26

Time in Mind

JULIAN KIVERSTEIN AND VALTTERI ARSTILA

1. Temporality in Neuroscience

Neural processing takes time, so it is unsurprising that time plays a significant role in neuroscientific explanation. For instance, when subjects are shown two visual stimuli that differ in luminosity, the brighter of the two stimuli activates the visual cortex up to 80 milliseconds before the less bright stimulus (Eagleman 2008). In addition to this latency effect, there is also substantial evidence that spatial attention speeds up the processing of stimuli (Spence and Parise 2010). Tracking the temporal properties of neural processing has also shed light on the role of so-called recurrent processing or feedback projections from higher cortical areas to lower areas. Finally, consider how all brain imaging techniques have a temporal resolution below which these techniques cannot distinguish activation in a given brain region. The limited temporal resolution of brain-imaging techniques restricts the kinds of data one can obtain using these methods.[1] This raises interesting questions for neuroscience about what psychological processes we can and cannot track using these methods, and hence how transparent the mind really is in brain-imaging.

These examples, and many others, demonstrate how time is undoubtedly an important variable in neuroscientific explanations of perceptual processing. However, there is little here to pique the interests of the philosopher of time. We think the place where the temporal issues in neuroscience genuinely have something to contribute to the philosophy of time is when we begin to think about the neural underpinnings of temporal experience.[2] Before we can show how this is so, we need to clear some ground.

Consider first how the temporal properties of events do not always match the temporal properties of our experiences of such events. This is to say that experiences representing temporal properties such as duration, simultaneity, or succession are not

A Companion to the Philosophy of Time, First Edition. Edited by Heather Dyke and Adrian Bardon.

always veridical. An experience may represent the temporal order of events incorrectly – we hear the thunder after we see the lightning although both happen simultaneously. An experience may represent an event – waiting in a line – as lasting for many minutes even though in reality the duration was much shorter. Both experiences are examples of what we will call "temporal illusions" – illusions in which our experiences misrepresent the temporal properties of events.

What the existence of illusions like these demonstrates is a *subjective* dimension to our experience of the temporal properties of events. Our experiences of temporal properties admit of an experience–reality distinction. Thus, we can distinguish temporal properties as they appear to us (or are represented by us) from the temporal properties we are representing. How are the temporal properties as they are experienced by us related to the temporal properties of the events we experience? Spatial properties like size, height, number, (and perhaps shape) are often supposed to be primary qualities. When our experiences correctly represent these properties there is a resemblance or match between what our experiences represent and properties of objects our experiences represent. The situation is different for secondary qualities like color, smell, or taste where there is substantial philosophical debate about the exact nature of these properties. Are temporal properties as they appear to us best understood as primary or secondary qualities? We might think by analogy with spatial properties that the answer is obvious and temporal properties as they figure in experience must be primary qualities: When our experiences correctly represent temporal properties there is a match between our experience of an event and the event we are experiencing. We will see below however that the answer is by no means obvious in the case of our subjective experience of temporal properties.

2. The Perplexing Nature of Apparent Motion

Temporal illusions occur when there is a mismatch between the temporal properties our experience represents an event as having and the temporal properties the event really possesses. Some temporal illusions have straightforward explanations. Consider, for instance, our earlier example of seeing the lightning before hearing the thunder. The explanation for this experience is due to light traveling faster than sound. Other examples of temporal illusions can be straightforwardly explained by the "normal" functioning of our nervous system and the time it takes the brain to process information in different sense modalities. When one's toe and nose are tapped at the same time, the sensory inputs from the nose have a much shorter distance to travel to somatosensory cortex than inputs from the toes. Hence they are not processed at the same time, or at least their processing does not begin at the same time. Some duration illusions are also likely to succumb to fairly straightforward explanations. The duration of a stimulus shown for a few milliseconds is always judged to have lasted longer. This stimulus can influence subsequent temporal integration for up to 100 milliseconds, a period that far exceeds the objective duration of the stimulus. Nevertheless, this may simply reflect the time our brains need to form a proper representation of the stimulus and for this representation to decay. Novel stimuli are estimated to have longer duration than more

familiar stimuli but this may simply be due to the fact that the novelty brings about an increased processing of the stimulus.

There are, however, other examples of temporal illusions that prove both philosophically and scientifically puzzling. One that has been much discussed by philosophers is the *phi phenomenon* (first described by Max Wertheimer 1912), an illusion of apparent motion induced by presenting two spatially distinct visual stimuli (henceforth *A* and *B*) separated by a short temporal gap.[3] If the temporal gap between *A* and *B* is too short (less than 50 milliseconds) we are likely to perceive the stimuli as simultaneous. If the interval is too long – 400 milliseconds or longer – the illusion is weak and *A* and *B* will usually be perceived as two distinct stimuli. When the temporal gap is between these limits (it works best when the interval is between 50 and 200 milliseconds), *A* and *B* are fused into a single stimulus *C* and we experience the illusory or apparent motion of *C* from the location at which *A* is initially presented to the location of *B*. Thus, if *B* is presented in the right way at time *t*, this can affect how we experience a stimulus *A* presented 50–200 milliseconds earlier than *t*. In particular the fine-grained details of the timing of *B* determines whether we experience *A* (presented at an earlier time) as stationary or as moving. Now the puzzle is to explain how a stimulus presented at time *t* can modulate what we experience to occur at a time prior to *t*? Obviously by the time we have our experience of *A* our brains must already have processed *B*. We will see later in the chapter (in sections 7 and 8), however, that saying this much still leaves a number of possibilities open as to how the brain's processing of B can influence our experience of A.

A tactile version of apparent motion raises similar questions. Geldard and Sherrick (1972) induced this tactile illusion (now known as the cutaneous rabbit) by placing small mechanical devices at evenly spaced locations on a subject's arm and shoulder, producing evenly spaced sequences of taps. Sometimes the device would produce taps at the same location on the arm, but on other occasions a sequence of taps would be delivered in rapid succession first to the wrist, then ten centimeters away from the wrist to the arm, then a further ten centimeters towards the elbow and so on. Under this condition what the subject reports experiencing are taps evenly distributed from the wrist up the arm like the sensations of an animal hopping up one's arm. Where the subject feels the tap (on the wrist or further up the arm) depends on the occurrence and location of subsequent taps. Just as with the phi phenomenon, we can ask how the subject's tactile experience of *A* at time *t* can depend on stimuli that are presented after *A*. We can agree that the brain must have already processed the later taps in order for them to have an impact on how we experience the earlier taps. Saying this much still leaves open how to explain the influence of stimuli at time *t* on our experience of stimuli presented at a time earlier than *t*.

One influential philosophical response to temporal illusions like these has been to insist upon a *content–vehicle* distinction for temporal representation in the brain (Dennett and Kinsbourne 1992/1997; Grush 2007). By the "vehicles" of temporal representation we mean the sub-personal, functional or neurobiological realizers of experience. The realizers of experience are understood as "vehicles" because they are the bearers of the contents available to the subject of experience. Dennett and Kinsbourne say we must distinguish time as it is represented in experience (content) from the timing of

representations (the vehicles that are doing the representing). The timing of neural processing in these examples of apparent motion seems to come apart from the timing of the events we experience. As already noted, we can only experience movement after the second stimulus has been processed sufficiently to influence our experience of the first. This is, however, not the order in which we experience events. We see a stimulus moving from location *A* to *B* or we feel an evenly spaced sequence of taps delivered to our arm. The temporal order in which events are processed by the brain doesn't in these cases match the temporal order in which events are experienced (because the movement is experienced before anything has been experienced at the location of the second stimulus).

Dennett and Kinsbourne (1992/1997) suggest temporal properties might be recorded by the brain in much the same way as letters are dated. When we receive multiple letters written on different days, we later reconstruct the order of the events recounted in the letters based on the dates of the letters, not on their arrival times. In a similar fashion, Dennett and Kinsbourne suggest that the temporal contents of experiences may not be determined by their arrival times (e.g., by the neural correlates of a token experience *A* occurring before or after the neural correlates of an experience *B*). Instead, temporal properties may be encoded in the contents of experience in much the same way as we work out what happened when in a series of letters by dating the events the letters describe.[4] The brain does not need to represent that *A* happens before *B* by tokening representations that occur in this order.

