is played very slowly, merely experiencing each note is not sufficient for experiencing them as successive.

However, James and Husserl seem to suggest that this claim holds across the board, for all experience. Their accounts suggest a picture according to which all phases of experience are strictly independent of preceding experience-phases. What we need, therefore, is some way in which our *current* experience can embrace temporal structure. As James puts it, the "feeling of past time" must be "a present feeling".[191] The thought is that at any time, the nature of one's perceptual experience must be based on the *single state available at that time*. Insofar as one is immediately aware (or apparently aware)[192] of succession or temporal structure, this must be due to that whole structure being represented or presented to you *at that very point in time*.[193]

---

[190] The translation has 'sensation' here but this seems to be a misprint.
[191] James 1890: 628; cf. Husserl 1964: 40 and Miller 1984: 108f. Husserl's lectures date from 1905.
[192] Throughout what is in question is our apparent awareness of succession and so the apparent objects of awareness. Nothing is meant to turn on awareness or perception being a success.
[193] This claim traces back to Kant. Kant makes it most clearly in the A-deduction. Thus at A99 he writes that a "representation, in so far as it is contained in a single moment, can never be anything but absolute unity". Guyer holds that this is "the fundamental premise of Kant's transcendental theory of experience" (1987: 171) and comments, "[What this claim] implies is precisely that although, of course, the manifold of subjective states occurs or is given successively, knowledge at any particular time that any particular succession of such states has occurred must be based on the single representational state available at that time. And this means that an interpretation of that state is necessary for the mind to determine the sequence of one impression upon another (as Kant puts it). In other words, the several members of a

Another less well-known exponent of exactly this principle is T.H. Green who, in the words of G.F. Stout, is one of a number of writers in whom we "find … a continual reiteration of the statement that the apprehension of succession cannot be itself succession – that in order to be aware of B as succeeding A we must have both A and B before consciousness at once" (1900a: 1). With his usual acumen Stout continues, "The necessity does indeed appear self-evident. But it is worth while to consider what is really involved in it, and in what way the actual process of consciousness satisfies this requirement which is imposed upon it *a priori*" (ibid.). I touch on Stout's own view below. But first I want to consider what makes this supposed *a priori* constraint so compelling for the thinkers mentioned above.

3. The Principle of Simultaneous Awareness

Following Miller (1984), let us call the idea given loose expression above – the claim that if one experiences succession or temporal structure at all, then one experiences it at a moment – the *Principle of Simultaneous Awareness* (PSA). Why do Husserl, James and others embrace this principle? I suggest two possible and related reasons.[194]

3.1 The Conceivability of Unawareness

Our question, recall, is whether there have to be further present tense facts which hold at *t* to ground the difference between *Cases A* and *B*, facts over and above the irreducibly temporal facts that clearly do distinguish the two cases. When I talk of 'irreducibly temporal facts' I mean facts the truth of which logically depends on states obtaining at times other than the present instant. Purely present tense facts merely causally depend (if

---

succession of states are indeed immediately perceived in succession, but there is nothing which counts as immediate perception of the succession" (1987: 171–2).

[194] Kelly opens his discussion of what he calls the puzzle of experience by posing it in the following terms: "How is it possible for us to have experiences as of continuous, dynamic, temporally structured, unified events given that we start with (what at least seems to be) a sequence of independent and static snapshots of the world at a time?" (2005: 210) It is quite perplexing why this would *seem* to be our *starting* point rather than something we are driven to by philosophical arguments such as those below.

at all) on states obtaining at times other than the present instant.[195] One line of thought motivating subscription to PSA is the following. There *must* be such facts, for there is a possible *Case C* which resembles *Case A* except that we have no experience of succession in *C*. Husserl is explicitly motivated by this thought, insisting that "it is conceivable that our sensations could endure or succeed one another without our being aware of it in the least" (1964: 31-2).

Case C

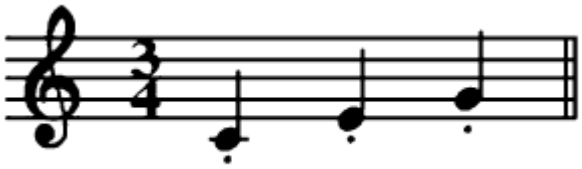

Case C, then, is like *Case A* – the same notes are sounded and heard at the same tempo and in the same order. However, so the line of thought goes, it is conceivable (and so possible) that one might experience each of the notes individually in *Case C* and yet fail to experience succession as an aspect of one's experience on hearing the final *G*. More generally, it is conceivable that one might fail to hear *any* relations of succession between the sounded notes over and above the individual sounds. Consequently, one would not be in a position at *t* to attend to the whole broken chord, spread out as it is, over a temporal interval. If, indeed, *Case C* is possible, it follows that a *merely* temporal difference between *A* and *B* cannot ground the difference in experience.

Even if this strategy is persuasive, however, it is crucial to note that the possibility of *Case C* does not establish that present tense facts (with respect to the time of hearing the final *G*) are *alone* sufficient for experience of succession. It may be that the obtaining of irreducibly temporal facts remains a necessary, though insufficient, condition. Thus, two strengths of PSA can be distinguished.

---

[195] Cf. Kripke 1978: Lecture V. Here logical dependence is not defined in terms of deducibility in any system of formal logic but is rather what is often termed "broad logical" possibility. On this point see Chalmers 1996: 35 which I follow.

Weak PSA Irreducibly temporal facts are insufficient to explain the difference between *Cases A* and *B*. There must, in addition, be present tense facts which hold at *t* to explain the difference. It is these facts which do not obtain in *Case C*.

Strong PSA Irreducibly temporal facts have no place in our explanation of the *A*/*B* contrast. Present tense psychological facts are necessary and sufficient to explain the contrast.

So far as I know, this distinction is not made in the literature. However, the authors mentioned above appear to assume the stronger principle. That leads one to suspect that they have additional motivations lurking in the background.

3.2 Russell Worlds

Another motivation, this time for Strong PSA, may be a thesis given vivid expression by Russell in *The Analysis of Mind*.

> There is no logical impossibility in the hypothesis that the world sprang into being five minutes ago, exactly as it then was, … There is no logically necessary connection between events at different times; … Hence the occurrences which are *called* knowledge of the past are logically independent of the past; they are wholly analysable into present contents, which might, theoretically, be just what they are even if no past had existed. (1921: Lecture IX)

Taken to its limit, this view holds that all presents facts, and hence facts about our conscious or mental lives, are compatible with the world's having been brought into sudden existence any finite time before the present moment and likewise being annihilated any finite time after the present moment. I focus on the past in what follows. If our mental lives are (logically speaking) wholly independent of our histories, then appeal to irreducibly temporal facts is evidently ruled out. Purely present tense facts must be sufficient to account for contrasts like that between *Cases A* and *B*.[196]

---

[196] Note that this thesis is *not* obviously entailed by presentism assuming that the presentist can appeal to irreducibly tensed relations. Nonetheless, it has no doubt often been assumed that such a thesis is a

In what follows, I refer to worlds newly created with the aim of perfectly matching some ordinary world at some moment as 'Russell worlds'. Whether Russell would have been happy taking his argument to the limit is far from clear. The data he proposes as the foundation for his construction in *Our Knowledge of the External World* are explicitly events with a finite duration (1914: 126; cf. 1915: 271 on the absurdity of instantaneous experience).[197] Nonetheless, Russell's basic thought concerning logical independence taken in its strongest form is a perennial one and it is not difficult to find adherents throughout the ages.

Hume is probably the best known proponent of the view that every moment of time is independent of every other. However, Hume's most likely source for the principle is Descartes. Such a principle plays a key role for Descartes in the *Meditations* where he relies on the following claim.

> A lifetime can be divided into countless parts, each completely independent of the others, so that it does not follow from the fact that I existed a little while ago that I must exist now … (1986: 33)[198]

And in one of his replies to Caterus he argues as follows.

> The separate divisions of time do not depend on each other; hence the fact that the body in question is supposed to have existed up till now 'from itself', that is, without a cause, is not sufficient to make it continue to exist in future, unless there is some power in it that as it were recreates it continuously. (1986: 88)

---

consequence of presentism. For example, Dainton (2001: 107-8) effectively makes this assumption in arguing that the presentist cannot account for our experience of succession.

[197] Interestingly, when Carnap (1967: §67) discusses the choice of basic element to be used in his logical construction of the world he chooses "experiences is their totality and undivided unity" (108) in "opposition to the 'atomising' school of thought in psychology and epistemology". He cites various writers on this point (e.g. Schlick, Schuppe and Cornelius) and notes in particular the work of Gestalt theorists. He does not mention temporal duration but there is every reason to think that the anti-atomist rationale ought to apply here too – thus elementary experiences for Carnap ought to be extended in time.

[198] Cf. Ismael nearly four-hundred years later: "We suppose that the mind decomposes into a set of momentary parts … the momentary parts of an extended self retain their psychological identity, they are psychologically separable from parts that lie along the same connected stream" (ms: 3-4).

Descartes does not argue for the claim here relied upon. He simply states that it should be "quite clear to anyone who attentively considers the nature of time" and declares it to be "one of the things that are evident by the natural light" (ibid.: 33).

Lennon in a wide-ranging discussion partly focused on what he calls the "staccato conception of time" (1995: 346) notes that discussion of these issues is "continuous and, at least implicitly, universal through the history of philosophy" (ibid.: 353). By way of supporting this contention Lennon draws our attention to the rich debate in eleventh and twelfth century Islamic philosophy. Of particular interest in the present context is the debate between Ghazali and Averroes as to whether there are any necessary connections between moments of time – Ghazali anticipates Descartes and Hume; Averroes anticipates the contrary view of Leibniz.[199]

More recently, we might also note Bergson's attack on what he calls the cinematographic conception of time (Bergson 1914), a metaphor for the very same idea that time can be analysed into logically independent moments. And similarly, Kripke (1978: *passim*) who argues against what he calls the "holographic or time instantaneous description picture" which he finds implicit in much contemporary (neo-Humean) metaphysics.

Russell's defence of the possibility of Russell worlds, relies on a brute conceivability claim. However, one might try and motivate that claim by appeal to further principles that have seemed independently attractive. For example, one might claim that the physical facts at a time were path-independent, that is held independently of physical facts at other times. If that were right, Russell's thought experiment would be consistent with respect to the physical world; there would be no logical impossibility in God creating the world in the state it is in at any moment, yet lacking any history. One might then add a supervenience claim to the effect that the mental facts at a time supervene on the physical facts at a time. This would then commit one to the view that all facts about our conscious or mental lives are compatible with the world's having been brought into sudden existence *any* finite time before the present moment.

---

[199] On Ghazali and Averroes see Lennon 1995: 355 who cites Van Den Bergh Vol. I p.185 and pp.318-9 respectively. On Leibniz's rejection of the Cartesian version of the principle see Lennon 1995: 351 who points us to Leibniz 1965: 30. In conversation, Christopher Martin detailed a history of not obviously successful attempts to trace the independence principle back to Aristotle.

I do not want to discuss these arguments here. All I want to do is to exhibit the pervasiveness of a certain way of thinking about time, however it is motivated. My main aim here is to consider why one might be *concerned* at being forced to accept Strong PSA and so why one might be inclined to resist any arguments that might be mustered in its defence.[200]

Accepting Strong PSA is tantamount to insisting that experience can be broken down into independent slices, each constitutively independent of any others. Thus, if Strong PSA is true, then experience must be homoeomerous down to instants. Conversely, if experience is not homoeomerous down to instants as I have suggested, we cannot accept Strong PSA. After all, if experience is anhomoeomerous, then it has temporal parts whose natures depend constitutively on the nature of experience over periods beyond the periods of the parts. In other words, it has parts whose nature can only be explained by appeal to irreducibly temporal facts. This is in direct conflict with the strictures of Strong PSA. Furthermore, if Strong PSA is false, then irreducibly temporal facts are required to explain certain aspects of experience. That suffices for anhomoeomery. Consequently, experience is homoeomerous just if Strong PSA holds.

Broadly speaking, two views of temporal experience have been suggested which attempt to account for temporal awareness within the constraints of Strong PSA. The first kind appeals to the specious present, the second to memory. I consider each account in turn. My conclusion is that for the same fundamental reason, neither succeeds in accounting for temporal awareness whilst respecting Strong PSA. As a result, we face a choice: reject Strong PSA or deny that we do in fact perceive temporal properties. Or, what is just the same point, reject homoeomery or deny that we do in fact perceive temporal properties.

---

[200] One obvious tension I will not explore here is between the staccato view of time and Naïve Realism. Naïve Realism holds that the phenomenal character of our mental lives is (at least partly) constituted by worldly objects and their properties. If successiveness is one such property which is actualised in our experience, then no such experience could be had in the absence of such successiveness. Thus, the Naïve Realist will need to deny that Russell's thought experiment conceives a genuine possibility. Indeed, given time-lag considerations, the incompatibility of Russell worlds and Naïve Realism goes beyond the presentation of temporal properties in experience.

4. Specious Present Theory

4.1 SPT and PSA

According to the specious present theory (SPT), at any instant we are aware of an extended period of time.[201] Thus, our experience at a moment literally embraces extended temporal structure. SPT provides a very clear account of how temporal experience is possible despite the constraints of Strong PSA.

However, various standard objections have been raised against this kind of theory. Recently, for example, Sean Kelly argues that our being aware of a duration at a moment "simply makes no sense" being "committed to claims about experience that have no sensible interpretation" (2005: 211). Kelly makes a very strong claim here: SPT is an *incoherent* response to the puzzle of temporal experience. Can that be right?

Two of Kelly's three reasons for thinking that SPT is incoherent rest on confusions. Firstly, Kelly thinks that SPT is committed to our being aware of the future and that awareness of the future is impossible. However, such a commitment (coherent or not) is not an essential part of SPT. Indeed, in the course of advocating his version of the theory, Broad contends that "to sense what has not yet become, would be literally to sense *nothing*" (1923: 358); likewise Russell writes, "It should be noted that there is no experience of the future. … What I mean is that there is no experience of anything *as* future" (1915: 227).

Secondly, Kelly suggests that SPT runs into difficulty in claiming that we are aware of the past. However, arguably, we are often aware of events that are no longer taking place because of time lag considerations. Few, for example, would deny that we see supernovae despite the fact that these occur long before the time of perception. Kelly seems to think that time lag considerations are out of place here, since, supposedly, such an appeal would involve abandoning the claim that we are directly aware of the present.

---

[201] See, in particular, James 1890, Broad 1923 and more recently Tye 2003. As with so many *loci classici*, James' account of the theory is problematic in many ways. Notoriously one is led to question how James could have thought the specious present quite so long as he did, arguing for example, that since "our maximum distinct *intuition* of duration hardly covers more than a dozen seconds (while our maximum vague intuition is probably not more than that of a minute or so), we must suppose that *this amount of duration is pictured fairly steadily in each passing instant of consciousness*" (1890: 630).

However, even if this were true, it is not an essential part of SPT that I am directly aware of the present (*contra* Kelly 2005: 219; again cf. Broad 1923: 358). The core claim of SPT is simply that at a moment we can be aware of an extended period of time.

Kelly's third objection is much more serious, however. He challenges SPT to make sense of the idea that a *momentary* experience might simultaneously present a number of successive states of affairs as successive. The specious present must make sense of this if it is to account for our awareness of succession and proponents of SPT certainly do make such claims. For example, Russell contends that "[s]uccession can occur within the specious present, of which we can distinguish some parts as earlier and others as later" (Russell 1921: 145; quoted in Kelly 2005: 220).[202]

Kelly's worry about this seems to be roughly the following: what I experience at a moment cannot be experienced *as* successive (as Russell claims). Experience of succession necessarily takes time. Rather, the most SPT can claim is that successive objects are experienced together. But this then faces the objection that our experience would be like a chord or cacophony rather than a genuine experience of succession.

Although I think this worry is ultimately fatal, Kelly's argument is too quick. More needs saying in order to explain what is wrong with the possibility of experiencing a genuine succession at a moment. For, if we distinguish between the temporal structure of experience and the temporal structure of the objects experienced, it is not obvious why the two should not come apart. That is, we need to ask what exactly is wrong with the following picture where an experience, e, at some one moment is an experience of events taking place at two separated times t_1 and t_2 as such (see below)? On the one hand, there seems nothing *theoretically* incoherent about this picture. On the other, it is hard not to feel the force of Kelly's worry.

---

[202] Likewise: "When I see a rapid movement, I am aware that one part of the movement is earlier than another, in spite of the whole being comprised within one specious present. If I were not aware of this, I should not know whether the movement had been from A to B or from B to A, or even that changes had occurred" (Russell 1948: 226). Dobbs and Broad (1951: 124-5) are especially emphatic about this aspect of the theory.

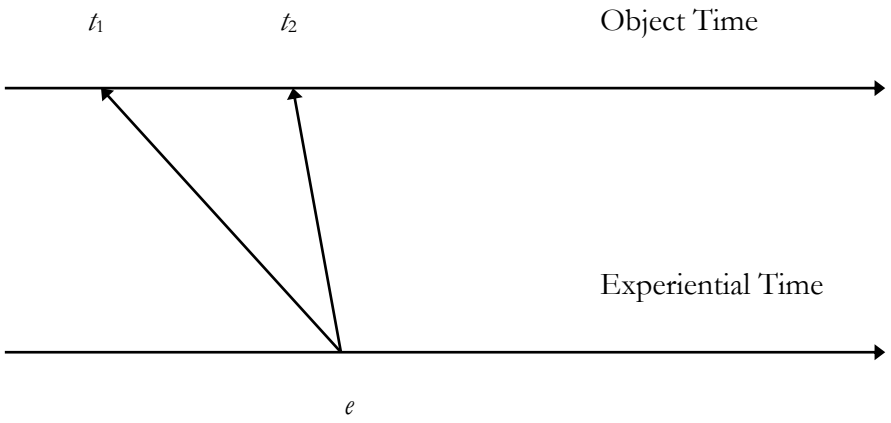

The argument of Part One shows what is wrong with this picture and provides the necessary additional considerations to drive Kelly's worry home. Thus, we can argue against the idea that we can experience succession (or any temporal interval) at a moment not on the basis that it is a theoretical impossibility but rather on the grounds that it is revealed to be impossible *when we reflect on the nature of our experience*. In particular we can deploy the following argument.

First it is noted that experience's temporal structure is experiential structure. We are in a position to know about it just in virtue of undergoing such experience. However, we then note the claim I labelled Temporal Transparency, viz., the idea that for any temporal determination of experience, I am only ever rationally positioned to judge my experience to be so determined on the basis of a judgement concerning the temporal structure of the apparent objects of perception. In other words, I am always in a position to know about my experience's temporal structure, but only by taking it to mirror the structure of its apparent objects. Consequently, our only rational strategy when undergoing experience *of* succession is to take our experience *itself* to be successive in temporal structure as opposed to instantaneous.

Now, specious present theorists might at this point complain that such a judgement, however rational, merely represents how things seem to subjects. However, this leaves them unable to explain how subjects are in a position to know about the temporal structure of their experience, something I have argued is essential to our conception of an experiential property. It also brings them into conflict with the principle I labelled Seems → Is above. According to that principle, if our experience has a phenomenal

property of a certain kind, then that experience cannot *systematically* seem to *differ* phenomenally in that regard. In particular, we cannot make sense of the idea that experience systematically seems to one's rational introspective reflection to possess a certain temporal structure, when it is not in fact genuinely so structured. Assuming Seems → Is and Temporal Transparency we must conclude that *contra* SPT, we cannot be systematically in error when we judge our experiences of succession to themselves be successive in temporal structure as opposed to instantaneous.

The two key principles of this argument have already been discussed and defended at length. The above argument shows how their joint application is inconsistent with traditional SPT. In other words, traditional SPT involves a rejection of the naïve and natural picture of experience in time developed in Part One. According to that picture we cannot be systematically in error when we judge our experiences of succession or duration to themselves be successive or possess duration as opposed to being instantaneous. That is a conclusion which is incompatible with SPT.[203]

We can approach the difficulty with SPT in a somewhat different way by considering what one might call coherence constraints on possible experiential lives. Imagine, for reductio, a person in a new Russell world – a world newly created with the aim of perfectly matching some ordinary world at some moment – who at the instant of creation enjoys experience as of a succession (*G* following on from *E* and *C*, say). Now imagine this person continues to live over time. If creation-instant, Russell world experience of succession is possible, and each instant of experience is logically independent of any other, we should be able to imagine that at a subsequent very close instant (indeed, at any instant no matter how close) the subject could undergo experience as of a note, say, a *G* again, but this time preceded by two beats of silence. Moreover, from a theoretical point of view, there is no reason the Specious Present Theorist has to offer against this being compatible with the subject veridically remembering his previous experience of apparent succession at this subsequent moment.[204] It is hard to see why this scenario should not be possible according to SPT. After all, according to Strong PSA

[203] Note that even severely weakening our epistemic standing with respect to our inner lives from that expressed by Self-Intimation and Seems → Is will not help SPT since the theory must posit a completely systematic illusion about the temporal structure of our experience.
[204] Note that given the timescales and the plausible claim that states of knowledge or epistemic standing cannot be merely instantaneous, a commitment to Self-Intimation (or anything like it) will require some commitment to perceptual experiences being briefly memorable. But this is being ruled out here.

only present tense facts are relevant to what we experience at some moment; thus separate instants of experience are independent.

Nonetheless, when we think about it, we don't seem equipped to make sense of such a case. How could one experience no succession whilst simultaneously remembering (quite rationally and correctly) that one had just experienced apparent succession? In other words, how could things seem both one way and also another incompatible way to you? If we reject the possibility sketched, we effectively impose coherence constraints on the kinds of possible experiences and memories one can combine. Yet it is not clear how we could justify the imposition of such constraints on experiential lives if one accepts Russell worlds as genuine possibilities. In other words, accepting Strong PSA commits us either to the possibility of seemingly incoherent experiential lives or to *ad hoc* constraints on the kinds of worlds which are possible.

Many objections have been raised against SPT. However, I want to suggest that the insights contained in many of the traditional objections to SPT are subsumed by the above argument. The above argument displays the *core* failing of SPT. Let me give two examples.