While it has been much more common to find philosophers, and recently also neuroscientists, using temporal illusions to argue along the lines of Dennett and Kinsbourne,[5] it is also possible to resist the inference from temporal illusions to a content–vehicle distinction for temporal experience. To reject such a distinction for temporal experience is to hold that the neural correlates of temporal experience treat time as its own representation.[6] Kelly (2005) cites the Gestalt psychologist Wolfgang Köhler as a proponent of such a view. Köhler writes that "Experienced order in time is always structurally identical with a functional order in the sequence of correlated brain processes." Consider an experience of a red flash followed by an experience of a green flash; Köhler's position says that the neural events ($n_1 \ldots n_n$) correlated with our experience of the red flash must occur before the neural events ($m_1 \ldots m_n$) correlated with our experience of the green flash. Now consider the phi phenomenon. We have already noted how, if we are to experience a single dot moving from position *A* to *B*, the brain must already know that a dot has been presented at location *B*. Still, this is quite consistent with the claim that the neural correlates ($n_1 \ldots n_n$) of our experience of the moving dot at location *A* occur before the neural correlates ($m_1 \ldots m_n$) of our experience of the moving dot at location *B*. The temporal properties of the neural correlates of experiences do not need to match the temporal properties of the unconscious processing stages preceding them. Thus, there could be a match of the kind Köhler describes between the order of our experiences in time and the order in which the sequence of correlated brain processes occur.[7]

There is thus a disagreement about whether temporal illusions require us to make a content–vehicle distinction for temporal experience. What we will call "the time as its own representation view" (henceforth TOR) says that in order to experience *A* followed

by *B*, the neural representation of *A* must precede the neural representation of *B*. Whereas views that distinguish contents from vehicles deny this. As Dennett and Kinsbourne write:

> What matters for the brain is not necessarily when individual representing events happen in various parts of the brain (as long as they happen in time to control the things that need controlling!) but their temporal content. That is, what matters is that the brain can proceed to control events "under the assumption that A happened before B."
>
> (Dennett and Kinsbourne 1992/1997, 150)

Certainly we must make a content–vehicle distinction for other types of property we experience, but is this also the case for temporal properties? Blue is not represented in our brains for instance by vehicles that are literally colored blue. More controversially, we can have representations whose contents exhibit systematicity without also having compositionally structured representational vehicles (Hurley 1998). There is, however, an important difference between these examples and the case of time. As Phillips (forthcoming) rightly stresses, both our experience and the objects of experience have temporal properties. We experience events that have temporal properties such as succession, simultaneity, and duration, and these very same temporal properties also attach to our experiences. Experiences take up a period of time, succeed one another or happen simultaneously with other experiences, and occur in a temporal order. The fact that experience and its objects both share temporal properties gives us a powerful motivation for resisting any content–vehicle distinction for temporal experience. To say that experience and its objects share temporal properties is just to say that there is no content–vehicle distinction for experience (at least on the assumption that experiences have contents). There is instead a match between the temporal properties an experience represents in virtue of its content and the temporal properties of the experience doing the representing.[8] In the next section, we will make this view a little more precise by considering what it means to posit such a match between the contents and vehicles of temporal experience.

3. The Structure of Temporal Experience

We have just claimed that episodes of experiences have temporal properties like duration, succession, and continuity. Experiences don't just represent temporal properties but have temporal properties in their own right. Let us briefly say something in defense of this claim. Consider first the way in which we experience continuity and change. We see moving things, we feel enduring sensations, and we hear organized sounds such as melodies. It might naturally be thought that for an experience to present us with any of these objects that experience must unfold over time. We cannot see a moving object (e.g., an approaching car) in a durationless instant. Movement occurs over an interval of time, and so, it might be thought, must our experience if is to present us with movement. The same goes for change: if we are to hear the changing notes in a melody our auditory experience must span a period of time long enough for us to hear the melody. This is just the jumping off point for a number of metaphysical questions about how, if at all, it is possible for experience to present us with events that take up whole inter-

vals of time. We will return to some of these questions below (in sections 5 and 6). The important point we wish to hold on to for now is that experiences somehow occupy whole intervals of time during which we can be presented with movement and change. Experiences are temporally extended processes or events.[9] To say experiences take up whole intervals of time is just another way of saying that experiences have temporal properties, they don't represent temporal properties by means of their contents.

If the above line of reasoning is along the right tracks, we can also say that experiences have a temporal structure. Structures are made up of elements or parts arranged and organized in a certain manner. We can describe this organization of the parts in terms of the relations between the parts. We have said that experiences are temporally extended events, and any event that extends through time can be divided into parts or stages or time slices. Each part will be datable in relation to other parts that make up an experience. It will also take up a certain amount of time so that when we put the parts together we get an event of a certain duration. It will stand in relations of temporal order (simultaneity and succession) to the other parts of which the experience is composed. On the basis of our claim that experiences are temporally extended events, we'll henceforth talk of experiences as having a *temporal structure.* The temporal structure of an experience consists of the parts that make up an experience and the temporal relations between those parts.[10]

We can now make more precise the idea of a match between the contents and vehicles of temporal experience. We suggest that the way to understand this match is in terms of an experience's temporal structure. First, there are the parts that make up a temporally extended experience and the monadic properties of those parts. There is a match between contents and vehicles when, for instance, an experience represents a sound (e.g., the "Do" in "Do Re Mi") as having a particular duration and it does so by having a part that lasts the very same duration. Second, there are the relations between the parts: if we experience "Do" as coming before "Re" this is because the vehicle that is doing the representing has a part that represents "Do" and a part that represents "Re" and the former occurs in time before the latter.

Returning now to the debate between TOR and its opponents, we have just seen how according to TOR the contents and vehicles of temporal experience have matching temporal structure. Opponents of TOR deny this. This they can do in two ways. The first option is to deny that experiences have a temporal structure. Chuard (2011), for instance, defends what he calls "temporal perceptual atomism" (TPA), a view which denies that we perceptually experience movement and change by means of temporally extended contents. Chuard's TPA claims that the contents of perceptual experience are confined to instants or very short intervals of time that do not persist long enough for us to perceive movement or change. If experiences lack temporally extended contents, they also lack temporal structure – this is, we take it, what Chuard means in referring to token experiences as *atomic*. We will consider TPA in the next section. A second strategy for resisting TOR accepts that experiences have a temporal structure but denies that there need be any match between the temporal structure of the events represented in experience and the vehicles that are doing the representing. Tye (2003) and Grush (2007) have both defended different versions of this strategy, but the view we will discuss below can be traced back to the phenomenological philosopher, Edmund Husserl. We will outline Husserl's position in section 5.

4. Temporal Perceptual Atomism

How do we come to be conscious of, and so have experiential access to, movement and change according to Chuard's TPA? He suggests we do so by means of a succession of distinct experiences E_1, E_2. . . E_n, each of which represents a part of a temporally extended process, and which occur in such close succession that we cannot notice the gaps between them. Experience-based (episodic) memory then provides us with information about past experiences on the basis of which we can make a judgment about movement and change. Chuard is in agreement with Thomas Reid (1850), who also argued that motion is not really experienced but only discerned with the aid of memory. Those who take experiences to have a temporal structure make the mistake of taking phenomenology at face value, says Chuard. Atomists by contrast allow that "phenomenological features can also result from cognitive, mnemonic, and introspective limitations of various sorts" (Chuard 2011, 11). Chuard is recommending then what he calls an "error theory" about the phenomenology of temporal experience.