Kelly's objection concerning SPT's capacity to account for experiences of succession is a traditional one. Broad attempts to handle it by arguing that, within the specious present, successive events are distinguished in terms of their degree of what he calls 'presentedness'. The problem with this move is that it is wholly unclear what this mysterious quality is (unless it is simply assimilated, implausibly, to vivacity). Thus, as Dummett puts it, in relation to motion perception,

> An attempt used to be made to account for [motion perception] within the specious present theory by saying that when, at some given instant, we had a visual impression of the object as at the position where it was at the time when light from it struck our eyes a short time before the present but within the specious present, we were aware of that visual impression *as past*, but nevertheless aware of it *after the mode of the present*. But what does that mean? It does not mean anything. It is merely a form of words concocted to conceal the fact that we cannot explain the phenomenon in question. (Dummett ms.: 6)

I wholly agree with Dummett here. Yet the objection in question (and indeed the objection Dummett himself raises in the paper just cited) does not get to the heart of the issue. Indeed, one can easily imagine modern representationalists reviving Broad's basic idea. If we are to move beyond such debates, we need to recognize the basic problem with SPT as traditionally conceived, namely, its wrenching apart of the temporal structure of experience from the temporal structure of the objects experienced.

A second major traditional worry for SPT exhibits the same superficiality. It is articulated by Dainton as follows.[205] If momentary acts of awareness present durations of time, then an event which occurs or a brief sound which is heard during the duration presented by one act will also be the potential object of other experiential acts (see diagram below). But we only hear such sounds or see such events once. Indeed, if we take momentary acts with extended durations as objects seriously it seems we will hear any sound infinitely many times. As Dainton says, '[t]his is ridiculous' (2000: 141).

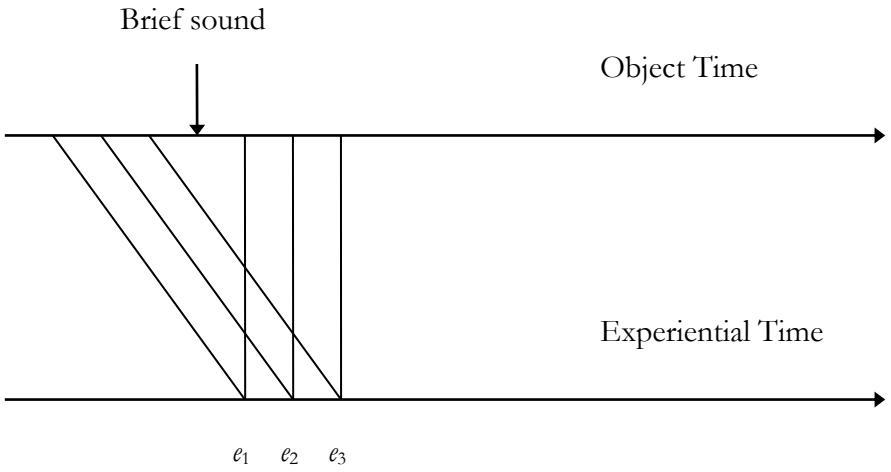

Michael Tye has recently defended a traditional SPT against this objection. He writes as follows. (I have adapted his example to fit the present discussion.)

> This objection is ineffective. Suppose that there are indeed two different token experiences of [a tone], one for each specious present … so that the [tone] is experienced at two different times. Still, it would be a mistake to infer from this that [the

[205] See also Mabbott 1951, Foster 1979 and Sprigge 1993.

> tone] is experienced *as* being at two different times or that I, the subject, have an experience as of two [tones].
>
> I hear [the tone] twice in that there are two times at which an act of hearing a [tone] occurs … But the [two] times have no time between them at which I experience that there is no [tone]. Indeed, there is no time between these two times at which anything is experientially represented by hearing. So, I do not hear a [tone] *as* occurring twice. It does not seem to me that there is a [tone] followed by a second [tone]. (2003: 94; cf. discussion of Tye in Chapter Two)

Tye's response is interesting because it shows how one can respond to the multiple soundings objection if (but only if) one is prepared to give up Temporal Transparency and/or Seems → Is. If one accepts those claims, his reply fails. According to Tye, despite our having two experiences of a tone, it nevertheless seems that there is only a single tone, occurring at a single time. Temporal Transparency tells us that rational judgements concerning the temporal structure of one's experience will match judgements about the temporal structure of the objects of one's experience. Thus, if it does seem as if there is only a single note, we will judge that there is only one time at which we experience a tone. However, Seems → Is then entails that we cannot systematically be misled concerning how our experience is; in other words, we cannot systematically be misguided when we judge that there is only one token experience of said sound. This conflicts with the claim that the specious present theorist is committed to, viz., that there are always multiple such experiences.

We should reject Tye's reply. However, if we do so, we can also leave aside the details of the multiple soundings objection. The objection again merely illustrates the more basic problem with SPT, namely, its attempt to pull apart the temporal structure of experience from the temporal structure of the objects experienced.

Theorists like Dainton and Foster, who are sympathetic to SPT, only avoid the two problems just raised (that of accounting for experienced succession and avoiding multiple soundings) by making two moves. Firstly, they insist that experiential acts are never *momentary* but rather always *extended*. Secondly, they claim that such extended acts overlap, that is, literally share common parts (see Foster 1979: 176; Dainton 2000: Ch. 7). Of course, the overlap claim only makes sense in light of the first claim, that is, if our

experiential primitives are extended acts. Two momentary experiences can only overlap by being strictly simultaneous. The diagram below illustrates the new picture and how it avoids Dainton's infinite soundings objection. Here, the event of the brief sound being heard in the experience e_1 is the *same* event as its being heard in later overlapping experiences such as e_2.

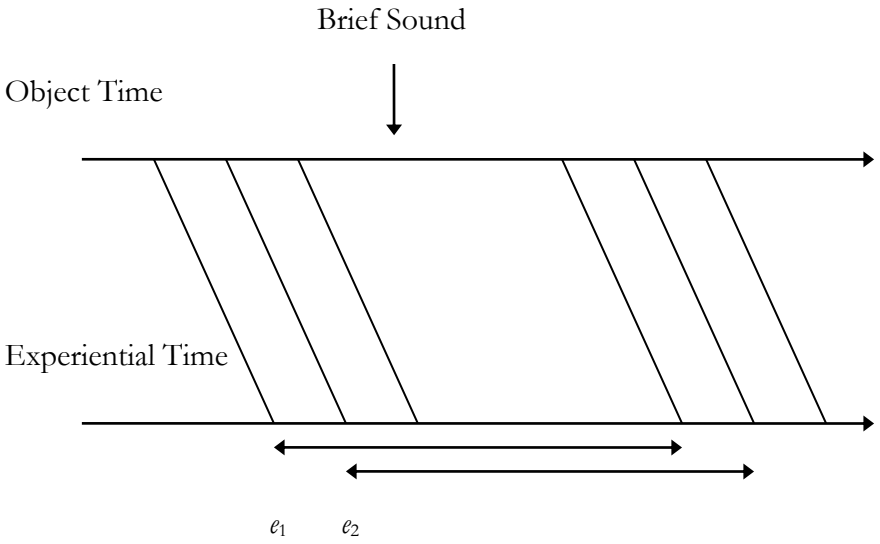

I consider these moves to be the only way to salvage SPT. For they represent the only way to avoid pulling apart the temporal structure of experience and the temporal structure of the objects experienced. The problem with these moves, however, is that the denial that awareness is "packaged into momentary acts" (Dainton 2000: 166; cf. Foster 1991: 249) straightforwardly flouts the thought that each new momentary phase of experience might (conceivably) be completely independent of preceding experience phases, i.e., Strong PSA.

Consider a situation where you are listening to a sound or melody. Take some moment during that period and consider the Russell world for that moment – a world newly created with the aim of perfectly matching the ordinary world at that moment. In order to match the ordinary world, the Russell world must allow that a sound or melody is being heard at the moment of creation. However, that involves the postulation of a non-momentary act of awareness and thus involves a commitment to there being a past or future beyond the instant in the Russell world. This conflicts with the supposed conceivability of qualitative identity of the world at a time without irreducibly temporal

matching. Given this, SPT holds no refuge for a defender of Strong PSA. The principle must be abandoned if we are to account for our perception of temporal structure.[206]

4.2 Foster and Dainton's Overlap Theory

Dainton, following Foster, does not simply reject Strong PSA. Rather he develops a sophisticated reworking of the SPT – the 'overlap theory' – designed to explain temporal awareness. We need to ask though: does this theory really have any work to do once we have rejected Strong PSA?

Recall James' objection to the simple theory, namely, that a succession of experiences does not amount to an experience of succession. Dainton offers the following response to this concern in the light of his overlap theory.

> The difference between an experience of succession and a succession of experiences poses no difficulty at all, for according to the overlap theory every temporally extended experience is an experience of succession. The experience [of successive tones, *Do-Re*] amounts to an experience of succession for two reasons: first because Do is co-conscious with Re (and vice-versa), and second, because the content of this experience is a phenomeno-temporal pattern, of *Do-flowing-into-Re*. There is no need to posit a point like awareness which encompasses both tones. (2000: 180)

Two notions are appealed to here: 'diachronic co-consciousness' and 'phenomeno-temporal patterns'. Let us briefly consider the role each plays in turn. Dainton tends to talk of co-consciousness both as a relation between acts and as a relation between objects. Thus, whereas in the passage quoted he is talking about the co-consciousness of objects, elsewhere he writes, "[t]hree total experiences X, Y and Z can be such that X is co-conscious with Y, and Y with Z, but X is not co-conscious with Z" (2000: 172). This, in turn, is perhaps explained by Dainton's ultimate rejection of an act-object conception of experience.[207] As a result there is a certain obscurity in his discussion. Nevertheless, as

---

[206] None of this is, as yet, to say that Weak PSA may not be well motivated by the conceivability argument. See §9, however.

[207] See Dainton 2000: 166. The basic thought seems to be that "it makes no sense to suppose that an act of awareness can apprehend a content of greater temporal duration than itself" (ibid.: 180) and so there is no motivation for the act-object view. I cannot see why this makes no sense – though that is not to say that it

I read it, the central problem with Dainton's appeal to diachronic co-consciousness as applied to either experiences or objects of experience is that it amounts to nothing more than a blunt denial of the claim that each new momentary phase of experience might be completely independent of preceding experience phases. Indeed, 'co-consciousness' seems no more than a piece of terminology with which to frame the rejection of Strong PSA. As Dainton puts it,

> Co-consciousness is a basic experiential relationship, one about which there is nothing more to be said, at least while we confine ourselves to describing how things seem. (2000: 84)

What about Dainton's notion of phenomeno-temporal patterns? Here, Dainton openly struggles with how to account for genuine experience of passage even given the overlap model. As he acknowledges, when hearing *Do-Re*,

> we experience the notes as occurring in a definite temporal order … hear *Do* giving way to *Re* … hear the first note *flow into* the second note. (2000: 173)

Likewise,

> an individual auditory sensation itself exhibits flow. For the short time it lasts, the tone seems to be extruding itself forward into the future. (ibid.)

Dainton sees trouble because,

> Since the temporal asymmetry is phenomenal, we cannot appeal to memory, and since co-consciousness is symmetrical with respect to time, co-consciousness cannot be the answer. (ibid.)

Now, Dainton is right to think that appeal to a symmetric notion of co-consciousness will not do here. Dainton's solution is to deny that it is the *job* of co-conscious experiences to solve the problem he points to. Rather, it is explained by "experience itself

---

is a genuine possibility. Moreover, it's not being a genuine possibility would not entail Dainton's conclusion. Perhaps, as suggested above, acts and objects necessarily match in temporal structure.

possessing an inherent direction," structure and flow. Co-consciousness of overlapping acts only explains why a "succession of notes is experienced as fully continuous" (176).

In other words, at this point, Dainton directly appeals to irreducibly temporal properties – the flow of experience itself. But, if it is legitimate to appeal to irreducibly temporal properties in the context of explaining our perception of order and flow, why not simply do the same with respect to our perception of temporal properties more generally? Why not simply reject Strong PSA and leave it at that? Once Strong PSA is abandoned, the whole overlap theory looks like unwholesome food served to a man already full.

5. The Appeal to Memory

If one remains convinced of the truth of Strong PSA, one needs to find a different way of making sense of our temporal awareness other than SPT. Let us therefore forget SPT and return to our original scenario of the *C* major triad. Consider the following question: when the final *G* sounds, what happens to our consciousness of the preceding notes? Assuming those notes are not sustained, I no longer hear the *E* or *C* which preceded it. If that is all there is to say about our experience, then it becomes hard to see how we can ever experience the relations two notes bear to each other, and so directly experience succession.

One option would be to say that the preceding notes simply persist in consciousness in the manner in which they were first presented. Of course, the problem here is, as Husserl puts it, that "instead of a melody we should have a chord of simultaneous notes or rather a disharmonious jumble of sounds" (1964: 30). Brentano sees that we can avoid this unfortunate result if we allow the preceding note(s) to remain in consciousness but in a different way, specifically, in memorial consciousness.[208] As we presently experience the present tone, we must, according to Brentano, remain aware of the preceding tone or

[208] I follow Husserl's exposition of Brentano's view which he uses as a stalking horse in the opening sections of *The Phenomenology of Internal Time-Consciousness* (1964). A full treatment of Brentano's view would need to take account of important changes in his position post-dating the theory Husserl sketches.

tones in memory. If we did not, "[i]n each moment we should have only the consciousness of the sensation just produced and nothing further" (Husserl 1964: 32).[209]

Early memory theorists such as Brentano assume that perception and memory are intrinsically the same kind of mental phenomenon. Thus, Brentano holds that "phantasy-presentations [i.e., imaginings and rememberings] … differ from sensations [i.e., perceptions] only in their [causal] origin, not in their [type of] content …" (1973: 316; quoted in Miller 1984: 105). Evidently, with this conception in play, no real progress has been made here in responding to the chord-cum-cacophony objection. Fortunately, better accounts of the relation between memory and perception are available. One attractive account is defended in Mike Martin's work. Martin moots a structural difference between sensory (episodic) memory and perception. Both types of mental event are directed towards particular events and objects. However, the relation is different in each case. Memory is "the representational recall of … an experiential encounter" (2001: 270) with a particular event or object. That is, when one sensorily remembers an *f*, one does so through recalling (imagining) a particular past occasion of consciously experiencing an *f* (cf. 2001: 273f.). In contrast, perception is a direct (representational or presentational) encounter with said particulars. With this account in play, there will be no straightforward chord-cum-cacophony objection since the present tone will be structurally privileged within our awareness.

Crude memory theories nonetheless face immediate difficulties and must be modified to cope with them. The crudest of theories will seek to explain our experience of succession in a triad case like *Case A* by claiming that the experience of succession we have at time *t*, when the *G* is heard, arises because when we hear the *G* sound, we simultaneously auditorily remember hearing the *E* and the *C*. However, as Dainton rightly points out, this theory is insufficient to account for the phenomenon since it is possible to hear a *G* and auditorily remember hearing a *C* and *E* heard many hours or days ago. No theory should predict that this would amount to an experience of succession. Thus, a constraint on any memory theory is that "[m]y memory must register the temporal distance between present and past experiences" (Dainton 2000: 124). Similarly, merely saying that we

[209] Cf. Hamilton's remark, "Without memory, each indivisible, each infinitesimal, moment in the mental succession would stand isolated from every other, – would constitute, in fact, a separate existence" (Hamilton 1861: vol.1, 205; quoted in Anderson and Grush forthcoming). See also Ismael ms: 9 which makes effectively the same point.

experience succession because we hear a *G* whilst remembering a *C* and an *E* won't do since that will fail to distinguish between more complex successions. It will fail, for example, to distinguish hearing *C*, *E*, *G* from hearing *E*, *C*, *G* or *C* & *E*, *G*.

The most promising idea that Dainton considers which meets these concerns is an appeal to "a distinctive sort of memory" which he terms "immediate short-term memory" (125). In contrast to "ordinary long-term experiential memory" such a faculty is supposed to provide memories which are "a lot more complete and accurate than our typical long-term experience-memories" (125) and will also be (wholly?) involuntary and automatic in contrast to ordinary long-term experiential memory which is "to a large degree voluntary … [and] subject to our will" (126).[210]

Such differences in kind, at least prima facie, avoid the initial objection since short-term memories are by their very nature, *short-term*. They are also said to be involuntary, which may meet worries to do with the passivity of temporal experience (cf. Husserl 1964: §20). This principally leaves us with the problem of complex successions. This is avoided if we allow that we can have memories of *experiences of succession*. Thus, the improved account runs as follows.

> First I hear *C*; I then hear [*E*], the experience of which is automatically accompanied by a short-term memory image corresponding to my hearing *C*; I then hear [*G*] and as I do so I have a short-term memory of an *experience* of succession: '*C*-being-followed-by-[*E*]'. (2000: 126)

This sophisticated memory theory has a certain *prima facie* plausibility. However, Dainton contends that the general claim that "memory is largely or wholly responsible for our experience of time" (2000: 123) is untenable.[211] In what follows, I argue that Dainton's first two objections can be avoided by adopting a non-standard version of the memory theory. However, if this non-standard theory is to avoid Dainton's third objection, it must reject Strong PSA. Thus, in the final analysis, memory theories cannot account for temporal experience within the confines of Strong PSA.

---

[210] The psychological literature standardly recognizes at least three kinds of memory systems, in relation to vision, specifically: iconic memory, visual short-term memory (VSTM) and long-term memory. Thus, memory theorists are hardly being speculative in positing distinct kinds of memory. On the other hand, there are evidently empirical factors to be considered if the theorist appeals to a particular system.
[211] Tye (2003: 88) also objects to memory theories along similar lines.

5.1 The Illusion Objection

Dainton's first objection goes back at least to Husserl's discussion of Brentano's version of the memory theory. There, Husserl suggests – and indeed tells us that Brentano concedes – that Brentano's theory is an error-theory. For, according to Husserl, Brentano himself does not see his theory as grounding the direct perception of succession and alteration at all. Rather, it explains why it *seems* as if we do so perceive.

> We believe that we hear a melody, that we still hear something that is certainly past. However, this is only an illusion which proceeds from the vivacity of primordial association. (Husserl 1964: 33)

A number of thoughts are conflated here. In particular, there seems no reason to think that we do not often see and hear things that occur in the past – a supernova or a thunder-clap, for example. Nevertheless, Husserl also seems to be pointing out that, on a memory theory such as Brentano's, all we ever *hear* is the current sound *independent* of other sounds. Nothing else is the object of a *perceptual* act. Given this, it seems that we have not accounted for direct experience of succession at all.

In reply to the objection, it might be argued that hearing a melody just is hearing the current note whilst remembering the past note. But this move does not seem to avoid a point which Dainton raises against Broad's similarly structured theory, viz., that his theory "has the consequence that awareness of change cannot be as immediate as awareness of simultaneity" (2000: 154). The central thought here is that the memory theorist cannot avoid the consequence that 'perception' of temporal relations, if it should be counted as perceptual at all, is a second rate or derivative kind of perceptual experience. That flouts the phenomenological datum we began with.

5.2 The Complexity Objection

Dainton continues his attack on memory theories by arguing that the more sophisticated, nested memory account is unacceptable on phenomenological grounds. As he puts it,

> the complexity of this proposal counts against it. Simply hearing the sequence *C-D-E* does not seem to involve intricate compound memories of the required sort. (2000: 127)

Even allowing for the specialness of short-term visual memory and its automaticity, there does surely seem something phenomenologically off-key about the appeal being made to such a form of memory in this context. Consider shutting your eyes during a gory scene in a horror film and being unable to avoid visually recalling what you have *just* seen. Assuming this is a case of short-term memory (if it is not, we need to press further on that notion), it seems clear that the phenomenology is very different from that of a perceptual act or of any act that takes place during ordinary experience. Worse still, the nested-memory theory is positing a plethora of such experiences. It is hard not to sympathize with a request to know where they are and what phenomenological reality they have when it comes to ordinary temporal experience.

One might object that there is a great difference between *mere* remembering and remembering in combination with perception. However, again, the thought is not very persuasive. If the two kinds of act are discrete and independent, something needs saying to explain our failure to introspect any acts of short-term memory during perception. And the memory theorist seemingly has nothing to say here.

5.3 A Common Theme

The two objections just raised are closely related. In particular, they arise because we make a certain assumption as to what is essential to an act's being an act of memory. Although Dainton allows that the memory theorist may appeal to a certain distinctive kind of short-term memory which is richly detailed and automatic, he nevertheless assumes that the appeal will be to *distinct, discrete acts of recall with their own objects* which occur alongside perceptual experiences. On that picture two things seem plausible.

(a) That it is only the *original* perceptual acts (which acts of recall are joined to or simply simultaneous with) that are genuine acts of direct perception. This grounding the illusion objection.

(b) That temporal experience should be effectively resolvable into a number of distinct and phenomenologically discernible acts which differ only qualitatively from ordinary acts of short-term recall. This grounding the complexity objection.

In the next section I sketch a way in which we might broaden our conception of how memory might be involved in temporal awareness which avoids these objections. Merely invoking automatic, short-term memory does not go nearly far enough.

6. Retentiveness without Reminiscence: a Non-Standard Memory Account

G.F. Stout considers a nice example of experience of succession, hearing a postman's daily 'rat-tat' on one's door. He remarks,

> There is successive experience: the first knock is heard before the second. There is also retentiveness: the individual's experience, when the second knock occurs, has a character which it would not have if he had not heard the first. Further there is *no reminiscence*. In actually experiencing the second sound the subject does not definitely discriminate it as a present occurrence from the first as a past occurrence. The second treads too closely on the heels of the first to admit of such discrimination between past and present as such. So there is no remembrance of the past as such. This just emerges subsequently when the whole experience of having heard the double knock is remembered. (1930: 170)

Stout makes two claims of importance in this passage. Both apply *mutatis mutandis* to any experience of succession.

(1) There is retentiveness: the individual's experience, when the second knock or note occurs, has a character which it would not have if the individual had not heard the first knock or note.
(2) There is no reminiscence (discrimination, remembrance).

How should we think of this combination of retentiveness without reminiscence?[212] The *Oxford English Dictionary* defines 'reminiscence' as:

> The act, process, or fact, of remembering or recollecting; sometimes *spec.* the act of recovering knowledge by mental effort (cf. *recollection*). (Simpson and Weiner 1989)

Thus, reminiscing is something one *does*. It is a distinct mental act with its own character. How should we distinguish retentiveness in contrast to this?