There are good reasons to reject an error theory about the phenomenology of temporal experiences. We will do so on the basis of the much quoted Jamesian (and Kantian) observation that a succession of experiences does not make an experience of succession[11] or to quote James (1890): "a succession of feelings, in and of itself, is not a feeling of succession." Consider C.D. Broad's example of looking at a clock and seeing the second hand smoothly move about the clock face as contrasted with seeing that the hour hand has moved.

> . . . to see a second hand moving is quite a different thing from "seeing" that an hour-hand has moved. In the one case we are concerned with something that happens within a single sensible field; in the other we are concerned with a comparison between the contents of two different sensible fields.
>
> (Broad 1923, 351)

There is a phenomenal difference between seeing the second hand moving and seeing that the hour hand has moved. In the latter case we do not see the hour hand moving, we look at one time and see it at location *A*, and then we look again at a later time and see it at location *B*. We infer from our memory of the hand being at location at *A* and our seeing it now at location *B* that the hour hand has moved. We do not need to go through any such explicit line of reasoning in the case of our seeing the second hand moving. We see the smooth and continuous movement of the hand around the clock face. We have here an example of what Susanna Siegel (2006) has called a phenomenal contrast. One of us has argued elsewhere that based on this phenomenal contrast, we can also argue for a difference in representational content of these two experiences (Kiverstein 2010). In the case of the second hand we see movement, whereas in the case of the hour hand we do not. In the latter case we have a succession of experiences but no experience of succession. When we see the second hand moving, we do not just have a succession of experiences but we also have an experience of succession. When we see the second hand moving we see it occupying successive positions over time.

Chuard firmly rejects any distinction between the experience of succession and the succession of experiences. He argues that such a distinction rests on the failure of the following supervenience thesis (SUP):

> Temporally extended whole experiences supervene on successions of shorter experiences.

Chuard considers a number of arguments for rejecting SUP all of which turn on the whole temporally extended experience having some property *F* which a succession of experiences lack. He shows that none of these arguments succeed in establishing the failure of SUP. We want to briefly describe an alternative strategy for rejecting TPA which doesn't directly turn on mereological considerations of the kind Chuard considers. The strategy we wish to outline (of which there are several proponents in the recent literature) denies that the parts of temporally extended experiences are short-lived experiences. Thus Christoph Hoerl proposes that we conceive of experiences as process-like rather than as states. He writes:

> . . . my experience of the movement of the object from *A* to *B* is itself a process that unfolds over time. It starts at t_1, when I see the object departing from *A*, and it finishes at t_{10}, when I see the object arriving at *B*. On this view, it is simply a mistake to ask what I am aware of at t_5, and to expect the answer to show us, all by itself, how we can be aware of a sequence of events . . .
>
> (Hoerl 2009, 8)

The reason this is the wrong question to ask, according to Hoerl, is that experiences are processes that take up whole chunks of time. My experience taken at a particular instant doesn't qualify as a token experience of continuity or change because it is only a part of a longer process and *qua* part of a longer process it is too short-lived to qualify as an experience of continuity or change. These kinds of temporal limits are familiar from a number of findings in the psychology of time perception. A minimum amount of time must elapse, for instance, in between the presentation of successive distinct stimuli in order for us to experience the presentations as successive. Below this threshold the distinct stimuli will be fused by our perceptual systems and perceived as simultaneous. Similarly a certain amount of time must elapse in order for us to perceive the order in which successive stimuli have been presented. Below this threshold we will perceive the presentation of the stimuli as successive but not be able to determine the order in which they were presented.[12] What Hoerl adds to these familiar findings is the idea of a minimal duration that must elapse before we can experience any continuity or change. Any intervals of time that fall below this interval are too short, he suggests, for us to experiences any continuity or change.

Ian Phillips (forthcoming) defends a similar claim. Phillips writes: "The most basic facts about our experiential lives are facts about extended stretches of the stream of consciousness, and what is true at an instant is true only in virtue of that instant being an instant during such a period of experience." A little earlier he tells us "there are certain durations of experience which are explanatorily or metaphysically prior to their temporal subparts" (ms, 17). Like Hoerl, Phillips denies we can experience continuity or change at an instant, arguing instead that the basic unit of experience is some period

of time δt. It is only as part of the events unfolding over this period of time that we can experience an event e at t.[13]

Let us return now to SUP. The reason that SUP fails is that experiencing succession, continuity, and change requires a certain amount of time to have elapsed. Anything below this threshold is too short-lived for it to amount to an experience of succession when taken in isolation from the larger temporally extended experience of which it is a part. SUP asks us to treat short-lived experiences as the building blocks out of which a temporally extended experience can be constructed. Consider how Chuard attempts to account for our experiences of succession as like the experience of the second hand moving round the clock face. He says a series of single short-lived experiences can present us with smooth and continuous movement. Hoerl and Phillips by contrast argue for a reverse dependence of each short-lived experience on a temporally extended experience as a whole. We can experience the second hand moving round the clock only because experiences take up a long enough period of time for us to experience the hand changing position. Each short-lived experience presents us with change only because it is a sub-part of a longer experience that presents us with change. The crucial issue between Chuard and proponents of TOR like Hoerl and Phillips is, therefore, whether a short-lived experience can give us an experience of continuity and change. If we say there are short-lived experiences of this type, this is of no help to TPA since it denies that we ever experience continuity and change. If we want to resist TOR, we have to look elsewhere than to TPA to views which allow that we can have short-lived experiences that present us with continuity and change. In the next section, we look to the phenomenological philosopher Edmund Husserl for such an account.

5. Husserl on Internal Time Consciousness

Consider what it is like to see a goal being scored in football. We see the ball follow a path from the boot of the player into the net. According to the traditional interpretation of the Husserlian view we wish to consider next, this is something we can experience because at each moment in time an experience has a three-part temporal structure comprising: (1) a primal impression that is constantly updated; (2) a component Husserl called "retention" that has the function of keeping hold of an experience that has just passed; and (3) a component Husserl labeled "protention" that anticipates future experience. The retentional component gives us an experience that enables us to track the ball as it travels towards the goal. The protentional component gives us our sense of expectation or otherwise of the goal: if we are surprised this is because the goal is contrary to what we were anticipating. Crucially it is each momentary token experience that has this threefold temporal structure. We experience where the ball has just been from one moment to the next because each token experience keeps a hold of what was represented in the entire previous phase of experience. We can be surprised now because we have retained in our current experience some sense of what was anticipated in our previous experience.

Husserl denies that there is any match between the temporal structure of experience and its objects and it is this claim that brings his account of temporal experience into conflict with TOR.[14] Husserl distinguishes between the retention and the object of

retention, which is a past stage of a temporally extended event. The retention is a part of the structure of my token experience now, but what is retained is a part of an event that occurred earlier in time (or if we want to put this in tensed terms in the past). Something analogous applies to protention, which is a part of the structure of an experience occurring now at the time the experience is tokened, but what the protention is directed towards is a later stage of a temporally extended event. (Again we can put this in tensed terms and say that what is anticipated is what is not yet now or what lies in the future.)

We can think of the retention and protention as "constituting the temporal horizon" for each token experience. Whatever is present to us in experience is present as a part of spatio-temporal field. Consider first how every object we perceive is perceived as part of a spatial field that goes beyond what is strictly speaking given to us from the point of view we occupy as an embodied perceiver. Each object of experience has hidden aspects that can be sensibly given to us by varying the point of view we occupy in relation to the object. What is visible to us of the apple is its front and facing side, yet the parts that are not currently visible are also phenomenally present to us. Alva Noë (2012) has labeled this aspect of our experience "presence in absence" – the hidden and occluded parts of an object of experience are *absent* insofar as they are hidden or occluded. They are, however, parts of the object that we experience as spatially present to us, hence they have *presence in absence*. Husserl would say that the hidden and occluded parts of an object including its inside form a part of the object's inner horizon. The same is true of the phases of a temporally extended event that the retentional and protentional phases of an experience are intentionally directed towards. Suppose you hear someone singing the notes Do Re Mi Fa So La: when you hear the person sing "Mi" the notes "Do" and "Re" are no longer sounding but they are still present in your experience. They are however not present in the same way as the "Mi" you are now hearing – these notes have what we are calling (following Noë) presence in absence. Something analogous holds for the notes "Fa," "So," and "La" – when you hear the "Mi" you anticipate what is about to come next. These notes are not sensibly present in your auditory experience – they have presence in absence.