Not by appeal to any standard classification of memory into, say, the procedural, semantic and episodic. These distinctions are distinctions amongst the grammatical *objects* picked out by the complement clauses in sentences of the form 'S remembers …'. Reminiscence and retention are not distinguished in terms of their objects. Nor is the distinction helpfully thought of in terms of the long-term versus short-term memory or indeed, in the visual case, between short-term visual memory and iconic memory. Crucially, retentiveness is not intended by Stout to be conceived as a distinct mental act or process of re-acquaintance with some particular object, event or event-phase (the 'rat' of the 'rat-tat', for example). If it is not a distinct act, what can we say positively about retention?

The key remark here is Stout's claim that "the individual's experience, when the second knock occurs, has a character which it would not have if he had not heard the first". Brian O'Shaughnessy echoes this claim in his recent discussion of the temporal properties of experience in general (which he takes to include intentional action). He then goes on to raise a question of obvious concern for us. "But why describe this as an exercise of *memory*?"

> The reason is, that had he not been acting [more generally: experiencing] thus in the past he would not be acting [experiencing] thus in the present, so that present experience

---

[212] The retention versus reminiscence/recollection distinction plays an important role in Hodgson's neglected work on time consciousness and is intimately related to Husserl's distinction between primary and secondary remembrance. See Hodgson 1898 and Husserl 1964. (Husserl also comes to use the term 'retention' as his thought develops.) Though arguably Stout is rather uncharitable in his reading of Hodgson, their debate in *Mind* (Stout 1900a, 1990b, Hodgson 1900) is instructive in pressing the question to what extent Hodgson rejects the traditional memory model. Precisely the same question needs pressing with respect to Husserl. See Anderson and Grush forthcoming for evidence that Husserl was aware of the *Mind* debate and Hodgson's work more generally.

> must both unite with and depend upon past experience. This means that the past must in some sense be *co-present with* the present, and such a co-presence is a mode of remembering. Doubtless it is a developmentally early form of memory, to be supplemented later by additional less primitive ways of relating to one's past, notably cognitive modes. What in effect we are concerned with here is the tendency on the part of experience and its given objects to unite across time to form determinate wholes. (2000: 56)

The important point in these passages is that both Stout and O'Shaughnessy insist

(a) That a subject's current experience can depend constitutively on how they have been experiencing in the recent past; and
(b) That the fact that an act has such a character is sufficient for it to count as a form of memory.

Thus, what it is like to undergo the experience one has of the postman's 'tat,' is constitutively dependent on the fact that it is a 'tat' which has been immediately preceded by experience of a 'rat'. One would not be experiencing thus in the present were it not for one's experience in the past.

O'Shaughnessy suggests that the constitutive link to the past involved in such experience is sufficient for us to think of such experience as an act of memory. Someone sympathetic to the constitutive claim may resist this further move. At this point, we need to consider a question raised by Mike Martin in his discussion of episodic memory, namely: what, if anything, "the varieties of memory that we mark out in natural language have in common that should make them all memories?" (2001: 261). Martin's limited positive answer to this question is that all forms of memory are ways of preserving cognitive contact. That is, we should think of memory in general as the retention of past psychological achievement.

For example, semantic memory is the preservation of past knowledge, episodic memory the preservation of past apprehension. As Martin puts it,

> Just as we can differentiate the kinds of cognitive contact and the objects they have, so too can we differentiate the kinds of memories that result. We can then conceive of memory in general as the preservation of cognitive contact in general … (2001: 266)[213]

On Stout and O'Shaughnessy's picture, in hearing succession, one's perception of a past tone does not merely leave a causal trace on current perception. Rather, current experience is constitutively dependent on past experience. This amounts to a retention of a past psychological achievement and thereby to an act of memory.

Martin distinguishes different forms of memory in terms of the different psychological successes retained. We cannot distinguish retention without reminiscence in this way as the object retained is the same as that retained in episodic memory: past perception. However, one possible way of fleshing out the concept of retention without reminiscence in contrast to episodic memory would be as follows. One first notes that listing the objects presented (or represented) in perception does not suffice uniquely to characterise our experience. To do that one needs to specify the *ways* in which things are presented in experience. These can make a difference to phenomenal content too. Thus, perceptual phenomenology cannot be exhaustively characterised in terms of the presentation of objects.

Consequently, experiences with qualitatively identical objects can still differ in phenomenal character so long as those objects are presented in different ways. For example, a 'tat' preceded by a 'rat' can be heard as a successor to a past 'rat', as the second part of a larger auditory event. On the other hand, a 'tat' not succeeding a 'rat' can be presented as emerging from past silence.

In this light, Stout's distinction can be understood as a distinction between two *ways* in which a single kind of cognitive contact can be preserved. In reminiscence, the retained or preserved contact is manifested in distinct acts of episodic memory. In retention, in contrast, past cognitive contact is preserved or retained as part of the character of a fresh act of acquaintance with a present object. The very encounter with the present itself is a

---

[213] Matt Soteriou suggested to me that it would be better to think of what is retained in episodic memory as an ability to reacquaint ourselves with the particular past episode of apprehension in question. In the context of the account below of retentiveness one might think of the perceptual act itself as, in part, a manifestation of this just-acquired and possibly very short-lived ability.

way in which cognitive contact with the past is preserved because the way in which the current object of perception is encountered is constitutively dependent upon past experience. When I experience the postman's 'tat', cognitive contact with the preceding 'rat' is preserved as part of the character of the presentation of the 'tat'.

7. Back to Dainton's Objections

Armed with our new understanding of memory as retention without reminiscence, we can now address the illusion and complexity objections raised by Dainton above. On the new model, acts of perception are not joined with distinct acts of recollection so as to 'constitute' temporal experience. Relations of preserved cognitive contact constitutive of memory are not provided from outside the perceptual experience itself, by a concurrent remembering or reminiscence. There is only one act that counts both as a perceptual act and as a manifestation of memory.

As a result, there is no complexity objection. We should precisely *not* expect to find phenomenology redolent of independent acts of episodic recall (short or long term) during perception. The kind of memory involved is quite different – retention *without* reminiscence. There is only one act and only one object. Of course, this object can be seen or heard in complex *ways*. A 'tat' can be heard as successor to a 'rat', a note as a culmination of a glissando, a word as quieter or louder than the preceding phrase.[214] This complexity is no problem; our experience *seems* to be complex in these ways. That is the datum.

The illusion objection arose because the standard memory account began with momentary 'direct' perceptions to which memories were joined. It was then hard to see how the joining of memories to perceptions could really account for our *direct perception* of temporal relations. By taking as a starting point genuinely perceptual acts plus memories, the memory theorist seems to be conceding that experiences of temporal structure are

[214] Indeed, the various ways one can hear notes is potentially limitless. For example, consider the following remark by Tenney and Polansky: "… for the musician, a piece of music does not consist merely of an inarticulate stream of elementary sounds, but a hierarchically ordered network of sounds, motives, phases, passages, sections, movements, etc. – i.e., time-spans whose perceptual boundaries are largely determined by the nature of the sounds and sound configurations occurring within them" (1980: 205).

not perceptual acts. At the very least the theorist must concede that perception of temporal relations is less basic than other forms of perception, that as Dainton put it, "awareness of change cannot be as immediate as awareness of simultaneity" (2000: 154).

The new single-act theory avoids the illusion object. On this account the perception of the current object of awareness *is* itself a manifestation of one's retained cognitive link to an object of one's past perception. The perceptual act is intrinsically an act of memory. However, there is also no sense in which it is not a perception proper. It is a retention *and* a perception, and, in virtue of being *both*, a perception of temporal structure. Thus, successiveness is as much a part of perceptual experience as the other ways in which objects are presented.

8. Sounds and Durations

Stout and O'Shaughnessy implicitly reject Strong PSA, holding that past perceptual experience is *constitutive* of current experience: one could not be experiencing the way one is now if one had not been experiencing so in the past.[215] However, it is not obvious why an adherent to Strong PSA might not agree that temporal experience must be understood in terms of the *way* the object currently being experienced is being experienced but deny that the individuation of such ways is *constitutively* dependent on past experience. I now argue that if the non-standard memory theory is to be applied to temporal experience in general, Strong PSA must be abandoned. In particular, I want to suggest that, even if the above account can explain our experience of succession (something I see no reason to grant at this stage), it cannot explain our experience of *individual sounds* themselves. To apply the above account to our experience of an individual tone one would need to claim that, at any moment, one's experience was of a 'tone-phase' heard as a continuing on of earlier tone phases – one's current experience being dependent (constitutively or otherwise) on one's prior experience of earlier phases. Whether or not we appeal to a constitutive dependence, this account can seem plausible if we think of how we hear notes sustained over some reasonable period – each current phase is heard as a part of a longer note. However, as I now argue, unless the relation is constitutive, the account

[215] However, both philosophers also seem to subscribe to the idea that there must be some present tense consequence of these irreducibly temporal relations. If this were not so, a simple-minded account would have sufficed. Thus, their theory fits best with Weak PSA.

cannot be a general one since it presupposes a form of temporal experience, namely, our perception of tone phases.

Dainton hints at the problem as follows.

> If phenomenal temporality is wholly the product of memory, … our experience of even a single brief tone must be explained in terms of involuntary short-term memories. But memories of what? The answer must be: a succession of strictly durationless experiences. (2000: 127)

According to Dainton this view "suffers from a very severe plausibility problem … it is hard to believe that we are not immediately aware of some duration in experience. Is a strictly durationless auditory experience even possible?" (ibid.). The central charge is that the memory theory must appeal to strictly durationless auditory experiences combined with a doubt as to whether such experiences are possible.[216]

What everyone (and I am no exception) agrees with is the claim that any auditory experience will present or represent an object as having a duration of some length. As John Foster writes, "it is inconceivable that there should be a sensation of sound which was not the sensation of a sound-filled period" (1982: 256). Likewise, Husserl insists, "[e]very tone itself has a temporal extension …" (1964: 43).[217, 218]

Now, assuming that we are happy with a distinction between the temporal structure of experience and the temporal structure of the objects experienced, it is not conceptually

---

[216] In a wonderful, brief note, entitled 'The Apprehension of Time,' Prichard makes just this objection to the Kantian view that "what we call hearing a sound really consists in the act of perceiving or hearing what is now given – *together with*… acts of remembering and thinking of certain things as related to what is now given" (1950a: 47). Prichard at first concludes, quite rightly, that "perceiving anything which has a duration – is absolutely ultimate and cannot be resolved into anything else" (48). Unfortunately, within a few lines he considers that this leads us into contradiction (due to his explicit commitment to Strong PSA: "the hearing of the whole note must take place at a definite moment") and declares that "it seems *impossible* that we hear a sound" (49).

[217] See also O'Callaghan 2007. Contrast the view implicit in this passage from Brough, "If a painting really could be seen in a single instant, its internal structure or content would be the visual equivalent of a musical composition consisting entirely of one sound, say, a crash of cymbals" (2000: 227).

[218] Although Husserl recognizes the problem here, I am not clear what his solution to it is. Indeed, he retains an ambivalence when, for example, he talks about "a sound (or a tonal phase) in the now point" (1964: 57). Indeed, in certain places he offers what sounds like a traditional, Reidian form of the memory theory. Thus, "Therefore at any given time I hear only the actually present phase of the tone, and the objectivity of the whole enduring tone is constituted in an act-continuum that is in part memory, in smallest punctual part perception, and in further part expectation" (1964: 24).

incoherent to think that we could *at some isolated instant* have an auditory experience, an instantaneous experience of a sound-filled period. However, we are now back to issues encountered in relation to SPT above. There I argued that there are good reasons to reject any theory which attempts to account for temporal experience in terms of our experiencing durations at an instant. When we reflect upon the nature of our experience, we come to appreciate that the apparent temporal structure of experience maps the apparent temporal structure of the world experienced. So any experience of a sound, something which must have a duration, will seem to rational reflection to itself possess a duration. However, in the domain of experience, it is, I claimed, not possible for experience systematically to seem some way to rational reflection, and yet not be such a way. Thus, we cannot systematically be in error when we judge that our experiences of sounds are not instantaneous.

As a result, a theorist who appeals to memory cannot apply that account to all temporal experience unless they abandon Strong PSA for just the same reasons we encountered in the case of SPT. Whichever way we turn, a general account of temporal experience cannot subscribe to Strong PSA.

9. Conclusions, Weak PSA and a Terminological Note

The traditional and contemporary debates on temporal experience assume that SPT and memory theories are clear rivals.[219] However, in the light of the argument of this chapter, they begin to look rather similar. The revised memory theory posits constitutive relations between past and present experience. In order to understand the nature of present perceptual experience, one must look beyond the instant. Similarly, Dainton's revised SPT posits a primitive relation of what he calls 'co-consciousness' between past and present experiences. Again, the thought is that, in order to understand the nature of present perceptual experience, one must look beyond the instant. If one thinks that such constitutive links are sufficient to think of temporal perception as a case of memory, then we can think in terms of a memory theory. On the other hand, it would not be inept to give the name 'the specious present' to that interval of experience standing in constitutive

[219] See especially, Dainton 2000 and Kelly 2005 for recent work that thinks in these terms.

relation to present experience. Thus, the theory in question can also be thought of as a SPT.

Both revised accounts reject Strong PSA. According to both, for example, if one has not been having an auditory experience in some temporal window around the current moment, one will simply not be in a position currently to have an auditory experience at all. However, this is not quite to say that there is no distinction between the two accounts. It may be that although they do not differ at short-time scales, the memory account can explain perceptual phenomena at longer time scales more adequately than a specious present theory can.

It is worth flagging at this point how my terminology relates to Dainton's recent talk of Extensionalist and Retentionalist views. According to the Retentionalist, "our experience of change is a product of simultaneously existing features of momentary episodes of consciousness" (2008: 626). The Extensionalist denies this, rejecting the idea that "our consciousness is confined to an instant" (ibid.). Clearly my argument so far rejects Retentionalism in this sense and so sides with Extensionalism. However, there are reasons for preferring the traditional terminology that I have used, at least for present purposes. One problem with the Extensionalist/Retentionalist dichotomy as glossed is that it fails to separate out acts and objects, no doubt because Dainton himself rejects the act/object distinction. However, the distinction is necessary to make sense of the range of views on the table. Thus, we need at the very least to consider a matrix of the kind below.

Now we might choose to call the left-hand column views, 'Retentionalist' since they reject extended *acts* of consciousness. However, the term 'Retentionalist' immediately suggests the involvement of *memory*. Thus, it is far from ideal as a label for a group of views which include views that think memory is not involved in temporal experience. Moreover, these views reject the involvement of memory for two very different reasons. On the one hand they might think that strictly speaking there is no such thing as temporal experience (as Reid does). On the other hand, they might believe in instantaneous acts with extended contents as part of these acts' specious presents. Perhaps it is better to think of the left-hand column views as all *Pointilliste* views and of some of them as Retentionalist.

		Is temporal experience the product of momentary *acts* of consciousness?	
		Yes.	No.
Is temporal experience the product of momentary *objects* of consciousness?	Yes.	Pure snapshot views – experience is made up of instantaneous awarenesses of instantaneous contents. Memory theory as conceived by Reid (and as conceded by Brentano).	No takers.
	No.	Traditional Specious Present Theory. Memory Theory as it is often characterised and arguably as Husserl and Hodgson ultimately conceive it.	Overlap Specious Present Theory (e.g., Foster and Dainton, though with act/object distinction). Revised Memory Theory (e.g., O'Shaughnessy).

There is another difficulty with the term 'Retentionalist', however. This is that some memory theorists who insist on the traditional distinction discussed above between retention and reminiscence wish to claim that temporal experience involves retention without reminiscence and that this allows them to place their view in the bottom right of the matrix. Thus, the term 'Retentionalist' is doubly problematic as used here and I have avoided it. It is far less objectionable to call the views in the bottom right of the matrix, 'Extensionalist', though, as noted, it is not entirely straightforward that the two views there mentioned are of a piece. However, my preference is to think in terms of the stream of consciousness first and foremost and not first off in terms of building the stream up from extended acts as the term 'Extensionalist' can suggest. One question this chapter began with was whether we really need anything beyond a very simple theory of temporal consciousness once Strong PSA is abandoned. To endorse an 'Extensionalist' view, even if only in name, is to prejudice the answer to that question. Thus, again, I have avoided the term.

By way of initiating further debate once it has been conceded that Strong PSA is false, I want to end this chapter by reconsidering how Weak PSA looks in the light of our rejection of Strong PSA. Recall that Weak PSA was the following claim.

> **Weak PSA** Irreducibly temporal facts are *necessary but insufficient* to explain the difference between *Cases A* and *B*. They need, in addition, to have present tense psychological consequences. It is these consequences which are absent in *Case C*.

The argument for Weak PSA was an appeal to a supposedly possible *Case C* where one experienced each note of a triad but yet failed to experience the final note as succeeding the tonic and major third heard before it. However, *Case C* does not motivate Weak PSA construed as a universal principle. Read as such, Weak PSA tells us that *merely* having *any* durational/successive auditory experience is insufficient for experiencing a duration/succession. What would support this general claim? Presumably the thought is that *C*-style cases can be conceived in *every* case of temporal experience.

In particular, presumably the Weak PSA theorist thinks we must be able to conceive of a distinction between the following two cases.

Case B

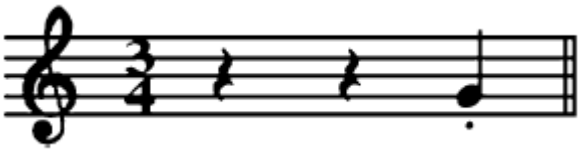

Case D

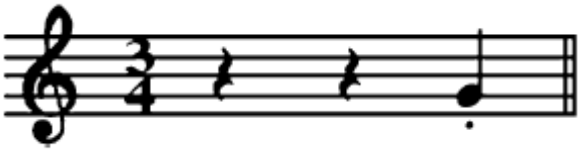

In *Case B*, I have what one might think of as an ordinary experience of a *G* natural. In *Case D*, in contrast, though I experience each individual phase of the note, my experience of the final phase is independent of my experience of preceding phases. If such a case were conceivable it would, as appeal to *Case C* was held to, show that at least Weak PSA

must be true. However, as should be obvious now, at some point this style of argument *must* break down. If it does not, we will be committed to a case where the final instant of our experience is independent of all our previous experience. This either directly flouts the claim that any auditory experience will be experience of a sound-filled *period* or runs into all the problems of the traditional SPT.

If this is right, then Weak PSA cannot be motivated all the way down on the basis of the conceivability of cases such as *D*. Once this is seen, it must be acknowledged that a simple-minded account must be correct *in at least some cases* – at least when the times in question are very short. In other words, even Weak PSA cannot be a universal principle. Moreover, this problem concerning the general validity of Weak PSA greatly complicates assessment of any argument such as the one we began with contrasting *Cases A* and *C*. As it was set out above, such an argument did not merely claim that pairs of cases such as *A* and *C* involved similar histories. Their histories were claimed to be similar *in all relevant respects* right up to the moment in question, *t*. But the simple-minded theorist now notes that our experience at *t* is not independent of (at least very nearby) surrounding experience. It is thus *not* obvious that one could – at the very moment the *G* sounds – enjoy experience which was not in part experience of succession unless there was some difference in history between the two cases. Thus, even if a simple-minded account is ultimately indefensible, more work needs doing to show that it is. On the other hand, what is certain is that Strong PSA is untenable if we do not wish to follow Kant, Reid and others in their denial that we perceive temporal properties. And that is to say that experience cannot be homoeomerous if we are to explain our perception of temporal properties.

Chapter Seven:

Indiscriminability and Experience of Change

Force 3: *Gentle Breeze – Leaves and small twigs in constant motion …*[220]

1. Overview

It is overwhelmingly plausible to hold that we often experience objects as being in constant motion. It is also very natural to appeal to the claim that there are changes too slight to be perceived in accounting for our experience of slow moving objects like hour-hands, and in explaining supposed failures of the transitivity of indiscriminability. However, Fara (2001) argues that this natural explanation precludes our experiencing constant motion. We therefore seem skewered on the horns of a dilemma. Here I argue that Fara's puzzle about constant motion can be resolved by thinking carefully about the anhomoeomery of experience in this context. As such this chapter forms the final part of my case against homoeomery. It also provides the opportunity to draw together hitherto disparate elements in the thesis, closing with a unified account of our experience of change.

The following course is charted. First, I introduce our natural explanation of non-transitivity in slow motion, 'hour-hand' cases and Fara's puzzle concerning that explanation. I raise two concerns about Fara's own proposed explanation of our experience in such 'hour-hand' cases. I then explain how a proper appreciation of the special conditions on experiencing change (as opposed to mere variation) allow one to respond to Fara's challenge by appealing to the anhomoeomery of experience in time. I thereby safeguard both our experience of constant motion as well as the natural explanation of our experience of slow changes. In contrast, as Fara in effect demonstrates, failure to appreciate the exceptional conditions in play with experience of

[220] *The Beaufort Scale: Land based specification*, UK Met. Office.

change renders it impossible to make sense of our experience of constant motion. In other words, unless we accept that experience is significantly anhomoeomerous, we cannot account for our experience of constant motion.

2. Experience of Slow Motion and Non-transitivity

The non-transitivity of indiscriminability in relation to phenomenal continua supposedly arises because it is thought possible to have three scenes, $S_{1\text{-}3}$, such that S_1 looks the same as S_2, S_2 the same as S_3, but nonetheless S_1 does not look the same as S_3.[221] A paradigm case is taken to be "phenomenal motion too slow to be momentarily perceived" (Goodman 1977: 203). As Dummett explicates:

> I look at something which is moving, but moving too slowly for me to be able to see that it is moving. After one second, it still looks to me as though it is in the same position; similarly after three seconds. After four seconds, however, I can recognize that it has moved from where it was at the start, i.e. four seconds ago. At this time, however, it does not look to me as though it is in a different position from that it was in one, or even three, seconds before. Do I not contradict myself in the very attempt to express how it looks to me? (1975: 316)[222]

The kind of case Goodman and Dummett have in mind is our experience of an hour-hand as it progresses slowly around a clock-face. After looking at the hour-hand for some time, one perceives the hand as being in a distinct position and so infers that it has moved. This kind of experience contrasts with a case in which we attend to a constantly moving second-hand and can straightforwardly see the second-hand moving around the clock-face. Here the change in position of the second-hand over time is an object of our experience; change itself is not experienced in the hour-hand case.

---

[221] In each case what is relevant is how things look to the same subject. Hellie (2005: §1.3) advises against discussing non-transitivity in terms of the vernacular of 'looking the same'. Certainly, we should be careful in unpacking that notion. However, we do talk this way and that needs accounting for. Moreover, Hellie himself offers a plausible understanding of 'looking the same' in terms of indiscriminability – his reading (4). For much more on indiscriminability, see below.

[222] See further discussion in Chapter Three above.

We should also contrast experiencing mere variation with experiencing genuine change. Looking at a full colour spectrum one perceives the variation in colour around the wheel but one does not perceive *change*. Change, as it is understood here, is variation of the properties of some object *over time*.[223] The properties of an ordinary colour spectrum do not change in this sense. An example of genuine change is variation in body temperature over the course of a day. Of course, most properties change and also merely vary. Body temperature is not the same in all parts of the body. Thus the body's temperature both changes over time and also merely varies between its parts.

3. Fara's Puzzle

In discussions of non-transitivity the fast, second-hand case is typically ignored. When it is explicitly considered, it is typically held to present no special problems. Wright does discuss the case of "seemingly continuous processes in time" (1975: 345), acknowledging, quite rightly, that they do not come "ready made out of finitely many stages" (ibid.). However, he treats the case as merely a less "artificial", more "dignified" version of the colour patch case. He achieves this by analysing it in terms of the discriminability or otherwise of *stages* of the process, where a stage is "an instantaneous exposure, as it were, of the process at [a] point [in time]" (ibid.). Likewise, Burgess distinguishes between fast, dynamic, and static cases but claims that "we can know *a priori* that [a continuous change] must be … divisible" (1990: 212) into "observed [temporal] stages of the process of change" (208). On this basis, he concludes that there is no essential difference between the cases.[224]

Yet, as Fara highlights, the second-hand case presents a serious difficulty for standard explanations of non-transitivity. The standard explanation of slow-motion, hour-hand cases, is that there are changes of position which are too slight to be perceived, the change in position of an hour-hand over the course of a few seconds being just such a change.[225] However, as Fara notes, this explanation is "very suspect … since it should

---

[223] Goodman (1977: 271-2) emphasises this ordinary distinction in a different context but does not apply it to his discussion of non-transitivity.
[224] Cf. De Clercq and Horsten, "… consider a process of change in respect of some observable property (a determinable such as colour, position or pitch). The process is composed of stages…" (2004: 440).
[225] This claim that there are changes too slight to be perceived is analysed in detail below.

leave us wondering why not every experience of motion is an experience of slow motion." As she puts it:

> If the reason that the hour-hand strikes us as still-looking for any twenty-second interval is that we cannot visually represent a change in position as small as, say, 1/6° (on a normal-size clock), then the second-hand should look still for any 1/36 second interval, for it changes its position only that amount during such an interval. But, when we watch the second-hand moving, it never looks still – it appears to be constantly moving. … Although the [proposed explanation] would explain why we experience slow motion, the explanation seems too strong, since it seems to preclude the possibility of experiences of constant motion. (2001: 926-7)

Think of Fara's challenge like this: the explanation of why there are changes of position which are too slight to be perceived carried across to the case of fast, second-hand motion conflicts with the existence of experience of *constant* motion.[226] Yet experience of constant motion is actual, so the explanation must be false.

Fara detects no relevant difference between experience of fast and slow motion.[227] Consequently, she must offer an alternative explanation as to why we don't perceive changes even in slow cases. On Fara's view, if the explanation did apply in the slow case, it would carry over to the fast. In the next section I raise two doubts about Fara's proposed explanation. However, at the heart of the chapter is the claim that perception of change is a fundamentally special case. The reason is that there are conditions on experience representing change that are not in play in hour-hand or colour spectrum cases, i.e., that do not apply to its representing distinctness or mere variation.[228] Appreciation of these conditions allows one to respond to Fara's challenge, safeguarding both our experience of constant motion and the plausible explanation of our experience in slow-change cases; failure to distinguish the two cases renders it impossible to make sense of our experience of constant motion.

---

[226] If I experience an object as *constantly* moving through some interval over some period, then there is no time during that period at which I see the object and yet it does not look to be in motion. Constant motion does not require constant velocity. It must also be distinguished from strictly *continuous* motion which involves experiencing motion through all proper parts of the interval.

[227] Thus: "It is intended that remarks in this discussion be taken to hold *mutatis mutandis* for perceived differences between two things at a single time, as well as for perceived changes in a single thing over time." (916, fn.14; cf. fns.20 and 21)

[228] As elsewhere I talk about what experience *represents*. Naïve realists can substitute 'presents' throughout without affecting the argument.

4. Fara's Explanation

Fara's own reply to the puzzle she raises is to suggest that the tiniest of slow changes may indeed be perceived but just not *noticed* (927-8). There are two immediate problems with this claim, one minor and one more serious.

Firstly, according to Fara, "*Noticing* the change in apparent position requires not only that there be an apparent change, but also that we believe there to be one" (928). It is not clear whether Fara thinks adding belief is sufficient or merely necessary for noticing. In any case, adding belief is neither necessary nor sufficient for noticing. It is not sufficient since I might believe that there was an apparent change for a congeries of better or worse reasons (or none at all). Perhaps my guru tells me that there has been an apparent change. Clearly, this does not mean that I notice it. Similarly, it is not necessary since we can make sense of someone genuinely noticing a change and yet judging that there was no apparent change and consequently forming no belief. Perhaps again the subject is convinced by his guru that there could not have been an apparent change, any appearance to the contrary being explained by a failure to properly attend to his experience. Certainly Fara (e.g., 927) is happy to allow that our judgements about the character of our experience can be mistaken; thus more needs saying about noticing.

There is a more pressing second concern. The explanation that Fara proposes is meant to explain why we *cannot* see hour-hand motion over a period of just a few seconds. To explain this whilst acquiescing that such motion is really apparent, one must argue that we *cannot* notice such apparent motion. However, whilst a distinction between what we experience and what we notice of the scene experienced seems in order, the idea that there are things we experience which we simply *cannot* notice seems extremely hard to make sense of.

At one point Fara notes that what she has "difficulty understanding is, if it is really the *look* of [two patches] which is different, why cannot this come across, so to speak, by looking just at *them*?" (915) What I have difficulty understanding is, if a change is really an *apparent* one, why this *cannot* come within our cognitive ken? To paraphrase her: there is a

heavy burden on the proponent of this explanation to persuade us that an *unnoticeable* apparent change is really an *apparent* change at all.

In glossing Fara's position on this point Martin suggests that her view involves the following commitment.

> Even if a subject may on occasion fail to notice the difference in look between adjacent samples, and indeed may be bound to fail to notice such a difference, nonetheless there is a difference to be noticed and which could be noticed. (2004: 77)

I find this commitment equally perplexing. If someone is *bound* to fail to notice a difference of some kind, how is it a difference which nonetheless *could* be noticed? Even embracing the distinction between impersonal and individual discriminability (see Chapter Three, §4) does not appear to help in this particular case, since no appeal is made to a specific deficit relative to which a subject or subjects in general fail to notice the difference despite it being impersonally noticeable. Any such deficit simply seems to be an aspect of vision or experience and as such will prevent any impersonal/individual distinction from gaining a useful grip here.

In what follows, I show that we can avoid these difficulties and retain the explanation of the slow motion case which Fara rejects.

5. Homogeneity

Focus on a case of reasonably rapid motion like that of a second-hand moving constantly around a clock-face from one to two o'clock, sweeping out an angle of 30°. You perceive the hand constantly moving from one o'clock to two o'clock. Now consider the following claim (cf. Fara's P2, 921).

> **Homogeneity** Variation in small enough spatial intervals appears homogeneous. That is, there is a small enough sector (with central angle $\Delta\theta$) which is such that all pointer positions lying within the sector appear the same.[229]

In the scenario as described, Homogeneity entails that variation in small enough *temporal* intervals appears homogeneous. That is, there is a small enough time (the time during which the minimal angle $\Delta\theta$ is swept out, i.e., $\Delta t = \Delta\theta/\omega$, where ω is angular velocity) which is such that all pointer positions during Δt appear the same.

Homogeneity has struck most philosophers as obvious.[230] In hour-hand and colour spectrum cases, Fara combines Homogeneity with the claim that phenomenal continua (i.e., cases of continuous looking variation) are possible to derive the conclusion that 'looks the same as' is not transitive. In the current context, we need not rely on the doubtful premise that we experience *strictly continuous* motion to generate a problematic conclusion. The claim that we experience *constant* motion suffices since, as already noted, this appears precluded by Homogeneity as formulated above.

In the hour-hand and colour spectrum cases, Fara accepts the 'Homogeneity plus phenomenal continua' argument as valid and instead questions its soundness. In particular, she argues that there are insufficient grounds for accepting either premise and moreover, that "the two premises are in so much tension with one another that their conjunction cannot be admitted as plausible" (922). Thus, in the end she affirms the "disjunctive claim: either there are no phenomenal continua after all, or it is not the case that sufficiently small parts of phenomenal continua are all homogeneous-looking" (ibid.). In the constant motion case, the disjunctive conclusion is rather that either there is no perception of constant motion or that Homogeneity is false. Given that we surely

---

[229] I have followed Fara in formulating Homogeneity. An important issue with this formulation is whether we should distinguish between: (i) pointer positions within $\Delta\theta$ look the same, and (ii) pointer positions within $\Delta\theta$ merely do not look distinct. (See Hellie 2005: 489, fn.9.) However, even if the weaker, second reading suffices to explain hour-hand cases, it does not avoid Fara's problem about constant motion. It would remain true that within a small enough sector no pointer positions would look distinct. Yet a constantly moving pointer looks to be constantly *changing* its position. And how could it look to be changing its position if over small enough sectors it did not look to occupy *distinct* positions (i.e., have changed position)? Note also that we need to give a precise sense to the notion of a position appearing the same. There is, for example, arguably a difference between saying that the pointer looks to be in the same position *relative to some fixed mark on the clock-face* and saying that it looks to be in the same position *relative to some previous position*. This nicety does not affect anything substantive in what follows.

[230] It is implicit in the 'dot' argument in Dummett 1975: 314-15. See also Wright 1975: 343 and, in relation to both Wright and Fara, De Clercq and Horsten 2004. However, contrast Hellie *op. cit.*.

do experience constant motion, this conclusion forces us to look more carefully at the motivation for Homogeneity.

6. Finite Powers of Discrimination

Homogeneity is typically motivated by the idea that our powers of discrimination are finite. As Wright avers, to reject Homogeneity would be "to suppose that we have infinite powers of discrimination" (1975: 346). However, intuitive support for the idea that our powers of discrimination are finite rests, I suggest, with the appeal of the claim that there are changes too small to be seen or, more generally, changes too small to be perceptually detected.

What exactly does this claim amount to? Fara distinguishes two readings:

> (a) For some sufficiently slight amount of change (in colour, sound, position, etc.), when we perceive an object for the entirety of an interval during which it changes by less than that amount, we perceive it as not having changed at all during that interval.
> (b) For some sufficiently slight amount of change (in colour, sound, position, etc.), we cannot perceive an object as having changed by less than that amount unless we perceive it as not having changed at all (as having changed by a zero amount). (917)

As Fara puts it, (a) tells us "that sufficiently small changes in our perceptual environment appear to us as no change," whereas (b) tells us "that there is a limit to how slight an apparent change can be" (917).

Fara goes on to argue that any acceptably modified version of (a) will collapse into a version of claim (b). In particular, as it stands (a) is unsatisfactory if it is meant to capture the thought that "there are changes too slight for us to detect by observation" because it entails that "whenever we perceived stasis we would perceive it as stasis" (919), something quite implausible in the light of several well-known illusions.[231] However, modifying (a) to avoid this consequence we arrive at a claim such as:

---

[231] See Kitaoka's striking examples at http://www.psy.ritsumei.ac.jp/~akitaoka/saishin17e.html.

> For some small amount of change ε, we cannot visually represent [or, more generally, have an experience as of] a change of any positive amount δ less than ε. (920)

And clearly "this is just to retreat to saying that our powers of discrimination are finite in the (b) sense" (920). I agree entirely and henceforth focus on (b).

7. A Crucial Distinction

In addition to the above statement of (b), Fara also writes the following.

> Looked at another way, what (b) denies is this: for every experience I could have as of a change, there is an experience I could have as of a lesser change (but for the exception just mentioned [that is, except for an experience of 'zero change']). (918)

However, it is not clear that this *is* what (b) as first stated denies. What the passage just cited denies is the following: for every experience as of a change that I might have, there is *some experience* I could undergo as of a lesser change. In other words, the passage denies the existence of a stand-alone experience which *only* represents a change below the threshold. Experiences *merely* as of very small changes are ruled out.

What is not ruled out by this passage, but may seem ruled out by (b) is the representation of change below a certain limit. Consider that an experience of a large change will *typically* involve the representation of sub-changes. My experience of a shooting star tracing an arc across the sky involves experiencing various parts of the motion – parts corresponding to the meteor's movement through proper segments of the whole arc traced. Read strictly, the passage above does not insist on any limit to how small these sub-changes might be represented as being. What it rules out is that very small changes might be the *sole* objects of a change-experience, i.e., that they might be experienced on their own and not as part of a larger change.[232] *More generally, just because we deny that slight*

[232] If we think of experience as a process or event then, plausibly, we might think of the process or event of experiencing the meteor's motion as involving sub-processes or sub-events corresponding to experiencing different parts of the motion. However, as I argue below, we should resist this idea in its full generality.

changes are imperceptible on their own, this should not be taken to commit us to claims about what is true in cases where changes over longer periods are perceived.

Thus, there are two different claims which we might want to make if we are trying to capture the thought that there are changes too small to be detected – two claims which Fara hesitates between in discussing (b).[233]

(b_1) For some sufficiently slight amount of change, we cannot have a perceptual experience as of an object merely changing by less than that amount unless we experience it as not changing at all.

(b_2) For some sufficiently slight amount of change, we simply cannot perceive an object as changing by less than that amount fullstop – even as part of a larger change – unless we perceive it as not changing at all.[234]

Even if Fara intends the (b_2) reading, the weaker reading (b_1) is sufficient to yield the denial she expresses in the quotation just given, viz., the denial that "for every experience I could have as of a change, there is an experience I could have as of a lesser change". In other words, (b_1) captures one clear sense in which there are changes too slight to be perceived; (b_1) tells us that only changes above a certain size, can in themselves be apparent changes. But the minimal size for a change experience need not be the same as the size of the minimal change detectable. Indeed, for all that (b_1) says, I may perceive the very finest degrees of some slight change and yet be incapable of having an experience which only represents a change smaller than that slight change.

---

[233] Compare and contrast: "(b) requires only that when we perceive an object to be changing , we perceive it to be changing by at least some given amount" (919) – reading (b_1) – and "the truth of (b) requires only that when we discern a change … there can be no interval during which the things appears to change … by some smaller amount than the threshold of discrimination" (919) – reading (b_2).

[234] John Morrison pointed out to me that talk of an absolute limit across all perceptual circumstances here is implausible, noting cases of viewing through microscopes and suggesting the notion of *change in one's visual field* to accommodate such cases. Even with this modification it may remain the case that some circumstances simply afford better discriminations than others. In the case in point, adding time-markings to the clock-face may change the perceptual circumstance in this way. Consequently, we will have to relativize (b_1) and (b_2) to perceptual contexts. I ignore this complication in what follows.

8. The New Account

If we accept (b_2), then, Fara argues, the non-transitivity of 'looking the same as' follows on the assumption that there are phenomenal continua. Likewise, Fara also argues that (b_2) precludes our experiencing constant motion. The reason is, she claims, that (b_2) commits us, *quite generally*, to there being finite intervals in which things look the same, i.e. to Homogeneity. Shortly I show that (b_2) does not, in fact, commit us to Homogeneity. However, first I want to show how by endorsing (b_1) but, for now, resisting (b_2) we can explain our basic intuitions concerning Homogeneity without precluding the existence of experience of constant motion. I first explain how this move is best understood and how it achieves the result just noted. I then show how the conclusions of previous chapters independently motivate the resulting picture of experience in the case of perceived change at hand.

If we only endorse (b_1), all that is ruled out is our *merely* perceiving an object moving some small finite amount, say seeing the second-hand sweeping out an angle of merely $\Delta\theta$ over a time Δt. On this proposal the visual system cannot *merely* represent a change of any positive amount less than or equal to $\Delta\theta$. However, experience of constant motion is not precluded. The reason is that there is no experience of motion of any kind in these situations since no change at all is being represented. On their own, changes so small are not apparent according to (b_1). The only possibility for phenomenal change and so experience of constant motion is over some larger interval strictly greater than $\Delta\theta$.

Plausible cases in which we experience motion (and so potentially constant motion) clearly involve intervals larger than $\Delta\theta$. In the example above, the clock is sweeping out a 30° angle and obviously $\Delta\theta << 30°$. However, (b_1) has no implications for cases where we perceive changes larger than $\Delta\theta$. In fact, (b_1) leaves open two main options concerning what we do perceive when we perceive changes larger than $\Delta\theta$.

One option is to insist that in virtue of perceiving the larger change one is able to perceive the tiniest of sub-changes within the relevant interval. If that were true, then Homogeneity would evidently fail to hold in such cases – indeed such cases would be cases of genuine, strict phenomenal continua. However, once we recognize the general point that just because we deny that slight changes are imperceptible on their own, this

need not commit us to claims about what is true in cases where changes over longer periods are perceived, we can account for constant motion without committing to perceiving the tiniest of sub-changes as with the above account. Indeed, whilst Fara leaves it open whether we might perceive the tiniest of changes (though not notice them), such a claim taken to its limit seems incredible. Is it really plausible to think of the visual system as affording discriminations down to 10^{-25}m say?[235]

The alternative approach denies that we perceive *all* the sub-changes of a perceived change. In some cases, certainly, we may perceive *certain* sub-changes in virtue of perceiving larger changes. But we do not see *all* sub-changes of perceived changes; there are some changes which we simply see *as such*. The content of experience is determinable in this respect. We can perceive a change over some reasonable temporal interval despite not perceiving all the sub-changes within that interval. This is not to say that we perceive the second-hand as stationary or unchanging through some sub-intervals, however. Rather it is to deny that we need say any more than that over these very small intervals we can be (in the process of) *experiencing change* over some larger interval.

In particular, we should insist that there are cases where *throughout* the course of some short period Δt, one is experiencing the second-hand sweeping out an angle of $\Delta\theta$ and yet: (i) it is not true that over any proper sub-period of Δt one *experiences* the hand sweeping out an angle of $\Delta\theta$; nor (ii) is it true that over any proper sub-period of Δt one experiences, or is experiencing, the hand sweeping out some *sub-interval* of $\Delta\theta$, say $\delta\theta$. Over sub-periods of Δt – indeed, throughout Δt – one is simply *experiencing* the second-hand sweeping out an angle of $\Delta\theta$.

On this account we can reject Homogeneity and respond to Fara's challenge regarding constant motion experience as follows. First, note how Fara's original argument had the following form. "If the reason that the hour-hand strikes us as still-looking for any twenty-second interval is that we cannot visually represent a change in position as small as, say, 1/6° (on a normal-size clock), then the second-hand should look still for any 1/36 second interval." However, even if a change of 1/6° and so a 1/36 second experience of the second-hand *on its own* would be an experience of no change, this does

[235] Comments from Tomas Bogardus helped me to clarify my position on this issue.

not mean that we cannot see constant change throughout, say, a one second period in which the hand turns 6°. Moreover, in contrast to the first account considered, saying this does not commit us to seeing changes of 1/6° during 1/36 second sub-periods of this second. Rather, as we probe the structure of the experience over the course of the second we must eventually stop and simply deny that experience has any finer structure. Again: over such timescales we simply see change. This means that we can allow for experience of *constant* motion throughout periods without insisting on experience of *continuous* motion over those periods in the strict sense that for every experience of change in position over an interval, there must be experience of some lesser change over some proper part of that interval.

Interestingly, on this new account it turns out that, although the $(b_1)/(b_2)$ contrast importantly highlights the possibility of there being certain changes which are only perceivable as part of larger changes, (b_2) need not be completely rejected after all. For, *pace* Fara, on the picture just sketched (b_2) does not straightforwardly entail Homogeneity. The above account denies that not being able to perceive an object as changing by less than some amount (unless it is perceived as not changing at all) commits us to pointer positions appearing *the same* throughout the brief periods where only such change is occurring. Instead it insists that we can be experiencing some larger change throughout a brief period in virtue of its being a sub-period of a larger change experience, and moreover, that that is all there is to be said about our experience in such cases.

In sum: Homogeneity fails precisely in those cases where there is experience of constant motion. Consequently, constant motion experience is safeguarded as well as the obvious truth cited in explanations of hour-hand cases, namely, that our powers of discrimination are finite. My account is attractive because of its ability to respond to Fara's challenge. What I now show is how the considerations we have encountered at length already in the thesis provide independent support for the picture of experience over time here exploited.

9. Anhomoeomery Again

Two key ideas were exploited in this chapter so far. First, that we may experience certain changes only as parts of larger changes; second, that change experience may be fundamentally determinable, in that we may experience change throughout some period without experiencing all the sub-changes of that change. Both ideas involve the thought that experience cannot remorselessly be broken down into independent parts. In the former case, the thought is that whether or not one sees certain slight changes depends upon what one experiences over some longer period. In the latter case, the thought is that whether one is *experiencing change* over some brief period depends upon whether one experiences change over some longer period.

As discussed at length in Chapter Four, it has traditionally been assumed that activities are homoeomerous where a process is homoeomerous just if none of its temporal parts has a nature (be that a matter of what activity the period is filled with or of what has been accomplished over the period) which depends constitutively on facts about the process beyond the period of the temporal part in question.[236] Following Taylor (1977), Soteriou (2007), and others, I argued that such an assumption was false of most ordinary processes. For example, I argued that walking and running are not homoeomerous down to instants or brief temporal parts. Recall that "walking can be characterized as an alternating sequence of single and double support" in contrast to, say, running which "involves alternating sequences of [single] support and nonsupport" (Enoka 2002: 179). Thus, a single support phase (which is all that will be going on during certain sub-periods of periods of walking) will not be sufficient an occurrence on its own to make it true that someone walked. Similarly, with talking one arguably has *at least* to utter a phoneme/distinctive feature (i.e., a genuine phonological unit) to have talked as opposed to have merely made noise.