Husserl would have us distinguish between the temporal structure of experience and its objects because retention and protention are parts of momentary short-lived experiences. What these components of short-lived experience refer to however is not a single instant in time as with TPA. The retention and what is retained and the protention and what is anticipated are not present in experience simultaneously. Retention refers to the earlier phases of a temporally extended event and protention refers to the later phases of an event that we anticipate are to come. In this way, a momentary experience can include in its content the earlier and later parts of a temporally extended event. Crucially, however, these earlier and later parts of an event are present in experience as the temporal horizon within which the current phases of an event are experienced as embedded. Retention and protention supply the temporal context against which new contents (in the form of the primal impression) are always presented. According to Husserl, these aspects of experience are also responsible for bringing about flow within consciousness. The temporal context of a momentary experience is continuously changing with each renewal and updating of retention and protention. The primal impression that is at the core of each experience is likewise undergoing constant

refreshing and refashioning. It is this continuous renewal of the retention-primal impression-protention structure that creates the appearance of dynamism and flow within consciousness. Husserl is very clear however that we need to distinguish this temporality within consciousness from the temporality of the objects of consciousness. This of course is very much in opposition with the core claim of TOR that experience and its objects share the same temporal structure.

6. Should We be Naïve Realists about Temporality?

In our introductory section we introduced the issue of whether temporal properties as they appear to us should be thought of as primary or secondary qualities. In this section we will briefly return to this issue since it will prove important for the position we end up with concerning the debate between TOR and its opponents. Let us begin with the views that Husserl challenges that take there to be match between the temporal structure of experience and its objects. On one version of this view, there is such a match because, at least in the good cases in which we are not the victim of an illusion or hallucination, experience puts a subject in relation to externally existing objects. There is a match then between the temporal structure of experience and its objects because experience is a relation to an object on this view. This is a *naïve realist* view of temporal properties that says at least when all goes well in experience there is a match between the temporal properties we experience and the temporal properties of objects in the world.

Husserl's view, by contrast, makes a distinction between the temporal properties of experience and the temporal properties of the objects of experience. Thus, he shares something in common with recent literature on tensed properties (e.g., the property of being present) which take it that these properties are tied in some way to our subjective or egocentric point of view[15]. We will call these philosophers "detensers." They argue that our experience of time passing and our experience of the presentness of now derive from our subjective point of view and not from time objectively passing or from events having the property of being present. Change on these views is just a matter of an object having different properties at different times. If change requires the coming into being and passing out of being of events this is entirely due to a dynamism to be found within experience. Detensers usually take experience to supervene on brain processes. They argue that the dynamism we find in experience is a product of the active interpretation of our brains and is projected onto the world. Whatever continuity and change we find in experience is not passively imposed on the brain by the world but is actively constructed. Robin Le Poidevin sums up the detensers view well when he writes:

> A-theoretic properties are not in the world, but are *projected* on to the world in response to certain features of our experience. This would be closely analogous to projectivist views of secondary qualities: the world itself is not colored, but certain properties of objects induce in us sensations which cause us to ascribe colors to them (see, e.g., Boghossian and Velleman 1989).
>
> (Le Poidevin 2007, 95)

Husserl would agree that there is a dynamism and flow in experiences that has its origin in our subjective perspective on the world. We talked above of retention and protention as "constituting" the temporal horizon against which new contents (or what Husserl calls the "primal impression") are always presented. The technical phenomenological notion of "constitution" means in part that we should think of this temporal horizon as *constructed* in experience rather than as a passive imposition on experience by the world. We should, however, be careful not to overstate the extent of Husserl's agreement with detensers. Detensers appeal to psychological and cognitive mechanisms to account for our living in a common now (Callender 2008) or our experience of the passage of time (Paul 2010). Husserl's claim is that dynamism and flow have their origin in the retentional and protentional structure of temporally extended consciousness and it is the relations between retention, primal impression, and protention which constitute the flow of experience. Dan Zahavi makes this point well when he writes that: "Inner time consciousness is not an object occurring in time but neither is it merely a consciousness of time; rather it is itself a form of temporality" (2007, 466). Inner time consciousness as a form of temporality makes it possible for temporally extended objects to appear or show themselves to me as a subject of experience. Where detensers look to explain tense in terms of sub-personal mechanisms, Husserl is offering a phenomenological description of personal-level subjective experience. His lectures on time consciousness are attempts at answering the question of how it is possible for us to be conscious of temporally extended events or objects such as melodies or birds in flight.

A theoretical assumption of our chapter is that we ought to take phenomenological descriptions of temporal experience at face value.[16] This means taking seriously the "how-possible" question Husserl raises. We are not supposing however that Husserlian phenomenology exhausts all there is to say about this question. Phenomenology can help guide scientific research by, for instance, ensuring that the conceptual framework that informs scientific theorizing is consistent with our best phenomenological descriptions of experience. Phenomenology can provide science with constitutive descriptions of experience that then act as guides as to what needs to be explained. However, the descriptions of experience phenomenology gives in answering how-possible questions also stand in need of explaining. We must make intelligible how they fit with what our best theories in the biological and cognitive sciences tell us about the nature of the human mind. These sciences provide us with descriptions of the sub-personal, functional, or neurobiological mechanisms that enable us as persons to experience and act as we do. There is, therefore, a constraint that runs from our best science of the mind back to phenomenology that our best phenomenological descriptions of experience cohere with our best science of the human mind.

We agree with detensers that our experience of continuity and change has its origins in cognitive and neural mechanisms that actively interpret and construct whatever temporality we find in the world. We deny, however, that sub-personal, neural explanations exhaust all there is to be said about our experience of continuity and change. We take phenomenology to answer how-possible questions that simply don't arise for the cognitive scientist. Phenomenology provides answers to these questions, answers which in turn provide a constraint on our scientific theorizing.

Ian Phillips argues that the how-possible question that concerns the Husserlian phenomenologist does not arise once we adopt the naïve realist position he favors. Phillips writes:

> it is not clear why we should accept that a genuine how-possible question arises. Cassam suggests that: "(t)o ask a how-possible question is to ask something which looks impossible given other things that one knows or believes is nevertheless possible" (2007, 1). But as yet we have no grounds for thinking that time consciousness is in any way mysterious, let alone for thinking it impossible.
>
> (Phillips ms., 4–5)

Phillips is of course right that once one accepts the view that there is a match between the temporal structure of experience and its objects then no "how-possible" question of the kind phenomenology raises needs to be addressed.[17] Phillips is effectively arguing that the Husserlian alternative to TOR we have outlined above is based on a philosophical question that from the perspective of TOR is unmotivated. However, what is really at issue here is whether it is possible to experience continuity and change in a short lived experience. Husserl shows how this is indeed possible. If TOR can be defended, Phillips dismissal of the how-possible question would seem entirely justified but the defensibility of TOR is what Husserl disputes.

We saw above how a relation of mutual constraint holds between phenomenology and cognitive science. Our best phenomenological descriptions of experience guide scientific explanation as to what needs to be explained, but at the same time phenomenological descriptions of experience must be consistent with our best science of the day. Our strategy in the remainder of the chapter will be to put this methodology into action. TOR and the Husserlian view offer competing phenomenological descriptions of experience. We will attempt to decide between them by asking which of these competing descriptions fits best with what we know from neuroscience about the perceptual processing underlying conscious experience.