Recall also that for all that it is natural to allow that if Paula Radcliffe was *running* throughout a two hours and fifteen minute marathon win, that she was running during all sub-periods of the two hours and fifteen minutes. However, in the light of the above, we

[236] As discussed in Chapter Four, I am inclined towards the view that experience is anhomoeomerous even when the condition of homoeomery is the following: an activity of Φ-ing by x is homoeomerous just if for any duration D which falls wholly within a duration during which Φ-ing by x goes on, D is a duration over which it is true that x Φ-ed. This detail need not detain us here, however.

are led to conclude that she must only count as running over certain sub-periods of time *in virtue of what was going on in surrounding intervals*. In isolation a single support phase isn't running; nonetheless, Radcliffe may be running during a time in which she only accomplishes a single support phase in virtue of what she accomplishes in surrounding times.[237]

The response offered to Fara above then turns out to be intimately connected to the denial of homoeomery where the relevant activity is experiencing change. Accounting for our experience of constant motion in the face of Fara's challenge has, in effect, led us to deny that change experience is homoeomerous down to very small temporal parts. Given the anhomoeomery of other processes this should perhaps not be an unexpected discovery about experience. However, as I argued above, there are other, quite specific reasons for thinking that experience of change could not be homoeomerous, reasons I now briefly recapitulate.

10. Temporal Minima

Firstly, change experience cannot be homoeomerous down to *instants*. That was the conclusion of Chapter Six where I argued that since change takes time, our experience of change must itself unfold over time. It follows from this that we should reject the idea that change-experience can be broken down into instants. We cannot experience change at an instant – though we may be in the midst of change experience at some instant and so be *experiencing* change then.[238]

This armchair argument establishes that there are experiential temporal minima in the relevant cases. Change experience is comprised of events or processes with *some* duration, and properties of change experience attach to those events or processes and not to instants. However, so far that argument only establishes that change experience is not

---

[237] On the closely related notion of dissectivity see Goodman 1977: 28, Armstrong 1961, and especially Hilbert 1987.

[238] The same applies to our experience of sounds. We experience sounds; any experience of a sound is an experience of a sound-filled period; the temporal structure of our experience matches the temporal structure apparently presented; thus, any sound experience itself will take place over time and so sound-experience cannot be broken down into instants. For all that, we can be *hearing* a sound at an instant if that instant falls during a period over which we hear a sound.

homoeomerous down to instants. For all that has been said it might be homoeomerous down to arbitrarily small temporal intervals. Armchair considerations will not help us here for the length of experiential temporal minima is plausibly an empirical matter. However, as we saw in the discussion of Dennett in Chapter Five, there are empirical reasons to think that experience cannot be broken down below quite significant durations. Thinking about these empirical findings also returns us to the idea that whether we experience some aspect of the environment depends crucially on the nature of *stretches* of experience.

Recall earlier discussion of visual masking experiments. There I argued that if we hold that experience is significantly anhomoeomerous, we can offer an attractive, realist account of masking phenomena which does not appeal to problematic delays in consciousness. The account notes that if experience does not break down into very short independent experiences, then there is no reason to assume that the experiential presence of a mask subsequent to the first, stimulus disc is irrelevant to answering the question whether the disc is perceived or not. So whilst it is true that if the masking ring hadn't intervened, the disc would have been reported, this has no bearing on the case where the ring was present. 'Has one seen the disc at *t*?' (where *t* is a moment briefly after the first stimulus presentation) is not a question that one can answer independently of one's experience during the surrounding period of time. To assume that it must be is to insist that experience is analysable down to very short periods. Certainly, the disc *would* have been seen had it not been for the mask. But it does not follow that either the disc *was* seen in the masking trial (cf. Dennett's Orwellianism) or, alternatively, that there is a slack in consciousness (cf. Dennett's Stalinism). This disjunction is only exhaustive if we assume a picture of experience analysable into independent instants or experiences of very short duration. On that picture, we cannot appeal to the nature of subsequent experience to distinguish the non-masked case from the masked case unless there is sufficient delay in consciousness for unconscious registering of the second stimulus to come into play.

Rejecting that picture, however, we can simply allow that in one case the first stimulus is seen and in the other not. As I put it above: different events have different perceptual appearances for us. The perceptual appearances of events are constrained by the fact that our experience itself has a certain coarse-grained temporal structure. These constraints

mean that we must think of the appearances of events as attaching first and foremost to durations above a certain minimum temporal threshold, in other words to temporally extended events rather than instantaneous or near instantaneous time-slices. In the masking case we should consider the events, (i) target followed by mask, and (ii) target followed by no mask, as different events with different appearances. What masking experiments show is that the appearances of sub-events of an event whole can contrast strikingly with the appearance of one of the sub-events presented on its own. Moreover, what masking suggests is that the temporal minima can be of the order of hundreds of milliseconds.

To return to the case of experienced change. Consider a change which takes place over a very short interval of time (e.g., Fara's case of the second-hand moving for a 1/36s ≈ 28ms period). What the discussion of masking suggests is that what happens in the subsequent period matters to our experience of this change. That fits well with the response given to Fara above since the idea here is that if the 28ms change is part of a larger change (i.e., the subsequent period involves more change) we may well see something different (indeed: *change* throughout the period) to what would be seen if the change were not a part of any larger change (quite possibly *no* change!). Moreover, the rejection of homoeomery makes room for the denial that we see the 28ms sub-changes of the changes which we perceive going on over longer periods. Rather we may simply be *experiencing* larger changes *as such* throughout these very short periods.

11. Temporal Limits

We have seen how to retain the intuitive explanation of slow motion, hour-hand cases whilst allowing for the perception of constant motion, hence solving Fara's puzzle. We do so by showing why Homogeneity fails in precisely those cases where we have genuine perception of motion. In addition we explain why failures of Homogeneity do not commit us to perceiving the tiniest of sub-changes of perceived changes even though we are experiencing constant change throughout the relevant periods.

The resulting picture is attractive given its ability to respond to Fara's worry. But it is also motivated independently since it rests on the denial that experience is homoeomerous down to instants or small durations. That denial has independent motivation arising from

the relation between experience itself as a temporal phenomenon and the temporal phenomena experienced, and further from the ability it provides to account for a variety of otherwise puzzling psychological phenomena.

In this and the following section I want to address two final concerns. Responding to them provides the opportunity to bring together various lines of thought developed in the thesis. The first worry can be put like this. The standard explanation of slow-motion, hour-hand cases, recall, is that there are changes of position that are too slight to be perceived, the change in position of an hour-hand over the course of a few seconds being just such a change. This explanation invites the question: But what if you're looking at the hour-hand for longer than a few seconds – say for an hour? Here the change in position of the hour-hand *isn't* too small to be perceived, yet despite this we don't see the hour-hand moving over that longer period of time. How is this to be explained?[239]

The explanation of why we don't see motion when we watch an hour-hand for an hour begins with the thought that the basic units of time perceived are of a certain duration. This duration is the key thing philosophers are alive to when they talk about the specious present or retentional aspects of experience.[240] One way of thinking about these durations is as periods over which we explain the properties of sub-parts in terms of the properties of the whole duration and not vice-versa. That is, periods where the explanatory direction runs from temporal whole to temporal parts. We have seen just this idea at work in the discussion above. We can be experiencing motion over some period in virtue of experiencing motion throughout some encompassing interval; we may fail to see some stimulus in virtue of the temporally surrounding experiential context.

The specious present does not last forever, however; nor do we retain all past experience in our current awareness (whatever Bergson may occasionally suggest). In other words, it is not just that the basic units of time perceived must be of a certain duration – there is

[239] Daniel Morgan and Matthew Soteriou both raised this concern with me independently.
[240] Cf. Russell 1948: 226: "'Pace perceiv'd' is only possible when the motion is so rapid that, though the beginning and end are noticeably different, the lapse of time is so short that both are parts of one sensation [i.e., fall within one specious present]". Conversely, "suppose that we are watching a chameleon gradually changing. We may be quite unable to 'see' a process of change, and yet able to know that, after a time, a change has taken place. This will occur if, supposing A and B to be the shades at the beginning and end of a specious present, A and B are indistinguishable, while A recollected is distinguishable from C when C occurs" (1927: 281).

also a *limit* to how long they can be. Beyond such a limit the explanatory direction is reversed; over longer time periods we explain our extended experience in terms of the various processes and events it is made up of. As a toy-model we can think of these basic durations of experience as windows and of there being limits to how wide or narrow windows can be. On the model we can fix the window width as 300ms, say. With the idea of a window in play we can say that one is only able to perceive change/movement if the clock-hand moves a sufficient distance within a 300ms window. In the hour-hand cases, the hand does not do this. We are well outside of such a window before sufficient movement takes place.

Beyond those windows the traditional thought that a succession of experience is not itself an experience of succession is correct.[241] Thus, in the hour-hand case simply having a course of experience which changes over some long period does not put one in a position to experience the succession itself. Put another way, in these long timescale cases we explain the course of experience in terms of its parts. None of those parts are of changes large enough to be perceived, so no change is perceived over the whole course of experience. In these cases the idea that sub-changes may be perceived in virtue of perceiving larger changes is not in play.

The idea that we could perceive sub-changes as parts of larger changes is principally a negative claim, viz., that just because we cannot see a change on its own, we should not assume that we cannot see it as a sub-change of a larger perceived change. That negative claim does not entail that we see any old change so long as it is part of a larger perceived change. As discussed, the very finest degrees of perceived changes are not plausibly picked up on. Thinking about hour-hand cases brings out a second important feature of the picture. This is that it is only sub-changes of changes falling within the basic, window duration which are such that they might be perceived in virtue of perceiving some larger change. For only in these cases does the explanatory direction run from temporal whole to parts. Thus, the appearance, and indeed presence, of events and processes which take place within windows depends upon what is being experienced over the duration of the whole window. Outside of windows, no such move is available.

---

[241] See, for example, James 1890: 629 and Husserl 1964: 31 discussed in Chapter Six.

To illustrate we can consider what would happen to us if our temporal window were suddenly to be altered. Imagine that instead of a 300ms window, one suddenly started experiencing with a 30ms window (if you like, a specious present of 30ms). Assuming that one's powers of discrimination remained fixed, a consequence would be that one could only perceive change/movement if a clock-hand moved a sufficient distance within that 30ms window. On many clocks, it is plausible that a second-hand does *not* move far enough within 30ms for any change to be detected. Thus, the effect of the reduction in window-width will be that second-hand experience becomes like hour-hand experience. We will see *that* the second-hand has moved, but cease to see its movement. Conversely, if our window suddenly expanded to 30s, we would begin to see the minute-hand's movement in just the way that we see the second-hand's movement.

12. Determinability and Resolution

A second concern with the account above is with its appeal to the idea that over very brief durations we simply experience motion *as such*. In itself this is not problematic. We should be wary of *assuming* that just because we experience motion over some short period that we must have experienced motion in the way that we normally experience it at longer timescales, i.e., by experiencing various sub-parts of the motion. This cannot plausibly be true all the way down. Insofar then as experience presents us with temporal structure and with properties and events logically related to time, it presents us with *determinable* not determinate structure and properties.

This general idea is rather common in representationalist accounts of experience alive to the limited discriminative power of ordinary vision. For example, the claim that we see determinables without seeing determinates is something representationalists will often claim in responding to alleged counter-examples such as blurry vision. Thus, Michael Tye argues that if you stare at a serrated stamp with 20/20 vision your experience will represent the stamp's relatively precise shape. However, if you take off your glasses, you will merely represent the squarish-ness of the stamp without representing the relatively determinate serrated shape that is the stamp's particular way of being squarish.[242] Given

[242] This example is from Pace 2007: 334 which I follow in my interpretation of Tye 2002: 149.

that determinable shape properties seem plausible content attributions, I see no particular reason for thinking that determinable temporal properties can't be invoked.

Note that the idea of determinable content in the absence of determinate content is not only available to the representationalist. As Pace comments,

> It is worth noting that a naïve realist could adapt Tye's solution to the problem of blurred vision by appealing to determinable spatial properties. For example, a naïve realist might claim that when one's vision is out of focus one is directly aware of the stamp and its squarishness (though not aware of its more determinate shape). (2007: 351)[243]

It must be acknowledged that such a move conflicts with the most extreme of 'naïve' intuitions about experience. On one picture, the naïve realist thinks of experience as actualizing the properties of objects and events in the world. If one thinks that all temporal properties in the world are determinately instantiated, then we must give up the claim that these determinate properties are literally actualized. Be this as it may, Pace is surely right that the core naïve realist idea can allow for our being related to determinable properties. Determinable properties are nonetheless properties of objects in the world and as such seem appropriate perceptual relata.[244]

In the context of the naïve picture outlined in Part One, however, there is a more striking consequence of determinable content that needs considering, one which links to the Dennettian idea of indeterminate perceptual content. On the naïve picture, recall, the temporal structure of experience itself mirrors the temporal structure of the apparent objects of experience. In particular, given Self-Intimation and Temporal Transparency, there can be no more to the experiential temporal structure of experience itself than is found in the structure of its objects. Thus, if the objects of experience have a determinable temporal structure so too does experience itself. As a result, what this discussion reveals is that although experience is structured in ordinary, physical time, its

---

[243] Note that Pace does not endorse either naïve or representationalist theory on this point.
[244] Some deny this claim. Thus Gillett and Rives hold that "the most reasonable position is that the world contains absolute determinate properties, but no determinable properties" (2005: 501). See also Armstrong 1978: Ch.22. Johnson is primarily focused on predicates rather than metaphysical issues about properties, but he also makes remarks which are suggestive in this regard, see his 1940: 175-6. Such a view does seem to me in tension with naïve realism/relational views of perceptual experience in this context.

experiential temporal structure is not as fine grained as ordinary, physical time. In other words, reflection on the nature of our experience reveals that although experience has gross temporal structure, there are no determinate facts about its fine-grained structure. Experience has *determinable* temporal structure but lacks *determinate* temporal structure. The structure of experience itself is essentially determinable. This immediately raises a worry for it has long been taken to be absolutely obvious that determinable properties can only be instantiated where some determinate property of that determinable is also instantiated.

Such a claim goes back at least to W.E. Johnson (1921) who provided the first systematic discussion of the determinable/determinate distinction in those terms. Thus, Johnson writes at the end of Ch.11, Part I of his *Logic*,

> The practical impossibility of literally determinate characterisation must be contrasted with the universally adopted postulate that the characters of things which we can only characterise more or less indeterminately, are, in actual fact, absolutely determinate.

As Sanford (2007) comments, "In saying it is a postulate, Johnson does not mean we merely assume it in order to deduce its consequences. He means rather that it is both obviously true and cannot be inferred from truths that are even more obvious."

Likewise, in his influential discussion of determinables and determinates, Armstrong lists "five features of the relation between determinable and determinate" (1997: 48). The first is Johnson's postulate.

> If an ordinary particular has a determinable property, then it is entailed that it has some determinate property, some more particular length or mass or colour, right down to the absolutely determinate lengths, masses or absolutely precise shades of colour. (ibid.)[245]

Finally, consider Funkhouser's recent paper on the determinable/determinate relation which commences with a section entitled 'Criteria for a Successful Analysis' (2006: §1).

---

[245] See also Gillett and Rives who endorse this principle of Johnson's and Armstrong's – "An individual satisfies a determinable predicate only if it satisfies some or other determinate predicate" (2005: 485) – before going on to argue for the much stronger, "pessimistic view, which rejects the existence of determinable properties" as default.

There he lists a series of "truisms about determinables and their determinates" which any "successful analysis of this relation should accord with" (ibid.: 548). Amongst these 'truisms' is the following claim.

> An object instantiating a determinable must also instantiate some determinate under that determinable. Coloured objects must be red or yellow or blue, etc. No object is merely coloured simpliciter. (ibid.: 549)

Call this principle *The Determinate Instantiation Principle* (DIP). If the principle applies universally, then the picture of experience defended in this thesis looks imperilled.

One reply to the DIP objection seeks refuge in an identity theory. The identity theorist can claim that the physical events or processes with which experience is identical *do* have determinate temporal properties. What is determinable is the temporal structure that is experiential structure, i.e., the temporal structure that contributes to what it is like, subjectively, to be the subject of the experience. It seems to me that this response is perfectly reasonable. However, it relies on the truth of some form of identity theory. In that sense, the picture I have defended appears to be hostage to a substantial philosophical theory, the truth of which one might reasonably be sceptical about.[246]

The only other option is to reject DIP. However, it must be acknowledged that DIP is highly intuitive. Consider: Could a person be tall (in some context) but lack a determinate height? Could a ball be flying through the air at some velocity but not at any determinate velocity? As Funkhouser suggests, could an object simply be coloured and yet not have any particular colour? Could an object be heavy but not have any determinate mass? These questions appear merely rhetorical. The answers all seem to be: *Obviously not*. One source of resistance to the picture of experience sketched above may well be precisely this commitment, viz., its commitment to experience unfolding over time but in no determinate fashion at very short-time scales.[247]

---

[246] For scepticism see, for example, Hornsby 1980-1, 1986.

[247] Note that it is not clear that it is right to say experience has a determinable temporal structure if that suggests that there are a number of determinate structures that experience *might* have and where the modality is nomological or metaphysical possibility. To this thought the naïve theorist can retort that human experience (as it is) is not the sort of process that *could* have a more determinate structure. Any experience has just the structure it does have, it just turns out that this structure lacks the fineness of grain

If DIP is strictly false, it will not do *simply* to dismiss it. We must explain its appeal either by restricting the principle or explaining why experience is a special case. Thus, though it is (perhaps) a central lesson of quantum physics that DIP fails at the Planck scale, pointing to that fact alone does not suffice to diagnose the force of DIP in relation to experience. Hence, whilst Funkhouser (2006: 566, fn.2) concedes, "quantum indeterminacy might allow for determinables with indeterminate determinates," he considers that "even granting quantum indeterminacy, in normal cases [the] criterion … still holds". Moreover, he points out that since "quantum indeterminacy is probabilistically qualified" we could rescue the criterion by amending it to: "every object instantiating a determinable also instantiates certain determinates to certain probabilities". Thirty milliseconds is hardly at the Plank scale; nor is any probabilistic treatment appropriate in this context. Thus, noting the alleged failure of DIP at the quantum level is not enough.

Sometimes DIP is explicitly restricted in its scope to physical objects. Thus, in earlier work, Armstrong makes the following claim.

> A physical object is determinate in all respects, it has a perfectly precise colour, temperature, size, etc. It makes no sense to say that a physical object is light-blue in colour, but is no definite shade of light blue. (1961: 59)

And he attempts to use this claim to refute phenomenalism on the grounds that "[p]hysical objects, which are determinate, cannot be constructions out of indeterminate sense-impressions" (1961: 58). Clearly the implication here is that sense-impressions are indeterminate. Let us agree with Armstrong that DIP does apply to physical objects (and events). Certainly that concession goes some way to account for the appeal of the general principle. Physicalism is plausible at least to the extent that it is plausible that fixing all the physical facts suffices to fix all the facts. Thus, the determinateness of the physical facts is the determinateness of the facts on which all other facts about the world supervene.[248]

---

we might have assumed it did and that physical time does have. Whilst this may be true, it seems at least that there are, conceptually speaking, determinate ways experience might have been.

[248] Dennett's 'Real patterns' (1991b) provides a model for thinking of emergent determinable facts (patterns) in a determinate world (the life world). Here all life world facts supervene on determinate facts.

Can we say anything that would motivate the failure of DIP in the case of experience itself? It seems to me that we already have. With ordinary physical processes there is no reason to think that, just because a process is temporally structured in some particular way, that we are in a position to know precisely what temporal structure it has. For example, if Paula runs for two hours, it may simply be impossible for us to know *precisely* how the relative speed of her first hour of running compared to that of her second. This, I take it, is the force of Johnson's remark above concerning "the practical impossibility of literally determinate characterisation". Nonetheless, this in itself is no reason to think that there *is* no answer to the question as to those relative speeds. Indeed, there may be powerful considerations for thinking that there must be such an answer, albeit beyond our ken.

On the other hand, if we accept Self-Intimation, this kind of picture cannot apply to experiential processes *in their experiential aspect*. Experiential properties are just those which we are in a position to know about. Thus, the nature of experience itself (unless physically realised or otherwise possessed of a non-experiential aspect) cannot outstrip our powers of self-knowledge. These, as revealed by the combination of Temporal Transparency with the finitude of our perceptual powers, are not limitless. Consequently experience itself must be determinable and DIP must fail in the experiential realm. The contrast Johnson tries to draw between the practical impossibility of determinate characterisation and actual determinacy fails in the experiential case. In the subjective world, the limits of determination are the limits of being.[249]

Bennett (2004: 171-2) insists not just that "our inner lives are temporally ordered" but that they are ordered "down to the finest detail". The failure of DIP should lead us to question this. In particular, the failure of DIP is closely related to the fact (if I am right) that experiential order is not *total* where experience exhibits a total order just if for all experiences/parts of experience in a given stream of consciousness, $e_i \leq e_j$ or $e_j \leq e_i$ where $\leq$ is the relation of precedence. Two cases: (i) you hear a flash and a bang, they do not appear simultaneous but you are in no position to tell which came first; (ii) you see

[249] One is reminded of the original version of the passage Armstrong quotes from Aristotle as 'A Warning' at the start of his 1997: "It is the mark of an educated man to look for precision in each class of things just so far as the nature of the subject admits." Armstrong has 'certainty' substituted for Aristotle's 'precision'.

motion between two points but cannot tell which direction the motion was in.[250] Only if we presuppose DIP will we assume that there must be a matter of fact concerning which was heard first, flash or bang or which direction the apparent motion was in. Without that assumption we can simply deny that either one was heard first or that the apparent motion was in one direction or other. One simply perceived an asynchronous flash and bang as such; one simply saw motion as such.[251]

Appeal to determinable structure arising from Self-Intimation allows us to respond to a recent puzzle raised by Lee (2007) which draws on claims from relativistic physics to argue against our inclination "to believe in a closer connection between the temporal layout of experience and its phenomenology than really exists" (367). This conclusion is, of course, diametrically opposed to a core claim of this thesis, namely, that the structure of experience in time is experiential structure which matches the temporal structure of what is apparently experienced. This is not the place to engage in detail with Lee's provocative and interesting argument. However, even leaving aside concerns about Lee's background "classical" conception of the stream of consciousness (e.g., 349) and its determination by a parallel sequence of momentary physical states, it is worth noting that Lee's conclusion rests on the crucial premise that "extremely fine-grained differences in the sequencing" of one's stream of consciousness – differences of no more than one nanosecond – "are still differences" (363). Insofar as such differences in sequencing pose a problem to our intuitive picture, they must be thought of as *phenomenological* differences. Differences in the realisation of the stream which have no impact on what it is like to undergo the stream are by-the-by. Yet it is utterly implausible that our introspective powers reveal so fine a grain. Thus, given Self-Intimation, there is every reason to deny the key premise. If we do so, we can happily exploit the determinability of experience just discussed to insist that experiential sequence is frame-independent since at phenomenologically relevant timescales (tens of milliseconds as opposed to nanoseconds) distinct experiences/parts of experience will all be time-like separated (cf. ibid.: 364).[252]

---

[250] See, for example, Pöppel 1978 and Hirsh and Sherrick 1961, both discussed by Hoerl 1998: 162.