7. Explaining Postdiction

The philosophically perplexing temporal illusions we introduced above are all examples of postdictive phenomena in which an event occurs at time *t* which then modulates how a stimulus is experienced at a time up to a few hundred milliseconds prior to *t*. Consider backward masking as another example of this kind of postdictive effect.[18] In this experimental paradigm subjects are briefly shown a target stimulus followed by a second masking stimulus. When the masking stimulus is presented 50–80 milliseconds *after* the target subjects will often report not seeing the target. A stimulus that would be highly visible when presented on its own is rendered invisible by a stimulus presented at a later time. There are a number of possible explanations for postdiction, some of which challenge TOR, as we will see. The problem for a proponent of TOR is therefore to say why we should prefer their explanation of postdiction over its rivals.

Consider first what Dennett (1991) calls the Stalinesque interpretation of postdiction. It claims that before any conscious experience can arise all information has to be gathered and presented at a "show trail" where a judgment is reached about what

happened, often based on misinformation (such as false testimonies and evidence that has been tampered with). The Stalinesque interpretation seems to be inconsistent with TOR since it allows that a momentary experience can, all in one go, represent a temporally extended event.[19] Michael Tye (2003) defends a version of this claim. He writes: "On this account, the experience of *A* followed by *B* casts an eye backward as it were, at what preceded it. The glance is all-in-one, however. It takes in a succession in the specious present – the time that, for the experiencer, is *now*" (2003, 88). The Stalinesque interpretation denies that there is a match between the temporal structure of experience and its object: while the object of our illusory experience is extended through time, the experience is located at the instant after the show trail.

A proponent of TOR will deny that we can, in a single glance, experience *A* being followed by *B*. They will argue that to experience *A* followed by *B*, sub-personal processing must take place, the stages of which are organized in a way that matches the temporal structure of experience. We can, of course, ask what a subject experiences at an instant but we should not treat this as a basic fact because what the subject experiences at an instant is always a part of a longer, temporally extended experience. Thus, the facts about what a subject experiences at an instant will depend upon the longer experience of which this instant is a part. This interpretation of postdiction differs from the Stalinesque interpretation in that it accounts for the influence of later stages of experience on earlier stages in terms of the temporally extended, process-like character of experience. Following Phillips (2011) we will call this the *extensionalist* account of postdiction.

Phillips rejects the Stalinesque interpretation of postdiction for "pulling apart the temporal structure of experience from the temporal structure of objects presented." He argues that postdiction establishes precisely the dependence of each instant of experience on the whole temporally extended experience of the kind we have seen proponents of TOR defend above. In backward masking we cannot treat a token experience of a target presented on its own and the token experience of the target when followed by the mask as tokens of the same type. They are different token experiences because the temporal context in which they occur is different.

Does this mean a victory for TOR? Such a conclusion would be premature since there is a further possible explanation for postdictive phenomena which Dennett has labeled the Orwellian interpretation. It rejects the idea of a temporal lag, a delay after which a Stalinesque show trail can occur, in explaining postdictive phenomena. Instead, it is argued that we have an ever-so-brief experience which subsequently gets forgotten. What we report is instead a constructed memory that rewrites history in the same manner as the Ministry of Truth in Orwell's 1984. Applying this idea to the phi phenomenon we get the following explanation. First, we have an ever-so-brief experience of the stationary dot at location *A* and then we experience a brief interval in which the screen is blank. However by the time we get to report on our experience, these initial experiences have been forgotten, overwritten by a false memory of the dot moving from location *A* to *B*. This is what we recall when probed for report, but the memory we report on is a confabulation: subjects misremember what happened; they experienced one thing but sincerely report experiencing another.

Dennett argues that the debate between Stalinesque and Orwellian accounts is in the end not substantive but "merely verbal."[20] At the same time, he suggests in passing

that the Orwellian account has at least one advantage over its rival. This is that it avoids a commitment to potentially costly processing delays of up to 200 milliseconds, which might really matter when it comes to fight or flight responses based on experience. This is also a point Rick Grush has repeatedly stressed in comparing his trajectory estimation model of temporal representation with the Stalinesque model (or what Grush calls "the smoothing model") that must posit a time lag of around 80–100 milliseconds in visual processing (Rao et al. 2001). Grush tells us this delay in perceptual processing is costly in two respects:

> First, there is the computational cost involved in smoothing over filtering. More data are involved and arriving at the smoothed estimate involves more processing than a merely filtered estimate, since the filtered estimate must be computed in route to computing a smoothed estimate. Second, there is a cost to the organism in terms of behavioural timeliness if percepts potentially crucial for action are delayed.
>
> (Grush 2005, 215–216)

We will unpack some of the technical vocabulary Grush employs in this passage in the next section of our paper. The crucial difference we want to highlight for now is that while TOR is committed to a potentially costly processing delay (of up to 100 milliseconds on the Rao et al. model)[21], the Orwellian account is only committed to what Grush describes as an "openness to revision" over a similar period of time. We have now pinned down a precise empirical question that might decide between TOR and the Husserlian view because we now have a genuine alternative to the extensionalist account of postdiction the proponent of TOR must endorse. He must argue that the later parts of a temporally extended experience can influence and modulate earlier parts because our brains introduce a delay in processing before reaching any firm interpretation of events. By positing such a delay, a proponent of TOR can account for temporal illusions in a way that is consistent with the idea of a match in temporal structure at the level of contents and vehicles. The Orwellian account, by contrast, denies that the brain introduces any such delay in perceptual processing, arguing instead that perceptual processing is subject to revision and reinterpretation over a short period of time. In the next section of our chapter we explore the latter possibility in a little more detail and consider the relationship between the Orwellian interpretation of postdiction and the Husserlian view of temporal experience.

8. The Predictive Inference Model

So far we have followed Dennett in presenting the Stalinesque and Orwellian interpretation as a theory of short-lived temporal *experience.* Might it be wiser to bracket questions about consciousness in looking for an explanation of postdiction? Postdiction is of course an experiential phenomenon, but we think it is important to note that the mechanisms that account for postdiction may not also explain consciousness. Indeed the explanation of postdiction we think very likely lies in unconscious perceptual processing. The neuroscientist David Eagleman has colorfully talked of us all "living in the past" by 100 milliseconds or so. This is a claim about personal-level experience. We

suggest that as such Eagleman might be making a mistake. The data on which he bases this claim might be better understood as relating to unconscious perceptual processing. Phillips is a little more careful, and we think advisedly so. In presenting his view of the cutaneous rabbit he writes:

> What does the subject feel at $t + \delta t$ in the second trail? If we want to resist the answer being, *the tap at the wrist* (as the Orwellian account claims) then we seem forced to claim that δt is a period of at least 240 ms (plus further processing time) – enough time for the tap to be re-localised in the light of information about the subsequent taps.
>
> (Phillips forthcoming, 18)

Later Phillips adds that nothing in his account "demands a delay" (forthcoming, 19) by which we take him to mean a delay in experience. The 240 milliseconds that is required to determine whether a tap is felt at the wrist or as belonging to a longer series of tactile sensations that feel like an animal hopping up the arm, is a delay in unconscious perceptual processing.

In arguing against Grush's trajectory estimation model (more on which in a moment), Phillips claims that "Grush's picture is . . . an Orwellian account since it posits conscious experiences which leave no lasting cognitive trace" (Phillips forthcoming, ms. 15). Elsewhere he argues against such a position on the grounds that it is committed to an "is/seems" distinction for experience, which, as he rightly points out, barely makes sense (Phillips 2011, 395). If Phillips' interpretation is correct, Grush is committed to saying that a subject *really* experienced one thing (a stationary dot at location and then a blank screen), but that it later *seems* to the subject as if he experienced something quite different (a single dot moving from *A* to *B*).[22] However, if Phillips can avoid an implausible delay in experience by arguing that postdiction happens in unconscious perceptual processing, this is a move that ought to be equally available to Grush. Indeed Grush hardly ever talks about "temporal experience" or "temporal consciousness" in presenting his model. Instead he talks about "perceptual" or "temporal" representation, which is of course ambiguous between conscious and unconscious perception.

Now consider what a version of the Orwellian interpretation would look like that took postdiction to occur in unconscious perceptual processing. This is most definitely not what Dennett had in mind when he presented the Orwellian interpretation. Dennett took the tampering with experience to occur in the memory of what we experienced, and it is our memory of what we experienced at an earlier time that gets revised over time. The interpretation of postdiction that we take Grush to be committed to takes the process of constructing a revisionary history in postdiction to be entirely unconscious. It is not experience that gets rewritten in the process of constructing a memory of what happened, but instead it is unconscious perception that is rewritten over a brief interval. Grush needn't claim that there is an experience that is quickly forgotten, thus introducing an appearance–reality gap into experience. Instead all of the revisions take place in unconscious perceptual processing. If we continue to read Grush as defending the Orwellian interpretation, the difference between the Stalinesque and Orwellian interpretation of postdiction does indeed begin to look like merely a verbal difference, as was argued by Dennett. Both models take postdiction to be the result of tampering

with perceptual hypotheses formed in preconscious perceptual processing. Still, as we will explain shortly, there remains a dispute between Grush and Phillips' extensionalist theory of postdiction which we described in the previous section. We will henceforth set aside the debate between Stalin and Orwell for the remainder of our chapter and turn our attention to this dispute instead.

We will begin with a brief review of Grush's (2005, 2007, 2008) trajectory estimation model of temporal representation. Grush argues that perceptual representations are the products of the three processes he labels: (1) filtering; (2) prediction; and (3) smoothing. Any given perceptual representation of the environment will be a product both of top-down expectations formed on the basis of the perceiver's knowledge of statistical regularities relating to a represented domain and of bottom-up generated sensory inputs. Sometimes our top-down expectations will be in conflict with driving perceptual input. When such conflicts arise the process of *filtering* kicks in, resolving the conflict either in favor of our knowledge based expectations or by overriding our expectations in favor of sensory inputs. *Prediction* is the process by which the perceptual system infers a model of the environment over a period of time from a model of the environment at a particular instant. Given knowledge of how the environment is at a time t and knowledge of how a process typically behaves over time, the perceptual system can produce estimates of how the environment was at earlier times and how it will be at later times. It is this part of the process that looks to play the same functional role as the structures of temporal experience Husserl describes in terms of retention and protention. (We will say a little more about this below.) Finally, by means of *smoothing*, later estimates can be combined with earlier estimates so as to update an original model of the environment over a period of time. It is at this stage that rewritings of past perceptions can occur, such as we find in postdiction. The perceptual system can in this way produce perceptual representations that keep track of the latest developments with regards to a represented domain. Insofar as the perceptual system is using what it has learned about the world to model things in advance of them happening, it can also keep abreast of the facts without having to incur any processing delays. This, we suggest, is a major advantage of Grush's model over Phillips extensionalist account which must posit a substantial delay in preconscious perceptual processing.

Grush's trajectory estimation model bears a close resemblance to computational theories in neuroscience that take perception to consist in predictive inference (Friston and Stephan 2007; Hohwy 2007; Clark forthcoming). We will henceforth refer to these models as "predictive inference models." According to these approaches, a central problem which our brains face in perceiving the world is that of working out what things in the environment are the causes of the sensory inputs to which our brains respond. One way of solving this problem would be to reason backwards from effects (sensory inputs) to the causes of those effects. Another strategy is to work with a model of the environment with the highest prior probability and then predict the causes of sensory input based on this model. If the prediction matches sensory inputs coming in from the world then the model is confirmed. If not, something analogous to Grush's process of filtering is called for and the model must be updated or the sensory signals that conflict with the model must be ignored. What we perceive on any given occasion will be determined by the cascade of models that best predicts sensory inputs at each of the levels in perceptual processing from the fast processing of low-level features, such

as edges and contours, to slower high-level processing associated with categorization. Perceptual processing has as its function the minimization of error at each of these levels of processing; errors occur when there is a mismatch between the internal models that encode our expectations about what is likely to happen in the world and actual inputs caused by things in the world. Perception, then, is largely a matter of what we expect to see, with sensory input from the world providing feedback that allows us to correct our expectations when we go wrong. If our internal models have been updated correctly based on past mismatches between predictions and input, the future likelihood of these models resulting in faulty predictions will be minimized.

So far we have been assuming perception is all about error minimization, but the process of error detection also involves uncertainty. Sensory inputs always come accompanied by a certain amount of noise and this noise must be factored in when determining whether or not to revise an internal model based on an error signal. A noisy prediction error should not lead to revision since it is likely to be the upshot of noise rather than an inaccuracy in one's model. Thus, an important aspect of what Grush calls "filtering" will be working out the accuracy and precision of the sensory signal. We do not need to concern ourselves here with the details of how our brains might determine the reliability of an error signal. What we wish to stress for now is that in addition to the accurate prediction of sensory inputs, a sensory system must also be able to work out what variation in sensory input is to be expected and to what extent this variation is due to noise in the signal. It must be able to determine the degree of uncertainty to attach to sensory signals themselves.

There are then at least three sources of variability in the contents of perceptual representations. When predictions are accurate the predicted inputs are cancelled out leaving the error signals to ascend to higher levels in the system. Thus, the first source of variability comes from the difference in the overall state of the sensory system before and after predictive inference. A second source of variability will come from the perceiver's own agency – her eye movements and other bodily movements for instance. Perceptual inference is *active* with the agents own movements bringing about changes in sensory input that best allow the brain to minimize prediction error. Finally, as we just noted, a perceptual system must be able to optimize the precision of the error signals by controlling the relative influence of error signals *vis a vis* the expectations encoded in its internal models. The relative influence of these error signals is also continuously in flux, leading to variation in the evidence on which perceptual inference is based in each cycle of processing.[23]

We suggest that the rewriting of the contents of perceptual representation that we find in postdiction is due to these types of variation. If this is right, our brains do not need to introduce delays in order to take advantage of the latest news when forming representations of the world. We can account for postdiction instead in terms of this openness of unconscious perception to revision in ways that reduce error in prediction.

We saw above how Phillips uses postdiction to argue for TOR. He takes postdiction to establish the dependence of the parts of an experience on the temporally extended experience taken as a whole. We can see now however that this dependence of the part on the whole does not necessarily decide in favor of TOR. The version of the Orwellian interpretation we have been developing can also claim that each moment of perception

(and by extension each moment of experience) has the content it does in part because of its surrounding temporal context. Husserl tells us that each primal impression has a meaning which is influenced by the previous phase of experience that is retained in consciousness and by the future course of experience the subject anticipates or protends. Each moment of experience has the meaning or significance it does because of earlier moments of experience that have presence in absence. The predictive neural mechanisms we have just described would seem to be able to do some of the same work as retention in contributing to the meaning of each primal impression. To see this, consider again our favorite example of postdiction, apparent motion. When we perceive a single dot moving from location *A* to *B* this is because the processing of the second dot at location *B* has been influenced by what has gone before. *A* and *B* are represented as a single moving dot rather than as two stationary dots because of the temporal context in which both are presented. Central to the idea of predictive inference is that our brains do not try to infer causes from effects but instead try to construct models that can best predict the causes of sensory input. What happens in apparent motion is that this model is updated so that now the brain's best interpretation of what is and has been happening is that the stimulus we are seeing on the screen is a single moving dot. The illusion we undergo is the brain's best bet at working out what the most likely cause of it sensory inputs might be given its statistical knowledge of how moving stimuli typically behave and the noisy sensory evidence that is currently available.