[251] Bennett adds: "We cannot think our way down to a level where time does not apply, because no parts of our experience, however small or odd, lie outside time" (ibid.: 172). Quite right: no parts of our experience lie *outside* time. However the fact that our experience is temporally structured through and through is compatible with that structure not being fully determinate.

[252] This does not address the important idea of frame-relative *dilation* of experience. Here I am inclined to agree with Lee that Relativity reveals that durations in the stream of consciousness are not absolute as we

13. Conclusions

At the beginning of Chapter Six, I quoted Goodman's acute remark, "We normally take experience in larger chunks, and if we try to pulverize it by focusing attention on particles within … we usually find ourselves puzzled and uncertain" (1977: 203). Focusing attention on very brief particles within the stream of consciousness, considered in independence from the stream from which they are abstracted, can now be seen to be the core mistake behind a whole series of puzzles about time and experience. Struggling with these puzzles under the yoke of that mistake leads philosophers to make highly counter-intuitive claims. It leads Fara to claim either that we experience unnoticeable changes, perhaps even strictly continuous ones, or that we do not experience constant change at all. It leads Dennett, at least on one reading, to a form of anti-realism about the temporal structure of our experience. And it leads Kant, Reid, Prichard and others to deny that, strictly, speaking we perceive succession and change fullstop.

We have seen now how insisting on the anhomoeomery of experience allows us to account for otherwise quite puzzling phenomena without abandoning any aspect of the natural picture of experience in time developed in Part One. Moreover, it also provides a positive framework within which to account for the nature of our temporal experience without recourse to the often baroque architectures of specious present or retention based theories.

---

might have thought, just as the durations of events perceived are not. Temporal Transparency is not endangered because metrical durations are not aspects of the content of experience and so amongst its experiential properties (Ch.1, §1.1).

Chapter Eight:

Seeing Movements

Billowing clouds
On the move
Before my umbrella-hat[253]

Though many aspects of our experience are logically connected to time, an overwhelmingly important aspect that humans have evolved to discriminate perceptually is the movement of objects in our environs. As Blake and Sekuler remark, "The experience of motion, in all its various forms, is so integral to daily activity that it is hard to imagine life without it" (2006: 317). Yet whilst motion-perception is evidently a fundamental aspect of our experience and in particular of our visual experience, philosophers and psychologists have long worried as to how it is possible. Movement takes place over time, hence philosophical concern over how motion perception is possible can certainly be an instance of the general worries about our perception of temporal properties discussed above. However, there seem to be relatively orthogonal concerns specific to understanding how visual motion perception is possible which have independently long exercised philosophers and psychologists.[254] It is these that this chapter takes up.

1. Prichard's 'Seeing Movements'

An extremely forceful expression of philosophical concern local to movement perception can be found in a provocative but neglected paper of Harold Prichard's called 'Seeing Movements'.[255] There Prichard argues that we can *never*, strictly speaking, see bodies move. This conclusion may at first sight seem too preposterous to be taken seriously.

[253] Kobayashi Issa (1805) D.G. Lanoue (trans.).
[254] Classic early discussions include Helmholtz 1910 and James 1890: vol. 2, 171f..
[255] Prichard 1950b. All references to Prichard in this chapter are to this paper.

Nevertheless, his argument deserves our attention because responding to it forces us to expose a deeply rooted but misguided way of thinking about vision. I come to the argument shortly. However, Prichard goes on to argue for an even more radical conclusion on the basis of his sceptical view of motion perception. As he puts it rather cryptically,

> [F]or the same reason it cannot be true that in the process so called we really see *bodies* – for we cannot separate the 'bodily-ness' of what is seen from the moving of what is seen. (43)[256]

It is interesting briefly to explore this spreading step. In particular, how should we understand the inseparability claim here? The remark might be read as the claim that we cannot see a moving body unless we see it as moving. There seems little to recommend this principle. There are many illusions which involve moving bodies appearing stationary (and vice-versa).[257] More plausibly, Prichard is insisting that we cannot see bodies unless we *can* (or, *in general, do*) correctly see them as moving. Though perhaps more plausible, the phenomenon of akinetopsia (motion blindness) forces us to abandon this view.

Localised brain damage can selectively disrupt and (at least in theory) outright destroy movement vision, a condition know as cerebral akinetopsia. Akinetopsic patients are extremely rare. Only a handful of definitive cases are documented.[258] In all, some very limited movement vision is spared. Thus, Zihl et al. (1991) describe the first patient with akinetopsia to be rigorously studied as follows,

> The patient [L.M.] had completely lost movement vision in depth, and could only discriminate between a stationary and a moving target in the periphery of her otherwise intact visual fields. In the central visual field some movement vision was spared for horizontally and vertically travelling targets, provided that target velocity did not exceed 10 deg/s. She possessed neither visual movement after-effects nor apparent (phi) visual movement. Visually guided eye and finger movements were also impaired. In contrast,

---

[256] Prichard is no doubt terse and cryptic here because he is already convinced of this view on the basis of more traditional arguments.
[257] See, for example, the optokinetic drum illusion discussed below.
[258] For a review of the relevant literature see Zeki 1991.

> her movement perception elicited by acoustic and tactile stimuli was not impaired. (2235)[259]

It is now widely accepted that "movement vision represents a separate and genuine visual function" (ibid.). Motion processing is not localised to just one area, hence the residual motion vision in patients such as L.M. Nevertheless, if the basic modular hypothesis is right, it is reasonable to postulate pure cases of akinetopsia in which patients lose movement vision completely.

As Pelak and Hoyt note, there are no standard psychometric tools with which to assess visual motion. However, in the current context, what is most interesting are the descriptions of the phenomenologies which accompany the condition in question. For example consider the following reports.

> [A patient reported by Pötzl and Redlich in 1911] described her perceptual experience of a moving target as if the visual stimulus remained stationary but appeared at different successive positions. (Zihl et al. 1983: 314)

> The visual disorder complained of by the patient [L.M.] was a loss of movement vision in all three dimensions. She had difficulty, for example, in pouring tea or coffee into a cup because the fluid appeared to be frozen, like a glacier. In addition, she could not stop pouring at the right time since she was unable to perceive the movement in the cup (or a pot) when the fluid rose. (ibid.: 315)

> [A former hunter] gave several examples of the difficulty he had with motion vision. For instance, while hunting, he was unable to see his dog move toward him or notice game approaching in the forest. Instead, his dog or game would appear in one location and then another, without any movement being seen between the two locations. … If he was driving a car, other cars in motion appeared to be stationary, and he occasionally perceived cars or objects to be moving when they were motionless. (Pelak and Hoyt 2005: 138)

The phenomenologies on record here refute the second reading of Prichard's claim. For these patients, moving stimuli *in general* (cars, targets and liquids alike) appear stationary.

---

[259] See also their pioneering 1983 study.

The objects are seen and yet their motion is not.[260] Since there is no reason not to take the phenomenologies at face value, we have a direct counter-example to the idea that we cannot in general see moving objects unless we can see them as moving.

However, Prichard's comment is strongly reminiscent of the following notorious contention of Locke's.

> Primary qualities of bodies [viz., solidity, extension, figure, motion or rest, and number] … are utterly inseparable from the body, in what state soever it be … and such as sense constantly finds in every particle of matter which has bulk enough to be perceived; and the mind finds inseparable from every particle of matter, though less than to make itself singly be perceived by our senses (1975: Bk II, Ch VIII, §9)

Locke's claim is weaker than Prichard's, for Locke only holds that one cannot see a body unless one also sees it in motion *or at rest*. This is enough for Prichard to generalize his argument since he is implicitly committed to the claim that we cannot perceive stationarity given the grounds on which he denies that we can perceive movement.[261] Moreover, there is greater plausibility to the Lockean claim; it is hard to imagine cases where one sees a body and yet fails to see it either as moving or as rest. Akinetopsia does not obviously provide such examples. However, there are cases. For instance, imagine catching a glimpse of something through the gaps in the floorboards. The light is poor and the gaps small such that one cannot discern what it is you see or what it is doing (if anything). All you can see is that there is something down there. I do not see why we should not say here that you *see* an object and yet do not see it as either in motion or at rest.[262] Thus, even if we felt compelled to accept Prichard's argument concerning motion,

---

[260] Some cases of akinetopsia *do* fit with what Prichard would have to say. Thus, a patient suffering from Alzheimer's disease is described as follows: "when an object began to move, it would 'disappear.' … He reported difficulty watching television when people or objects were moving, and this was especially true while viewing sporting events. For example, football players on the television screen would seem to 'disappear' as soon as they started to run. While watching television movies or other action-filled television shows, he frequently commented to his wife that he could not 'see' anything going on. But he could watch the news being broadcast (during which no movement or significant action occurred) without difficulty" (Pelak and Hoyt 2005: 139-40).

[261] In effect, there is no reason not to extend the neutrality claim (see below) to include neutrality between cases where both observer and object are stationary and cases where both observer and object are moving yet not moving relative to each other.

[262] Indeed, one might well use the same example to deny Locke's other inseparability claims. Another case: imagine seeing an object on the distant horizon, too far away to determine its colour, its shape or whether it is moving but close enough to see that *something* is there.

we would not obviously be forced to accept any more general conclusion. I now turn to the argument against motion perception.

2. The Argument Against Seeing Motion

Prichard's argument can be condensed as follows.

(1) Seeing a body move is quite distinct from seeing *that* a body is moving.
(2) Movement is absolute.
(3) However, "two bodies moving differently or one body moving and the other at rest may inevitably through a given time present precisely the same appearance" (42).
(4) Thus, "what we call seeing a body move cannot really be seeing a body *move*" (ibid.).

Premise (1) is firmly grounded in our visual phenomenology: witness the standard example of the difference between seeing a second hand move round the clock, in contrast to seeing *that* the hour hand has moved.[263] More generally, it is plausible to think that 'seeing that' locutions track our epistemic standing, i.e., what we are in a position to know about the world in virtue of our perceptual experience. Thus, given a suitable epistemic context, one can, for example, see that a train is moving without seeing the train at all, just its smoke, or the rustling leaves, or the station monitor.

Premise (2) may strike one as obviously false, however. Whatever may have been true in 1925, few scientifically informed people these days *do* think, let alone assume, that movement is absolute. Nevertheless, the premise can easily be amended. For, even if we abandon the idea of an absolute frame of reference, we can fix a frame of reference with respect to which we can treat motion as 'absolute' and distinguish one object moving away from another 'stationary' object and vice-versa. This is, in practice, how

[263] See James (ibid.). Note that an inferential account of movement perception which reduces motion perception to judging, or being positioned to judge, that motion is occurring was historically the only game in town despite its obvious inadequacy. Indeed, Kellman comments that "To describe this as the prevailing view would be an understatement since until recent years an alternative view had scarcely been articulated" (1995: 351).

astronomers proceed in considering planetary motion. Roughly speaking, an approximately inertial frame is fixed with reference to the 'fixed stars,' whose centre is the centre of mass of the solar system. In our own case, an obvious choice of reference frame is fixed by the earth's surface and its gravitational field. With respect to such a frame we can distinguish 'absolute' from 'merely relative' motion. Moreover, it is very plausible that we *do* visually discriminate 'absolute' (as opposed to merely 'relative' motion) – this is exactly how things seem to us and how psychophysicists talk (e.g., Wertheim 1999). Indeed, this is hardly surprising: humans have evolved to discriminate motion in an earth-bound context, for such detection is what matters to our sublunary survival and prosperity. Consequently, (2) can be amended to,

(2*) Movement is 'absolute'.[264]

So amended, the argument still goes through since premise (3) remains unaffected.

The only remaining response is to deny that the same appearance *is* inevitably presented by the situations which differ only in that in (a) an object moves and I am stationary and in (b) the object is stationary and I move relative to it. Call the claim to be denied the *neutrality claim*. Two tasks need accomplishing: we must diagnose the attraction that the claim has had for Prichard and others. We must then develop a positive account which rejects it.[265]

3. The Video Model of Vision

It is common to think of vision as simply a way in which we are sensitive to light and of the visual system as a sophisticated light-sensitive mechanism. These claims can sound like utterly basic scientific facts. Nevertheless, they are false and appreciating this frees us from the trap of the neutrality claim. Vision is not a way in which we are sensitive

[264] I occasionally leave off scare-quotes in what follows.

[265] Note how the same kind of argument can be given with respect to the gravitational vertical and yet it is clear that we see things as oriented with respect to the gravitational vertical. Clément and Eckardt produce some fascinating results which show, in their words, that "gravitational reference plays a significant role in … 'visual' illusions" (2005: 911). The only thing to disagree with here is the use of scare quotes which implicitly allows for an analogous neutrality claim concerning the 'strictly visual'. See also Rock's classic 1956.

merely, or even typically, to light, it is a way in which we are sensitive to *our surroundings*. Visual awareness is possible because our *retinas* are sensitive to light. But we must not conflate that basic scientific fact with a false claim about the nature of vision itself.

Encouraged by thinking of vision as simply as a sensitivity to light, we think of ourselves as equipped with something like a video camera.[266] One need not fall prey to a homunculus fallacy or endorse a traditional representationalism to think that an analogy between vision and video photography is appropriate. In particular, one can think of visual content as analogous to photographic content in the sense that just as the character of a video or photograph *is* purely determined by light registration on some sensitive film or receptor array,[267] the character of visual experience supervenes on light stimulation: no qualitative difference in visual experience without some difference in light stimulation.

I now argue that underlying the intuitive pull of the neutrality claim is the video model, itself grounded on the idea that vision is simply a way in which we are sensitive to light. Consider first this passage from a recent paper of Richard Price's where he argues that we should deny that objects look to the left or right of you.

> Call S_1 the situation in which object O, against a plain white background, moves from your left to your right, and you keep your head still. Call S_2 the situation in which O stays still, and you move your head from right to left. From a purely visual point of view, S_1 and S_2 seem qualitatively identical. The change in look of O in S_1 is identical with the change in look of O in S_2. Of course, you may be able know whether you are in S_1 or S_2 by proprioception, since proprioception may tell you whether you have moved your head. But if we focus on the purely visual aspects of S_1 and S_2 (to do this, we can imagine that your proprioceptive awareness has been numbed), it seems hard to deny that O's change in look in S_1 is identical with its change in look in S_2. (2005: 369-70)

The idea that we can consider the two situations from "a purely visual point of view" plays a crucial role in this argument. As Price readily acknowledges, it is only from this

---

[266] Before the video-camera was part of everyday life, a simple camera obscura could equally have served as a model for vision.

[267] The photographer Philip Greenspun *defines* photography as "the recording of light rays" (see http://www.photo.net/learn/making-photographs/light). One might take issue with this claim. However, for present purposes I see no reason to demur.

point of view that the two situations can even be claimed to seem identical. Yet we ought to ask what is meant here by "a purely visual point of view" and why the two situations are supposedly indistinguishable from it. If one tries the experiment, the object in S_1 is seen to move. In S_2, it is seen to remain stationary. Moreover, the two situations are distinguished in their final state by the fact that one's head is shunted to the side in S_2 unlike S_1.[268] *You* are in a different place. And one might think this underpinned a difference in how things *seem* to you visually. So we need to ask why it seems so obvious to Price that the situations are identical "from a purely visual point of view".

Similarly consider Prichard:

> 'When in the lift', it may be said, 'I may, according to circumstances, see the shaft going up or see the shaft stationary.' Here the physical conditions so far as seeing the shaft is concerned are the same; what they give rise to is certain appearances, which therefore must be the same, and therefore the difference between 'seeing the shaft going up' and 'seeing the shaft stationary' *must* solely lie in a difference of interpretation of identical appearances. (44)[269]

This passage reveals a picture of the physical conditions of vision appropriate to the way a camera records the world. There will be no difference in the recorded videos whether it is the lift that goes up or the camera that goes down. Watching the video, only contextual clues can be used to determine whether the camera was stationary or moving.[270] Appearances, like bits of film footage, are "determined by the position [and trajectory] of the point [of view/camera position] relatively to the body [seen/shot]" (42). If we think of vision this way, it will indeed be obscure how we could genuinely *see* 'absolute' motion. Yet we do, and so this must be the wrong way to think about vision. How should we think about it?

---

[268] As we shall see shortly, neck muscle feedback plays an important role in our visual awareness of motion and relative position.

[269] Prichard is not endorsing the conclusion (which he takes to be akin to a reductio of a certain natural, 'apprehension' view of perception). He *is* endorsing the principle that moves us there.

[270] This fact is made use of in special effects using a camera mounted on a dolly track, hydraulic arm or gyroscopically balanced SteadiCam. Of course, movement relative to obvious background features must be avoided; as Walton notes, "With carefully arranged camera movement, a stationary horse image on a changing background might depict the horse in motion, or a changing background might depict it as stationary" (2008: 163).

4. The Neuropsychology of Motion Perception

It has long been held "evident that perception of motion is normally visual-kinaesthetic" (Alexander 1914: 283).[271] However, there are various ways of understanding this claim. According to Alexander, on the basis of vision and our kinaesthetic sense, we can *tell* that objects are in motion. Visual experience positions us to know a disjunctive truth: either that the body we are seeing is moving or that we are, and then "it is the kinaesthetic element which tells us" which (ibid.). However, this is not an account of motion *perception*. As Prichard rightly insists, there is a great gulf between *judging* and *perceiving*.[272] Again: we need to explain how we *see* things move – not just how we see *that* they are moving.

Modern neuropsychology suggests another way of understanding the relation between proprioception and visual awareness. In particular, Wertheim (1999) proposes the following theory. Information from three sources combines into a "reference signal" which serves an evaluative function in relation to data drawn exclusively from the retinal image. This combined signal subserves motion perception. The three sources are as follows.

(a) An "efference copy signal" which carries information about oculomotor commands, used to determine eye velocity in relation to the head.
(b) A "vestibular afferent signal" (the vestibular system is a combination of the equilibrium system and the somato-sensory kinaesthetic system) which provides information related to the location and movement of the head in space.
(c) A "visual afferent signal" which carries information directly extracted from the retinal image flows.

Together, these inputs are used to determine 'absolute' eye-movement. If we assume that neural processing subserves perceptual consciousness by making accessible to us (i.e., manifest in perceptual consciousness) various features of the mind-independent

[271] Though contrast H.H. Price 1932: 255.
[272] This is the central point he makes against the sense-datum theorist, insisting that "to *see* something is just not to interpret or judge something else seen as that thing" (44).

empirical world,[273] we can understand the processing which combines the "reference signal" with a *purely* retinal data-stream as making accessible to us the motion properties of objects around us. Intuitively, the information in the reference stream allows the brain to 'factor-out' self-motion and so make manifest only 'absolute' motion.

Wertheim's theory provides *subpersonal* explanations of a range of phenomena. Wertheim's central example is the optokinetic drum illusion where a subject sits inside a large drum painted with vertical black and white stripes. The drum is rotated at sixty degrees per second around the seated observer. If the lights are switched off and then suddenly turned on, the subject first feels stationary and sees the drum moving as it is. Over a period of about six seconds the drum appears to slow down and eventually appears stationary; concurrently the subject experiences increasing vection – a visually induced sense of self-motion.[274] Here the *relative* motion *both actual and apparent* is the same throughout the six second period. Thus, the illusion provides a clear case (if we needed one) of phenomenological difference despite sameness of relative motion – *contra* Prichard's neutrality claim. Wertheim's theory predicts this result by appeal to the nature of the visual afferent component in the reference stream.[275] Other components are required to explain (and hence are partly evidenced by) the Filehne Illusion (where during short periods of smooth pursuit eye movement made in the dark across a small stationary stimulus, the stimulus is perceived as moving in the opposite direction to the eyes), centre-surround induced motion and movement after-effects.[276]

There is, importantly, also at least one case reported by Haarmeier et al. (1997) where a subject, RW, is reported as having suffered bilateral extrastriate cortex lesions apparently impairing the efference copy signal and leaving his visual world "grossly unstable unless the eyes are stationary" (Bruce et al. 2003: 264). Bruce et al. comment that "RW must also be living proof of the incorrectness of Gibson's assertion that optic array

---

[273] See, for example, Campbell 2002: 118-20 and Brewer 2004: 69-70.
[274] For details see Wertheim 1994.
[275] Only when a visual (i.e., retinal) pattern is large, has relatively low spatial frequency characteristics, and moves relatively slowly across the retinas for some extended period will stimuli generate 'optokinetic' signals. Thus, though drum rotation immediately generates a *retinal* signal, "vection develops only gradually, due to the low temporal bandpass characteristics of the gating mechanism in the optokinetic pathway. Therefore, a (visually induced) reference signal is not immediately present. Hence, initially the drum is correctly perceived as moving" (Wertheim 1994).
[276] On this last illusion, see Wertheim 1987. Of course, we are familiar with many such illusions in everyday life involving trains, lifts and fairground rides.

information alone is sufficient for active vision in the world. People [or rather their visual systems] really do take account of their own eye movements" (ibid.).[277]

5. The Simple View of Motion Perception

The cognitive neuroscience just sketched strongly supports the idea that kinaesthetic information is used in perceptual *processing* itself. This still leaves open the link postulated between proprioception and vision, for there are two ways of understanding the involvement of the reference signal in subserving conscious awareness of movement. One possibility is to insist that motion perception is an essentially *cross-modal* phenomenon. That is to say, we never *visually* perceive motion although we do perceive it. Our experience of motion is *sui generis*.[278]

This approach is unattractive for three related reasons: it is obscure, unmotivated and phenomenologically off-key. It is obscure because nothing has been said about the nature of this type of sensory awareness other than that it is *not* analysable into visual awareness and kinaesthetic awareness. It is unmotivated given the availability of the simple and unmysterious view I sketch in the next paragraph. Finally, it is phenomenologically off-key for, when one reflects upon one's experience of motion, it *does* seem that we genuinely *see* movement as well as feel it. One is not at all inclined to retract such claims and concede that strictly speaking one did not *see* motion but rather had some *sui generis* experience of it. Moreover, when one has done with reporting one's visual and sensational experience, one has done all that needs doing in reporting what it is like to perceive motion.