At this point it might be objected that while this explanation may work for some postdictive phenomena it does not work for them all. Indeed, in some cases the idea of a delay in processing seems to make much more sense than an explanation in terms of updating of generative models based on error signals. Consider the flash lag effect as one such example: the illusion consists of a moving stimulus *A* and a flash, which is perceived to lag behind the moving stimulus even though in reality it is presented at the same location as *A*. It might be thought that the predictive inference interpretation is committed to what has come to be called the "motion extrapolation explanation" of the flash lag (Nijhawan 1994). According to this explanation, the flash lags the moving stimulus because we perceive it where we extrapolate it should be given the amount of time it takes to process the moving stimulus. We perceive the moving stimulus where we predict it to be, which is slightly ahead of the flash based on processing delays. The shortcoming of this explanation is that if *A* stops its movement at the time of flash, they are perceived to occur at the same location (there is no extrapolation) (Eagleman and Sejnowski 2000). Another shortcoming is that when the moving stimulus changes the direction of its movement at the time of flash, the flash lags again, but now in the direction of the new movement and not of the original movement as implied by the motion extrapolation explanation (Eagleman and Sejnowski 2000).This last finding is particularly interesting because the change in the direction can only be determined by the location of the moving stimulus after the flash occurred. The best explanation of this postdictive phenomenon would seem to be one that introduces a processing delay.

It is not however immediately obvious that a predictive inference model must endorse the motion extrapolation explanation of the flash lag illusion. Such models must agree

that the trajectory of *A* relative to the position of the flash must be estimated and this looks a lot like an extrapolation. However these models will also allow for the correction and *updating* of this estimate over a short interval in a way that would seem to, at least potentially, account for the data that motivates the appeal to a processing delay. Consider two trails, one in which *A* continues moving in the same direction at the time of the flash and the second in which *A* changes its direction at the time of the flash. The location we perceive the flash at will depend on the direction of *A*, but this is something the brain can know only after the flash has occurred. Predictive models can allow, however, that the brain makes use of the information that *A* has changed direction in forming its percept. The brain generates the stimuli it expects based on past learning but it also corrects its expectations on the basis of feedback in the form of incoming stimuli. The information that *A* has changed direction forms a part of the incoming feedback the brain can use to adjust its estimation of what happened at an earlier time. It is not only prediction (or extrapolation) that is doing the work in this account of postdiction. Equally, if not more important, is the ongoing process of revising and correcting perceptual interpretations of the most likely causes of sensory inputs based on prediction error. We can conclude then that accounts of postdictive phenomena that appeal to predictive inference do equally as well as their rivals, such as the extensionalist theory, that posit a processing delay. Both accounts allow that some time must pass for postdictive phenomena to occur. What is in dispute is whether during this time there are models of the world under construction that are subject to revision or whether instead the brain waits some time before reaching any firm conclusions about what just happened.

Conclusion

It might appear that all of this is rather orthogonal to the debate between TOR and its opponents, but this is not the case. To see why not we need to now make a distinction between two versions of TOR:

> TOR-c: The vehicles and contents of conscious temporal perception share an identical temporal structure.
>
> TOR-u: The vehicles and contents of unconscious perceptual states share an identical temporal structure.

The extensionalist theory of postdiction claims that TOR-c is true, but we take it as neutral on the truth of TOR-u. Once the idea of a processing delay is introduced there is no requirement that the unconscious perceptual representations have the very same temporal properties as the events they represent. Indeed, the very idea of a processing delay implies there is no match between the timing of representations and the time of the events represented in perception. Suppose we accept a delay in processing of the kind a proponent of TOR-c accepts – *A* occurs at t_1 and *B* at t_2, but the representations of a single moving dot occurs only after *A* and B have both been presented. If TOR-u

holds then *A* must be processed first followed by *B*. However, once we insert a delay there is no longer a requirement that there is a match in the order in which items are processed and the order in which they are represented, since this could all be sorted out during the processing delay in something analogous to an integration stage. TOR-c thus seems to be independent of TOR-u.

The predictive inference theory we have outlined rejects TOR-u. We have suggested that it may also explain how experience could have a threefold retention-primal impression-protention structure of the type we find described in Husserl's writings on time consciousness, which is incompatible with TOR-c positing as it does short-lived temporal experiences that present us with continuity and change. We suggest then that the predictive inference account, in common with the Stalinesque theory, will reject both TOR-u and TOR-c.[24]

We began our chapter by asking whether temporal properties as they are experienced by us are primary or secondary qualities. We cannot claim to have fully answered this question. Instead, we have shown how experience would need to be organized if temporal properties were to turn out to be primary qualities: experiences would need to have matching temporal structure at the level of contents and vehicles. We have outlined a rival theory which denies that there is any such match. On this account temporality as we experience it is best understood as something that is constructed in experience, rather than a passive imposition of the world onto our minds. We have also outlined a theory of neural mechanisms which would seem to support the latter view. We saw in the last section how the active interpretation and reinterpretation of sensory input is absolutely central in this account of temporal representation. Moreover, insofar as the predictive inference model requires us to reject TOR-c, it supports a view of perceived temporality as a projection of our minds onto the world. Temporality as it is experienced by us is the product of the brain's active interpretation of events in the world. Both the views we have outlined can agree that unconscious brain processes enable us as perceiving subjects to stand in a perceptual relation to temporally extended events. The important difference is that while on TOR-c this perceptual relation involves a match between our minds and reality, on the Husserlian view this perceptual relation is an active construction.

Notes

1 Functional Magnetic Resonance Imaging technique (fMRI), for example, has a temporal resolution of a few seconds. Hence it cannot track the details of fast movements or the processes that relate to understanding speech. Techniques that have a better temporal resolution, such as magnetoencephalography (MEG), on the other hand, often have poor spatial resolution especially on the deeper brain structures.

2 This relates closely to the topic of when we perceive a stimulus. Interestingly, two of the main methods for determining the timing of perception experimentally do not always yield the same results. The two methods are reaction time studies, in which subjects are asked to react as fast as they can to predetermined stimulus, and simultaneity or temporal order tasks, in which subjects are asked to judge whether distinct stimuli occurred at the same time or the order in which the stimuli appeared. When the properties of the stimulus are varied the subjects' results in these tasks are not always influenced in the same way. It seems

likely that there are many mechanisms that use different types of information – different neural time stamps – to determine the timing of perception. This would explain why we get different answers depending on the tasks and questions that we pose to the subjects (Jaśkowski 1996).

3 Nelson Goodman (1978) as well as Daniel Dennett and Marcel Kinsbourne (1992) describe a special case of phi phenomenon called the color phi phenomenon. To induce this phenomenon, the two stimuli need to be of different colors, and when the two stimuli are separated by a duration of the right interval subjects report seeing a single moving stimulus that changes color around the middle of its illusory path. We ignore this twist in what follows, however, because it doesn't raise any additional questions that aren't already raised by the original phi phenomenon. Dennett and Kinsbourne (1992, 186) seem to agree that changing color is merely a complication as they write: "The same question [concerning the temporally perplexing nature of phi phenomenon] can of course be raised about any phi, but the color-switch in mid-passage vividly brings out the problem".

4 Dennett and Kinsbourne describe a second possibility using the example of a movie in which the audio and visual tracks are running out of synch. This situation can be amended simply by fitting the audio track to the visual track by first determining the temporal structures of events in both tracks and then fitting the two tracks together. Dennett and Kinsbourne call this "content-sensitive settling," and they suggest that the brain may use a comparable strategy to infer the temporal properties of events. On this view the temporal properties we experience are the result of an inference "drawn by comparing the (low-level) content of several data arrays . . . having drawn inferences from these juxtapositions of temporal information, the brain can go on to represent the results in any format that fits its needs and resources – not necessarily in a format in which 'time is used to represent time'" (1992/1997, 151).

5 See Dennett and Kinsbourne 1992/1997; Eagleman and Sejnowski 2000; Johnston and Nishida 2001; Grush 2007.

6 For a defense of such a view in the neuroscience literature see Vibell et al. (2007).

7 Again, an explanation is owed as to how the brain can know about the presentation of the dot at location B at the time when we *experience* the dot begin to move from location A. A natural way to explain this is to introduce a processing delay in between the presentation of visual stimuli and the occurrence of the neural events that underpin a visual experience of those stimuli. We'll return to this possibility in sections 7 and 8.