A much more attractive approach is just to say that we *see* movement. On this account, the reference stream is understood as subserving *visual* perception of movement. Call this the simple view. The view may look as if it has not discharged its explanatory burdens. But what needs explaining? Our experience as of a stationary object that we are moving relative to *seems* very different to our experience as of an object moving relative to us when we are stationary. Furthermore, we have a compelling neuropsychological account

[277] For Gibson's initially promising view see his 1966: 256.
[278] Cf. O'Callaghan 2007: Ch.11 on the audio-visual case.

of the mechanisms underpinning the difference. Surely the burden is on anyone who wishes to reject the simple account. Thus, we are brought back to the neutrality claim which stands in stark opposition to this simple view.

Firstly, note that, in order to elucidate the idea of "a purely visual point of view," Price asks us to consider how things would be if our proprioceptive awareness were numbed. Now, of course, on the simple account of motion perception, proprioceptive *information* is crucial to visual motion detection. Tampering with our proprioceptive system would interfere with our visual system. Thus, causally speaking, there is no straightforward way of separating out the two forms of awareness. If we 'numb' proprioceptive awareness, we remove information from Wertheim's reference stream and this will doubtless affect how things seem *visually*.

However, we could conceivably numb proprioception whilst maintaining the reference signal 'down-stream' in some way. But what reason is there now for thinking that so numbing proprioception would render Price's two situations, S_1 and S_2, indistinguishable? 'Absolute' motion is, *prima facie*, part of our purely *visual* phenomenology. The thought that proprioception must really be doing the work only arises because we are already convinced of something like the video model.

Similarly, consider Prichard's idea that the "physical conditions so far as seeing the shaft is concerned are the same" when you are stationary and the shaft moves, and vice-versa. This is simply false on Wertheim's model, for the physical conditions so far as seeing the shaft are concerned include information relating to the location and movement of one's body in space. Unlike a camera which does not register such information but merely light, subpersonal perceptual processes do make use of this information. Thus, there is good reason to think that such information can be employed to make manifest 'absolute' motion in such a way that my experience is not neutral in the way the video is.

What can be hard to see is how this information can add anything to an already seemingly complete piece of, as it were, camera footage. Certainly, no further visual *object* need be present – we do not need to see the earth's surface which partly defines our reference frame, for example, and nor would it help if we did. Whilst we do often see the relevant frame, this could only play a role in our *interpretation* according to the camera-

model. Such a worry betrays a failure properly to give up the video model. We must take seriously the idea of vision as a way in which absolute *motion* can be made manifest and so of the idea of an absolute motion appearance. We see things *as moving*. There is no better answer to the question, 'What is added to a motion neutral experience?' than, *'Motion!'* ('What are akinetopsics blind to?' *'Motion!'*)

We experience motion, then. We do not merely infer it. But such phenomenology – as Prichard rightly points out – is incompatible with thinking of vision simply as a light sensitive mechanism and hence with the video model of perceptual experience. Note that the point is not quite that information internal to the visual array is always insufficient to determine whether an object is moving or not. Information from the vestibular system and efferent ocular copy is typically relevant to what we see. However, purely visual information can sometimes be sufficient to determine whether an object is moving or not. What Prichard is misled by is the correct thought that if the information is contained in the visual array *and is extracted through normal vision*, this at most can put us in a position to *judge* that we are moving or are stationary. And judging is not seeing. That was premise (1). Thus, if my experience is akin to a piece of video footage, I will only be able to *judge* that it is the object moving and not me (i.e., the 'camera') on the basis of contextual information supplied by my visual experience. I have been urging that we perform a modus tollens and conclude that our experience is not like this.

In short then, in order to resist Prichard's radical conclusion that we do not see movement we need to deny the claim that visually speaking our experience is neutral across cases where relative motion between us and an object seems the same. This neutrality claim garners its support from the video model of visual appearances which arises when we think of the visual system as simply a light-sensitive mechanism. Tempting as it may be, it is a fundamental mistake to think of vision in such a way.

ACKNOWLEDGEMENTS

This thesis has accrued many specific debts. A version of Chapter Five was presented at a seminar in Warwick in June 2008. Thanks to the audience there for their usual hospitality, stimulating questions and suggestions. Chapter Six is a revised and expanded version of my paper, 'Perceiving Temporal Properties,' forthcoming in the *European Journal of Philosophy*. Early ancestors were presented in Oxford, Warwick, London and Bristol, and a revised version at a *Consciousness and Time* conference in Edinburgh and a departmental colloquium at UT Austin. I am grateful to the audiences on those occasions and in particular to Bill Brewer and Matthew Manning for their thoughtful and thought-provoking responses. I am also grateful to Robin Le Poidevin who in his capacity as *EJP* referee provided me with most helpful comments as well as a pre-print of some of his own recent work on the topic. John Campbell very kindly took the time to give me detailed comments on a draft whilst I was in Berkeley.

Chapter Seven began as a paper presented at the NYU/Columbia Graduate Philosophy Conference in March 2008 and the UT Austin Graduate Philosophy Conference in April 2008. Once again I received extremely helpful comments from the audiences there and especially my respective commentators John Morrison and Tomas Bogardus. Versions of both Chapters Six and Seven have also been put through the trial by fire that is the Vicious Circle – I am much indebted to all its members.

Apart from those specific occasions, I have benefitted from the kind attention of five supervisors and innumerable advisors. In my first term at UCL, sketches of material received great care and encouragement from Paul Snowdon. In my final year Mark Eli Kalderon helped me revise, prune, sharpen, qualify, clarify and immeasurably improve almost all my material. And in my last term, Chris Peacocke read a complete draft, offering hugely helpful and acute comments on myriad aspects of the argument. Matt Soteriou, though not officially a supervisor, must be counted as one. From the earliest days of working with him as an undergraduate in Magdalen he impressed upon me the importance of ontology in the philosophy of mind, and over the years his work and comments have illuminated issues time and again. Finally, Mike Martin occupied centre stage in helping me develop my views over the course of many rich and stimulating

supervisions and conversations. His influence extends well beyond my thinking about experience and time to my whole conception of how philosophy can be done.

Amongst my advisors special thanks to four friends, Geoff Lee – long may the fruitful disagreement continue! – Hemdat Lerman, Hanna Pickard and Maja Spener for many hours of coffee and detailed comments between them. Thanks also to John Bengson, Tim Crane, Amanda Cupples, James Genone, Anil Gomes, Susan Hurley, Gabriel Lakeman, Rory Madden, Elijah Millgram, Derek Parfit, Joe Perkins, Richard Price, John Schwenkler, James Stazicker, Scott Sturgeon, Lee Walters, Jonah Wilberg, and many others – you know who you are. Late in the day, Barry Dainton read and commented in detail on the entire manuscript, helping me clarify and correct many points in the argument. I'm most grateful for that kindness.

Finally, my great thanks to my excellent examiners, David Chalmers and Charles Travis. My future work will long continue to benefit from their many insightful questions and comments.

This thesis was made possible by the generous support of the AHRC and All Souls College, Oxford.

Everything was made possible by my parents.

I.B.P.
January, 2009

BIBLIOGRAPHY

Alexander, H.B. (1914) 'The Perception of Motion' *The Journal of Philosophy, Psychology and Scientific Methods* 11 (11) pp.281-290.

Alpern, M. (1953) 'Metacontrast' *Journal of the Optical Society of America* 43 pp.648-657.

Andersen, H. and Grush, R. (forthcoming) 'A Brief History of Time-Consciousness: Historical Precursors to James and Husserl' *Journal of the History of Philosophy*.

Ansorge, U. and Heumann, M. (2006) 'Shifts of visuospatial attention to invisible (metacontrast-masked) singletons: Clues from reaction times and event-related potential' *Advances in Cognitive Psychology* 2 (1) pp.61-76.

Anthony, M.V. (1992) 'The where and when of what?' *Behavioral and Brain Sciences* 15 pp.201-202.

Aristotle (1922) *On Coming-to-be and Passing-Away* [more commonly known as: *On Generation and Corruption*] H.H. Joachim (ed.) Oxford: Clarendon Press.

Armstrong, D.M. (1961) *Perception and the Physical World* Humanities Press.

Armstrong, D.M. (1978) *A Theory of Universals* Cambridge: CUP.

Armstrong, D.M. (1997) *A World of States of Affairs* Cambridge: CUP.

Aronson, J., Dietrich, E., and Way, E. (1992) 'Throwing the conscious baby out with the Cartesian bath water' *Behavioral and Brain Sciences* 15 pp.202-203.

Bennett, J. (2004) 'Time in Human Experience' *Philosophy* 79 pp.165-183.

Bergson, H. (1914) *Creative Evolution* A. Mitchell (trans.) London: MacMillan & co. Ltd.

Blake, R. and Sekuler, R. (2006) *Perception* McGraw-Hill.

Block, N. (1980) 'Troubles with Functionalism' in N. Block (ed.) *Readings in the Philosophy of Psychology* Cambridge, MA.: Harvard University Press pp. 268-306.

Block, N. (1992) 'Begging the Question Against Phenomenal Consciousness' *Behavioral and Brain Sciences* 15 pp.205-206.

Block, N. (1993) 'Review of *Consciousness Explained*' *The Journal of Philosophy* 90 (4) pp.181-193.

Block, N. (1995) 'On a confusion about a function of consciousness' *Behavioral and Brain Sciences* 18 (2) pp.227-47.

Block, N. (2007) 'Consciousness, accessibility, and the mesh between psychology and neuroscience' *Behavioral and Brain Sciences* 30 (5/6) pp.481-499.

Block, R. and Zakay, D. (2001) 'Retrospective and prospective timing: Memory, attention, and consciousness' in C. Hoerl and T. McCormack (eds.) *Time and memory: Issues in philosophy and psychology* Oxford: OUP pp.59-76.

Borges, J.L. (1964) 'A New Refutation of Time' in his *Labyrinths: Selected Stories and Other Writings* N.Y.: New Directions Pub. Co.

Bouhassira, D., Kern, D., Rouaud, J., Pelle-Lancien, E., and Morain F. (2005) 'Investigation of the paradoxical painful sensation ('illusion of pain') produced by a thermal grill' *Pain* 114 (1) pp.160-167.

Bregman, A. and Campbell, J. (1971) 'Primary auditory stream segregation and perception of order in rapid sequences of tones' *Journal of Experimental Psychology* 89 pp.244-249.

Brentano, F. (1973) *Psychology from and Empirical Standpoint* A.C. Rancurello, D.C. Terrell and L.L. McAlister (trans.) London: Routledge & Kegan Paul.

Brewer, B. (2004) 'Realism and the Nature of Perceptual Experience' *Philosophical Issues* 14 pp.61-77.

Broad, C.D. (1923) *Scientific Thought* London: Routledge & Kegan Paul.

Broad, C.D. (1929) *The Mind and its Place in Nature* London: Kegan Paul, Trench, Trübner & Co. Ltd.

Brough, J.B. (2000) 'Plastic Time: Time and the Visual Arts' in J.B. Brough and L. Embree (eds.) *The Many Faces of Time* Kluwer Academic Publishers pp.223-244.

Brown, S.W. (1995) 'Time, change, and motion: The effects of stimulus movement on temporal perception' *Perception and Psychophysics* 57 (1) pp.105-16.

Bruce, V., Green, P., and Georgeson, M. (2003) *Visual Perception* Psychology Press.

Burge, T. (2003) 'Perceptual Entitlement' *Philosophy and Phenomenological Research* 67 (3) pp. 503-548.

Burgess, J.A. (1990) 'Phenomenal Qualities and the Nontransitivity of Matching' *Australasian Journal of Philosophy* 68 (2) pp.206-220.

Byrne, A. (1997) 'Some Like It HOT: Consciousness And Higher-Order Thoughts' *Philosophical Studies* 86 pp.103-129.

Byrne, A. (2006) 'Color and the Mind-Body Problem' *Dialectica* 60 pp.223-44.

Byrne, A., Hilbert, D.R., and Siegel, S. (2007) 'Do we see more than we can access?' *Behavioral and Brain Sciences* 30 (5/6) pp.501-2.

Campbell, J. (2002) *Reference and consciousness* Oxford: OUP.

Carmen, T. (2007) 'Dennett on seeming' *Phenomenology and the Cognitive Sciences* 6 (1-2) pp.99-106.

Carnap, R. (1967) *The Logical Structure of the World* London: Routledge & Kegan Paul.

Carroll, L. (2001) *Alice's Adventures in Wonderland* Courier Dover Publications.

Carruthers, P. (2005a) *Consciousness: Essays from a Higher-Order Perspective* Oxford: OUP.

Carruthers, P. (2005b) 'Why the Question of Animal Consciousness Might Not Matter Very Much' *Philosophical Psychology* 18 (1) pp.83-102.

Chalmers, D.J. (1996) *The Conscious Mind* Oxford: OUP.

Chalmers, D.J. (2004) 'The Representational Character of Experience' in B. Leiter (ed.) *The Future for Philosophy* Oxford: OUP pp.153-81.

Chalmers, D.J. (2006) 'Perception and the Fall from Eden' in T. Gendler and J. Hawthorne (eds.) *Perceptual Experience* Oxford: OUP pp.49-125.

Clark, A. (1989) 'The particulate instantiation of homogeneous pink' *Synthese* 80 pp.277-304.

Clark, A. (1992) 'Experiential facts?' *Behavioral and Brain Sciences* 15 pp.207-208.

Clark, A. (2000) *A Theory of Sentience* Oxford: OUP.

Clément, G. and Eckardt, J. (2005) 'Influence of the gravitational vertical on geometric visual illusions' *Acta Astronautica* 56 pp.911-917.

Crane, T. (1998) 'Intentionality as the mark of the mental' in A. O'Hear (ed.) *Contemporary Issues in the Philosophy of Mind* (Royal Institute of Philosophy Supplement 43) Cambridge: CUP pp.229-251.

Crane, T. (2006) 'Is There a Perceptual Relation?' in T. Gendler and J. Hawthorne (eds.) *Perceptual Experience* Oxford: OUP pp.126-146.

Currie, G. (1995) *Image and Mind. Film, Philosophy, and Cognitive Science* Cambridge: CUP.

Dainton, B. (2000) *Stream of Consciousness* London: Routledge & Kegan Paul.

Dainton, B. (2001) *Time and Space* Acumen.

Dainton, B. (2008) 'The Experience of Time and Change' *Philosophy Compass* doi:10.1111/j.1747-9991.2008.00153.x.

De Clercq, R. and Horsten, L. (2004) 'Perceptual Indiscriminability: in Defence of Wright's Proof' *The Philosophical Quarterly* 54 (216) pp.439-444.

Defrin, R., Ohry, A., Blumen, N., and Urca, G. (2002) 'Sensory determinants of thermal pain' *Brain* 125 (3) pp.501-510.

Dennett, D.C. (1979) 'On the Absence of Phenomenology' in D. Gustafson and B. Tapscott (eds.) *Body, Mind and Method: Essays in Honor of Virgil C. Aldrich* Dordrecht: D. Reidel pp.93-114.

Dennett, D.C. (1991a) *Consciousness Explained* Boston MA: Little Brown & Co.

Dennett, D.C. (1991b) 'Real Patterns' *Journal of Philosophy* 88 pp.27-51.

Dennett, D.C. (1993) 'Caveat Emptor' *Consciousness and Cognition* 2 (1) pp.48-57.

Dennett, D.C. (1995) 'Is Perception the "Leading Edge" of Memory?' in A. Spadafora (ed.) *Iride: Luoghi della memoria e dell'oblio* 8 (14) pp.59-78 (Italian), English version at URL = http://ase.tufts.edu/cogstud/papers/ispercep.htm.

Dennett, D.C. and Kinsbourne, M. (1992a) 'Time and the observer: the where and when of consciousness in the brain' *Behavioral and Brain Sciences* 15 pp.183-200.

Dennett, D.C. and Kinsbourne, M. (1992b) 'Escape from the Cartesian Theater' (response to commentators) *Behavioral and Brain Sciences* 15 pp.234-47.

Descartes, R. (1986) *Meditations on First Philosophy* J. Cottingham (trans.) Cambridge: CUP.

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

Dretske, F. (2003) 'How Do You Know You are Not a Zombie?' in B. Gertler (ed.) *Privileged Access: Philosophical Accounts of Self-Knowledge* Ashgate pp.1-13.

Dretske, F. (2006) 'Perception without Awareness' in T.S. Gendler and J. Hawthorne (eds.) *Perceptual Experience* Oxford: OUP pp.147-180.

Dobbs, H.A.C. and Broad, C.D. (1951) 'The Relation between the Time of Psychology and the Time of Physics. Part I' *The British Journal for the Philosophy of Science* 2 (6) pp.122-141.

Dummett, M. (1975) 'Wang's Paradox' *Synthese* 30 pp.301-24.

Dummett, M. (ms.) 'Seeing Things Move' Unpublished typescript.

Emerson, R.W. (1867) *May Day and Other Pieces* London: G. Routledge & Sons.

Enns, J.T. and Di Lollo, V. (2000) 'What's new in visual masking?' *Trends in Cognitive Sciences* 4 (9) pp.345-352.

Enoka, R.M. (2002) *Neuromechanics of Human Movement* Human Kinetics.

Evans, G. (1980) 'Things Without the Mind' reprinted in his *Collected Papers* Oxford: Clarendon Press 1985 pp.249-90.

Fara, D.G. (2001) 'Phenomenal Continua and the Sorites' *Mind* 110 (440) pp.905-935, originally published under the name "Delia Graff".

Farah, M.J. (1992) 'The distributed pineal gland' *Behavioral and Brain Sciences* 15 p.209.

Foster, J. (1979) 'In Self-Defence' in G.F. Macdonald (ed.) *Perception and Identity* London: Macmillan.

Foster, J. (1982) *The Case for Idealism* London: Routledge & Kegan Paul.

Foster, J. (1991) *The Immaterial Self* London: Routledge.

Foster, J. (2000) *The Nature of Perception* Oxford: OUP.

Funkhouser, E. (2006) 'The Determinable-Determinate Relation' *Noûs* 40 (3) pp.548-569.

Gallagher, S. (1998) *The Inordinance of Time* Evanston: Northwestern University Press.

Geldard, F. and Sherrick, C. (1972) 'The Cutaneous "Rabbit": a Perceptual Illusion' *Science* 178 (57) pp.178-9.

Gennette, G. (1979) *Narrative Discourse: An Essay in Method* J.E. Lewin (trans.) Ithaca N.Y.: Cornell University Press.

Gibson, J.J. (1966) *The Senses Considered as Perceptual Systems* Boston: Houghton Mifflin.

Gillett, C. and Rives, B. (2005) 'The Non-Existence of Determinables: Or, a World of Absolute Determinates as Default Hypothesis' *Noûs* 39 (3) pp.483-504.

Gombrich, E.H. (1964) 'Moment and Movement in Art' *Journal of the Warburg and Courtauld Institutes* 27 pp.293-306.

Goodman, N. (1977) *The Structure of Appearance* Boston: D. Reidel, reprint of Goodman, N. (1951) *The Structure of Appearance* Cambridge, MA.: Harvard University Press.

Grice, H.P. (1961) 'The Causal Theory of Perception' *Supplementary Proceedings of the Aristotelian Society* 35 pp.121-52.

Grice, P. (1962) 'Some Remarks about the Senses' reprinted in his *Studies in the Way of Words* Harvard University Press 1989 pp.248-26

Guyer, P. (1987) *Kant and the Claims of Knowledge* Cambridge: CUP.

Haarmeier, T., Thier, P., Repnow, M., and Petersen, D. (1997) 'False perception of motion in a patient who cannot compensate for eye movements' *Nature* 389 pp.849-852.

Hamilton, A. (ms.) 'Rhythm and Time: a Major and Almost Entirely Neglected Philosophical Problem' Unpublished typescript.

Hamilton, W. (1861) *Lectures on Metaphysics and Logic* (2 vols.) Edinburgh and London: Blackwood and Sons.

Harman, G. (1990) 'The Intrinsic Quality of Experience' *Philosophical Perspectives* 4 pp.31-52.

Hellie, B. (2005) 'Noise and Perceptual Indiscriminability' *Mind* 114 pp.481-508.

Hilbert, D.R. (1987) *Color and Color Perception: A Study in Anthropocentric Realism* Stanford: CSLI.

Hirsh, I.J. and Sherrick Jr., C.E. (1961) 'Perceived order in different sense modalities' *Journal of Experimental Psychology* 62 pp.423-432.

Hodgson, S.H. (1898) *Metaphysic of Experience* (4 vols.) London: Longman, Green, and Co.

Hodgson, S.H. (1900) 'Perception of Change and Duration – A Reply' *Mind* 9 (34) pp.240-243.

Hoerl, C. (1998) 'The Perception of Time and the Notion of a Point of View' *European Journal of Philosophy* 6 (2) pp.156-171.

Hopkins, R. (1998) *Picture, Image & Experience* Cambridge: CUP.

Horgan, T. and Kriegel, U. (2007) 'Phenomenal Epistemology: What is Consciousness that We may Know It so Well?' *Philosophical Issues* 17 (1) pp.123-144.

Horgan, T. and Tienson, J. (2002) 'The Intentionality of Phenomenology and the Phenomenology of Intentionality' in D.J. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings* Oxford: OUP pp.520-33.

Hornsby, J. (1980-1) 'Which Physical Events Are Mental Events?' *Proceedings of the Aristotelian Society* 81 pp.73-92.

Hornsby, J. (1986) 'Physicalism, Events and Part-Whole Relations' in E. LePore and B. McLaughlin (eds.) *Actions and Events* Oxford: Blackwell pp.444-58.

Hurley, S. (1992) 'Some mistakes about consciousness and their motivation' *Behavioral and Brain Sciences* 15 pp.211-2.

Hurley, S. (1998) *Consciousness in Action* Harvard University Press.

Husserl, E. (1964) *The Phenomenology of Internal Time-Consciousness* M. Heidegger (ed.) and J. S. Churchill (trans.) Bloomington: Indiana University Press.

Ismael, J. (ms) 'Memory' Unpublished typescript URL = http://www.usyd.edu.au/time/ismael/papers/7.memory.pdf.

Jackson, F. (1993) 'Review: Appendix A (For Philosophers)' *Philosophy and Phenomenological Research* 53 (4) pp.899-903.

James, W. (1890) *The Principles of Psychology* New York: H. Holt and company.

Johnson, W.E. (1921) *Logic* Vol. 1 Cambridge: CUP.

Kanai, R., Paffen, C., Hogendoorn, H., and Verstraten, F. (2006) 'Time dilation in dynamic visual display' *Journal of Vision* 6 pp.1421-1430

Kant, I. (2003) *Critique of Pure Reason* N. Kemp Smith (trans.) New York: Palgrave Macmillan.

Kanwisher, N.G. (1987) 'Repetition blindness: Type recognition without token individuation' *Cognition* 27 pp.117-143.

Kellman, P.J. (1995) 'Ontogenesis of Space and Motion Perception' in W. Epstein and S. Rogers (eds.) *Perception of Space and Motion*, *Handbook of Perception and Cognition* Academic Press pp.327-364.

Kelly, S. (2005) 'The Puzzle of Temporal Experience' in A. Brook and K. Akins (eds.) *Cognition and the Brain* Cambridge: CUP pp.208-240.

Kolers, P. and von Grünau, M. (1976) 'Shape and color in apparent motion' *Vision Research* 14 pp.329-335.

Kotchoubey, B., Lang, S., Winter S., and Birbaumer, N. (2003) 'Cognitive processing in completely paralyzed patients with amyotrophic lateral sclerosis' *European Journal of Neurology* 10 pp.551-558.

Kripke, S. (1978) *Time and Identity* Unpublished typescript.

Kurthen, M., Moskopp, D., Linke, D.B., and Reuter, B.M. (1991) 'The Locked-in Syndrome and the Behaviorist Epistemology of Other Minds' *Theoretical Medicine* 12 pp.69-79.

Lachter, J., Durgin, F., and Washington, T. (2000) 'Disappearing percepts: Evidence for retention failure in metacontrast masking' *Visual Cognition* 7 pp.269-279.

Landman, R., Spekreijse, H., and Lamme, V.A.F. (2003) 'Large capacity storage of integrated objects before change blindness' *Vision Research* 43 (2) pp.149-64.

Le Poidevin, R. (2007) *The Images of Time* Oxford: OUP.

Lee, G. (2007) 'Consciousness In A Space-Time World' *Philosophical Perspectives* 21 (1) pp.341-374.

Leibniz, G. (1965) 'Critical Remarks Concerning the General Part of Descartes' Principles' in his *Monadology and Other Philosophical Essays* P.S. and A.M. Schrecker (trans.) Indianapolis: Bobbs-Merrill.

Lennon, T.M. (1995) '*Veritas Filia Temporis*: Hume on Time and Causation' in S. Tweyman (ed.) *David Hume: Critical Assessments* Vol. III Routledge pp.344-360.

Lewis, D. (1980) 'Veridical Hallucination and Prosthetic Vision' in his *Philosophical Papers Volume II* Oxford: OUP 1987 pp.273-291.

Lloyd, D. (1992) 'Towards an identity theory of consciousness' *Behavioral and Brain Sciences* 15 pp.215-6.

Locke, J. (1975) *An Essay Concerning Human Understanding* P.H. Nidditch (ed.) Oxford: Clarendon Press.

Lockwood, M. (1989) *Mind, Brain and Quantum* Oxford: Basil Blackwell.

Lycan, W. (1992) 'UnCartesian materialism and Lockean introspection' *Behavioral and Brain Sciences* 15 pp.216-17.

Mabbott, J.D. (1951) 'Our Direct Experience of Time' *Mind* 60 pp.153-167.

MacFarlane, J. (2003) 'Future Contingents and Relative Truth' *The Philosophical Quarterly* 53 pp.321-36.

MacFarlane, J. (forthcoming) 'Truth in the Garden of Forking Paths' in M. Kolbel and M. Garcia-Carpintero (eds.) *Relative Truth* Oxford: OUP.

MacFarlane, J. *Philosophical Glossary* URL = http://johnmacfarlane.net/135/glossary.html.

Macken, W., Tremblay, S., Houghton, R., Nicholls, A., and Jones, D. (2003) 'Does auditory streaming require attention? Evidence from attentional selectivity in short-term memory' *Journal of Experimental Psychology* 29 pp.43-51.

Martin, M.G.F. (1998a) 'An Eye Directed Outward' in C. Wright, B. Smith, and C. Macdonald (eds.) *Knowing Our Own Minds* Oxford: OUP pp.99-123.

Martin, M.G.F. (1998b) 'Setting things before the mind' in A. O'Hear (ed.) *Current issues in philosophy of mind* Cambridge: CUP pp.157-180.

Martin, M.G.F. (2001) 'Out of the Past: Episodic Recall as Retained Acquaintance' in C. Hoerl and T. McCormack (eds.) *Time and Memory: Issues in Philosophy and Psychology* Oxford: OUP pp.257-284.

Martin, M.G.F. (2002) 'The transparency of experience' *Mind & Language* 17 (4) pp.376-425.

Martin, M.G.F. (2004) 'The Limits of Self-Awareness' *Philosophical Studies* 120 pp.7-89.

Martin, M.G.F. (2006) 'On Being Alienated' in T. Gendler and J. Hawthorne (eds.) *Perceptual Experience* Oxford: OUP pp.354-410.

McGinn, C. (1982) *The Character of Mind* Oxford: OPUS.

McGinn, C. (1995) 'Consciousness Evaded: Comments on Dennett' *Philosophical Perspectives: AI, Connectionism and Philosophical Psychology* 9 pp.241-249.

Mellor, D. H. (1981) *Real Time* Cambridge: CUP.

Merleau-Ponty, M. (2002) *Phenomenology of Perception* C. Smith (trans.) London and New York: Routledge.

Micheyl, C., Tian, B., Carlyon, R., and Rauschecker, J. (2005) 'Perceptual Organization of Tone Sequences in the Auditory Cortex of Awake Macaques' *Neuron* 48 (1) pp.139-148.

Millar, A. (2000) 'The Scope of Perceptual Knowledge' *Philosophy* 75 pp.75-88.

Miller, I. (1984) *Husserl, Perception, and Temporal Awareness* Cambridge, MA.: MIT Press.

Mills, E. (2002) 'Fallibility and the Phenomenal Sorites' *Noûs* 36 (3) pp.384-407.

Moore, G.E. (1903) 'The Refutation of Idealism' *Mind* 12 (48) pp.433-453.

Mourelatos, A.P.D. (1978) 'Events, Processes and States' *Linguistics and Philosophy* 2 pp.415-34.

Näätänen, R., Tervaniemi, M., Sussman, E., Paavilainen, P., and Winkler, I. (2001) '"Primitive intelligence" in the auditory cortex' *Trends in Neuroscience* 24 pp.283-288.

Nagel, T. (1974) 'What is it like to be a Bat?' *The Philosophical Revie*w 83 (4) pp.435-50.

Nagel, T. (1986) *The View from Nowhere* Oxford: OUP.

Neumann, O. and Klotz, W. (1994) 'Motor Response to Nonreportable, Masked Stimuli: Where is the Limit of Direct Parameter Specification' in C. Umilta and M. Moskovitch (eds.) *Attention and Performance* 15 Cambridge, MA.: MIT Press pp.123-150.

Noë, A. (2002) (ed.) *Is the Visual World a Grand Illusion?* Special Issue, *Journal of Consciousness Studies* 9 (5-6).

Noreika, V. (forthcoming) *Consciousness and the Brain in Dreaming: Empirical and Theoretical Studies* PhD. Department of Psychology, University of Turku.

O'Callaghan, C. (2007) *Sounds* Oxford: OUP.

O'Shaughnessy, B. (2000) *Consciousness and the World* Oxford: OUP.

Pace, M. (2007) 'Blurred Vision And The Transparency Of Experience' *Pacific Philosophical Quarterly* 88 (3) pp.328-354.

Papineau, D. (2007) 'Reuniting (scene) phenomenology with (scene) access' *Behavioral and Brain Sciences* 30 (5/6) p.521.

Peacocke, C. (1985) 'Imagination, experience and possibility' in J. Foster and H. Robinson (eds.) *Essays on Berkeley* Oxford: Clarendon Press pp.19-35.

Peacocke, C. (2003) 'Conscious Attitudes, Attention and Self-Knowledge' in B. Gertler (ed.) *Privileged Access: Philosophical Accounts of Self-Knowledge* Ashgate pp.83-110.

Pelak, V. and Hoyt, W. (2005) 'Symptoms of Akinetopsia Associated with Traumatic Brain Injury and Alzheimer's Disease' *Neuro-Ophthalmology* 29 pp.137-142.

Pickard, H. (2001) 'Knowing and Owning a Body' D.Phil. Thesis, University of Oxford.

Pöppel, E. (1978) 'Time Perception' in R. Held, H. Leibowitz and H. Teuber (eds.) *Perception* Berlin: Springer-Verlag pp.713-729.

Price, R. (2005) 'Content Ascriptions and the Reversibility Constraint' *Philosophical Perspectives* 19 (1) pp.353-374.

Prichard, H.A. (1938) 'The Sense-Datum Fallacy' reprinted in his *Knowledge and Perception* Oxford: Clarendon Press (1950) pp.200-214.

Prichard, H.A. (1950a) 'The Apprehension of Time' in his *Knowledge and Perception* Oxford: Clarendon Press pp.47-51.

Prichard, H.A. (1950b) 'Seeing Movements' in his *Knowledge and Perception* Oxford: Clarendon Press pp.41-46.

Pylyshyn, Z.W. (1979) 'Do mental events have durations?' *Behavioral and Brain Sciences* 2 (2) pp.277-8.

Reid, T. (1827) *Essays on the Powers of the Human Mind; to which are added, An Essay on Quantity, and An Analysis of Aristotle's Logic* London: printed for Thomas Tegg.

Reingold, E.M. (1992) 'Conscious versus unconscious processes: Are they qualitatively different?' *Behavioral and Brain Sciences* 15 pp.218-9.

Roache, R. (1999) 'Mellor and Dennett on the Perception of Temporal Order' *Philosophical Quarterly* 49 pp.231-238.

Rock, I. (1956) 'The orientation of forms on the retina and in the environment' *American Journal of Psychology* 69 pp.513-528.

Rollins, M. (1992) 'Content and conformation: Isomorphism in the neural sway' *Brain and Behavioral Sciences* 15 (2) pp.219-220.

Rosenthal, D.M. (1990) 'A Theory of Consciousness' *Report No. 40*, Research Group on Mind and Brain. ZiF, University of Bielefeld, Bielefeld, Germany, reprinted in N. Block, O. Flanagan and G. Güzeldere (eds.) *The Nature of Consciousness* Cambridge, MA.: MIT Press 1997 pp.729-753.

Rosenthal, D. (1993) 'Thinking that One Thinks' in M. Davies and G. Humphreys (eds.) *Consciousness* Oxford: Blackwell pp.197-223.

Rosenthal, D.M (2002) 'Explaining Consciousness' in D.J. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings* Oxford: OUP pp.406-421.

Rosenthal, D.M. (2005) *Consciousness and Mind* Oxford: OUP.

Rothstein, S. (2004) *Structuring Events* Blackwell Publishing.

Ruben, D-H. (2008) 'Disjunctive Theories of Perception and Action' in A. Haddock and F. Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge* Oxford: OUP pp.227-244.

Russell, B. (1912) *The Problems of Philosophy* Oxford: OUP.

Russell, B. (1914) *Our Knowledge of the External World as a Field for Scientific Method in Philosophy* Chicago and London: Open Court Publishing.

Russell, B. (1915) 'On the Experience of Time' *The Monist* 25 pp.212-233.

Russell, B. (1921) *The Analysis of Mind* London: G. Allen and Unwin.

Russell, B. (1927) *The Analysis of Matter* London: K. Paul, Trench, Trübner & Co. Ltd.

Russell, B. (1948) *Human Knowledge: its scope and limits* London: Allen and Unwin.

Sadie, S. (1980) (ed.) *The New Grove Dictionary of Music and Musicians* Oxford: OUP.

Sanford, D.H. (2007) 'Determinates vs. Determinables' in E.N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2007 Edition) URL = http://plato.stanford.edu/archives/fall2007/entries/determinate-determinables/.

Schacter, D.L. (2001) *The seven sins of memory: how the mind forgets and remembers* Boston, MA: Houghton Mifflin.

Seager, W. (1999) *Theories of Consciousness: An Introduction and Assessment* London: Routledge.

Searle, J.R. (1990) 'Consciousness, Explanatory Inversion, and Cognitive Science' *Behavioral and Brain Sciences* 13 pp.585-595.

Searle, J.R. (1995) '"The Mystery of Consciousness": An Exchange' *New York Review of Books* 42 (20) December 21, 1995. URL = http://www.nybooks.com/articles/1680.

Sellars, R.W. (1917) 'The Status of Epistemology' *The Journal of Philosophy, Psychology and Scientific Methods* 14 (25) pp.673-680.

Shepherd, R.N. (1992) 'Mental representation: Always delayed but not always ephemeral' *Behavioral and Brain Sciences* 15 pp.223-4.

Sherwood, M.E.W. (1897) *An Epistle to Posterity* New York: Harper.

Shoemaker, S. (1990) 'Qualities and qualia: What's in the mind?' reprinted in his *The First-Person Perspective, and Other Essays* Cambridge: CUP 1996 pp.97-120.

Shoemaker, S. (1994) 'Self-Knowledge and "Inner Sense": Lectures I-III' *Philosophy and Phenomenological Research* 54 (2) pp.249-314.

Siegel, S. (2004) 'Indiscriminability and the Phenomenal' *Philosophical Studies* 120 pp.90-112.

Siegel, S. (2006) 'Which properties are represented in perception?' in T. Gendler and J. Hawthorne (eds.) *Perceptual Experience* Oxford: OUP pp.481-503.

Siegel, S. (2008) 'The visual experience of causation' *Philosophical Quarterly*.

Siewert, C. (1998) *The Significance of Consciousness* Princeton University Press.

Simpson, J. and Weiner, E. (1989) (eds.) *Oxford English Dictionary* Oxford: OUP.

Sligte, I.G., Lamme, V.A.F., and Scholte, H.S. (2006) 'Capacity limits to awareness' paper presented at the *Association for the Scientific Study of Consciousness* Oxford, U.K.

Sligte, I.G., Scholte, H.S., and Lamme, V.A.F. (2008) 'Are there multiple visual short-term memory stores?' *PLoS ONE* 3 (2) pp.1-9.

Smith, A.D. (2002) *The Problem of Perception* Harvard University Press.

Smith, E. and Delargy, M. (2005) 'Locked-in syndrome' *British Medical Journal* 330 pp.406-409.

Snowdon, P. (2006) 'Radical Externalisms' *Journal of Consciousness Studies* 13 (7-8) pp.187-198.

Soteriou, M. (2005) 'The Subjective View Of Experience And Its Objective Commitments' *Proceedings of the Aristotelian Society* 105 (2) pp.193-206.

Soteriou, M. (2007) 'Content and the Stream of Consciousness' *Philosophical Perspectives* 21 (1) pp.543-568.

Sperling, G. (1960) 'The Information Available in Brief Visual Presentations' *Psychological Monographs* 74 (11).

Sprigge, T. (1993) *James and Bradley: American Truth and British Reality* Chicago: Open Court.

Sprigge, T. (2002) 'Idealism' in R. Gale (ed.) *The Blackwell Guide to Metaphysics* Malden, MA.: Blackwell pp.291-241.

Staddon, J.E.R. (2000) 'Consciousness and Theoretical Behaviorism' *American Zoologist* 40 pp.874-882.

Steinmann, R., Pizlo, Z., and Pizlo, F. (2000) 'Phi is not beta, and why Wertheimer's discovery launched the Gestalt revolution' *Vision Research* 40 pp.2257-2264.

Stetson, C., Fiesta, M.P., and Eagleman, D.M. (2007) 'Does Time Really Slow Down during a Frightening Event?' *PLoS ONE* 2 (12) pp.1-3.

Steward, H. (1997) *The Ontology of Mind: Events, Processes, and States* Oxford: OUP.

Stoljar, D. (2004) 'The Argument from Diaphanousness' in M. Ezcurdia, R. Stainton, and C. Viger (eds.) *New Essays in the Philosophy of Language and Mind*, *Supplemental volume of The Canadian Journal of Philosophy* Calgary: University of Calgary Press pp.341-390.

Stout, G.F. (1900a) 'Perception of Change and Duration' *Mind* 9 (33) pp.1-7, reprinted in Stout 1930.

Stout, G.F. (1900b) 'Reply to Hodgson' *Mind* 9 (34) p.243.

Stout, G.F. (1930) *Studies in Philosophy and Psychology* London: Macmillan & co.

Strawson, P.F. (1966) *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason* London: Methuen.

Strawson, P.F. (1979) 'Perception and its objects' in G. McDonald (ed.) *Perception and Identity: Essays presented to A. J. Ayer* London, Basingstoke: Macmillan pp.41-60.

Susser, E., Sprecher, E., and Yarnitsky, D. (1999) 'Paradoxical heat sensation in healthy subjects' *Brain* 122 pp.239-246.

Sussman, E. Ritter, W., and Vaughan, H. Jr. (1999) 'An investigation of the auditory streaming effect using event-related brain potentials' *Psychophysiology* 36 pp.22-34.

Taylor, B. (1977) 'Tense and Continuity' *Linguistics and Philosophy* 1 pp.199-220.

Taylor, B. (1985) *Modes of Occurrence* Oxford: Basil Blackwell.

Tenney, J. and Polansky, L. (1980) 'Temporal Gestalt Perception in Music' *Journal of Music Theory* 24 pp.205-241.

Thau, M. (2002) *Consciousness and Cognition* Oxford: OUP.

Todd, S.J. (2006) 'Unmasking Multiple Drafts' *Philosophical Psychology* 19 (4) pp.477-494.

Tomlinson, C. (1997) *Selected Poems* Oxford: OUP.

Travis, C. (2004) 'The Silence of the Senses' *Mind* 113 (449) pp.57-94.

Treanor, N. (ms.) 'The Ontology of Experience' URL = http://www.qifp.com/content-text/Papers/The%20Ontology%20of%20Experience%20--%20Treanor.pdf.

Treisman, M. (1992) 'Does the perception of temporal sequence throw light on consciousness' *Behavioral and Brain Sciences* 15 pp.225-8.

Tye, M. (1992) 'Visual Qualia and Visual Content' in T. Crane (ed.) *The Contents of Experience* Cambridge: CUP pp.158-76.

Tye, M. (1993) 'Reflections on Dennett and Consciousness' *Philosophy and Phenomenological Research* 53 (4) pp.893-898.

Tye, M. (1995) *Ten Problems of Consciousness: A Representational Theory of the Phenomenal* Cambridge, MA.: MIT Press.

Tye, M. (2002) 'Representationalism and the Transparency of Experience' *Noûs* 36 (1) pp.137-151.

Tye, M. (2003) *Consciousness and Persons* Cambridge, MA.: MIT Press.

Tye, M. (2006) 'Content, Richness, and Fineness of Grain' in T.S. Gendler and J. Hawthorne (eds.) *Perceptual Experience* Oxford: OUP pp.504-530.

Van Den Bergh, S. (1954) *Averroes' Tahafut Al-Tahafut* Vol. I Oxford: OUP.

Van Gulick, R. (1992) 'Time for more alternatives' *Behavioral and Brain Sciences* 15 pp.228-9.

van Inwagen, P. (1993) *Metaphysics* Westview Press.

Vendler, Z. (1957) 'Verbs and Times' *The Philosophical Review* 66 (2) pp.143-160.

von Fieandt, K. (1966) *The World of Perception* Homewood, IL: The Dorsey Press.

von Helmholtz, H.L.F. (1910) *Treatise on Physiological Optics* New York: Dover.

Walton, K.L. (2008) 'Experiencing Still Photographs: What Do You See and How Long Do You See It?' in his *Marvellous Images* Oxford: OUP pp.157-192.

Warren, R., Obusek, C., Farmer, R., and Warren, R. (1969) 'Auditory sequence: confusion of patterns other than speech or music' *Science* 164 pp.586-7.

Way, J. (2007) 'Self-knowledge and the limits of transparency' *Analysis* 67 (3) pp.223-230.

Wertheim, A.H. (1994) 'Motion Perception During Self-Motion: The Direct versus Inferential controversy revisited' *Behavioral and Brain Sciences* 17 pp.293-355.

Wertheim, A.H. (1987) 'Retinal and extra-retinal information in movement perception: how to invert the Filehne illusion' *Perception* 16 (3) pp.299-308.

Wertheim, A.H. (1999) 'Motion percepts: "Sense specific," "kinematic," or …?' *Behavioral and Brain Sciences* 22 pp.338-340.

Wertheimer, M. (1912) 'Experimentelle Studien über das Sehen von Bewegung' *Zeitschrift für Psychologie* 61 pp.161-265.

Williamson, T. (1996) 'Cognitive homelessness' *The Journal of Philosophy* 93 (11) pp. 554-573.

Williamson, T. (2000) *Knowledge and its Limits* Oxford: OUP.

Wollheim, R. (1987) *Painting as an Art* Princeton: Princeton University Press.

Wright, C. (1975) 'On the Coherence of Vague Predicates' *Synthese* 30 pp.325-365.

Wright, C. (1992) *Truth and Objectivity* Cambridge, MA.: Harvard.

Wright, C. (1998) 'Self-Knowledge: the Wittgensteinian Legacy' in A. O'Hear (ed.) *Current Issues in Philosophy of Mind* Cambridge: CUP pp.101-122.

Wright, C. (2002) 'What Could Antirealism about Ordinary Psychology Possibly Be?' *The Philosophical Review* 111 (2) pp.205-233.

Yaffe, G. (2003) 'Time in the Movies' in P.A. French and H.K. Wettstein (eds.) *Midwest Studies in Philosophy XXVII: Meaning in the Arts* Oxford: Blackwell pp.115-138.

Zeki, S. (1991) 'Cerebral Akinetopsia (Visual Motion Blindness)' *Brain* 114 pp.811-824.

Zihl, J. et al. (1983) 'Selective disturbance of movement vision after bilateral posterior brain damage' *Brain* 106 pp.313-340.

Zihl, J. et al. (1991) 'Disturbance of movement vision after bilateral posterior brain damage: further evidence and follow up observations' *Brain* 114 pp.2235-2252.