8 This inference might be resisted. It might instead be argued that an experience and its objects can have a matching temporal structure, even though there is a mismatch in the temporal structure of an experience's representational vehicle and its content. It is not immediately obvious to us however how an experience could have temporal properties that match those of its objects without the vehicle also having temporal properties that match the object it represents. We will henceforth assume that this is not possible.

9 Some philosophers have argued that short-term memory might play a role in our experiences of continuous or changing events. See for instance Le Poidevin (2007, ch. 5) and Phillips (2010) for two different versions of this view. We will not have much to say about the role of short-term memory in this paper but see Kiverstein (2010) for some critical discussion of Le Poidevin.

10 We leave it open for now how to conceive of the relationship between the temporally extended experience and the individual momentary experiences of which it is composed. We will return to this question later in Sections 4–6.

11 Chuard is of course familiar with this Jamesian point and much of his paper is dedicated to carefully rebutting several different arguments that aim to drive a wedge between experiencing succession and a succession of experiences, more on which in a moment.

12 For a review of these findings see Wittmann (1999) and Pöppel (1988). For an interesting discussion of these findings in relation to the metaphysics of time see Callender (2008).

13 Hoerl and Phillips' claim that episodes of experiencing have some minimum duration receives support from psychophysical studies such as Robert Efron (1970, 57). Efron argued that "visual perceptions have a minimum duration which lies in the range of 120–240 msec, and that auditory perceptions have a minimum duration which lies in the range of 120–170 msec."

14 Husserl does not deny that content and object can share a matching temporal structure. He accepts that experiences can have contents that present us with continuity and change. What he denies is that experiences have to be extended through time in order to have contents that present us with temporally extended events.

15 See, e.g., Le Poidevin (2007); Callender (2008); and Paul (2010). We do not mean to attribute to proponents of TOR a tensed view of reality. Hoerl (2009), for instance, argues that experiences have tenseless contents that can match reality because reality is likewise tenseless. However, it is an open question how you can explain the dynamism and flow within experience by invoking experiences with a tenseless content that match a tenseless reality. Husserl's theory would seem to have an advantage here in that it is specifically designed to account for the temporality or flow we find within consciousness.

16 Clearly this is an assumption in need of defense. See Kiverstein (2012) for an attempt at discharging this obligation. We cannot say much about this here but see our discussion of Chuard's TPA above for an application of this principle and the short discussion to follow will say something further about what it means in practice.

17 From a phenomenological perspective it is precisely the validity of a naïve realist theory of perception that ought to be questioned. The naïve realist theory of perception is part and parcel of what Husserl labels the natural attitude, and his phenomenological philosophy begins with a suspension of the natural attitude. There are of course many important critical questions one could and should raise about Husserl's phenomenological method. Many of these questions have been extensively discussed by phenomenological philosophers who followed Husserl, most notably Heidegger and Merleau-Ponty. Phillips of course has strong arguments for rejecting the how-possible question but all of these arguments are motivated by his commitment to TOR which from a Husserlian perspective is at best questionable.

18 For a review of this paradigm see Breitmeyer and Öğmen (2006).

19 We follow Tye and Phillips here who both maintain that the Stalinesque interpretation of postdiction is inconsistent with TOR. It should be mentioned, however, that it is not obvious that this is correct. As an analogy, consider the case that some portable CD players do not play the song on the headphones immediately as they read it from the disc, but instead the reading slightly precedes the playing. One reason for this is that the player can interpolate the small parts that are missing due to scratches based on the information before and after the missing part. This means that the postdictive effect takes place before the sounds are heard (Stalinesque revision) and yet the sounds are played one by one and in the order in which they are in the record (TOR).

20 Dennett dismisses the debate as merely verbal because both the Stalinesque and Orwellian interpretations are consistent with what the subject says and does, and he seems to want to infer from this that there is no fact of the matter about which of these theories is correct. Any decision about the "moment of processing in the brain" we judge to be "the moment of consciousness" would be an arbitrary decision according to Dennett. Since this is what the debate between the Stalinesque and Orwellian accounts is about, Dennett concludes there is nothing to decide between these two accounts. Perhaps, however, Dennett is relying

on a verificationist view of consciousness, which we should resist here. In any case we think there is more to be said.

21 Even longer delays are required to account for the 200 millisecond inter-stimuli intervals in apparent motion experiments.

22 We disagree with Phillips' interpretation of Grush's view, however, because the interpretation does not take into account that in Grush's case the Orwellian tampering does not concern memories but (possibly unconscious) perceptual states. That is, our perception of a stationary stimulus followed by a blank screen is rewritten by the perception of a moving stimulus. Furthermore, both the original and the latter experiences are equally consequences of emulation processes.

23 We are basing our discussion here on Hohwy (2011, § 5).

24 We do not take ourselves to have exhausted all of the positions it is possible to occupy in this landscape. One could defend TOR-u but reject TOR-c for instance, a combination of views we have not considered since our concern has been with TOR-c. See Arstila (ms) for a theory of apparent motion with this profile.

References

Arstila, V. (ms). *Apparent Motion for Philosophers*.

Boghossian, P. and Velleman, D. (1989). Colour as a Secondary Quality. *Mind* 98, 81–103.

Breitmeyer, B.G. and Ö men, H. (2006). *Visual Masking: Time Slices Through Conscious and Unconscious Vision*. Oxford: Oxford University Press.

Broad, C. (1923). *Scientific Thought*. New York: Harcourt Brace.

Callender, C. (2008). The Common Now. *Philosophical Issues* 18: Interdisciplinary Core Philosophy, 339–381.

Cassam, Q. (2007). *The Possibility of Knowledge*. Oxford: Oxford University Press.

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

Clark, A. (forthcoming). Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. *Behavioral and Brain Sciences.*

Dennett, D.C. (1991). *Consciousness Explained*. London: Penguin Books.

Dennett, D.C. and Kinsbourne, M. (1992/1997). Time and the Observer. *Behavioral and Brain Sciences* 15(2), 183–247.

Eagleman, D.M. (2008). Human Time Perception and its Illusions. *Current Opinions in Neurobiology* 18 (2), 131–136.

Eagleman, D.M. and Sejnowski, T.J. (2000). Motion Integration and Postdiction in Visual Awareness. *Science* 287(5460), 2036–2038.

Efron, R. (1970). The Minimum Duration of a Perception. *Neuropsychologia* 8, 57–63.

Friston, K. and Stephan, K.E. (2007). Free Energy and the Brain. *Synthese* 159, 417–458.

Geldard, F.A. and Sherrick, C.E. (1972). The Cutaneous "Rabbit": A Perceptual Illusion. *Science* 178(4057), 178–179.

Goodman, N. (1978). *Ways of Worldmaking*. Hassocks, Sussex: Harvester.

Grush, R. (2008). Temporal Representation and Dynamics. *New Ideas in Psychology* 26, 146–157.

Grush, R. (2007). Time and Experience. In T. Müller (ed.). *Philosophie der Zeit* (1–18). Frankfurt: Klosterman.

Grush, R. (2005). Internal Models and the Construction of Time: Generalizing from State Estimation to Trajectory Estimation to Address Temporal Features of Perception, Including Temporal Illusions. *Journal of Neural Engineering* 2(3), S209–218. doi:10.1088/1741–2560/2/3/S05

Hoerl, C. (2009). Time and Tense in Perceptual Experience. *Philosopher's Imprint* 9 (12).
Hohwy, J. (2011). Phenomenal Variability and Introspective Reliability. *Mind and Language* 26 (3), 261–286.
Hohwy, J. (2007). Functional Integration and the Mind. *Synthese* 59 (3), 315–328.
Hurley, S. (1998). Vehicles, Contents Conceptual Structure and Externalism. *Analysis* 58 (1), 1–6.
James, W. (1890). *The Principles of Psychology*. New York: Dover.
Jaśkowski, P. (1996). Simple Reaction Time and Perception of Temporal Order: Dissociations and Hypotheses. *Perceptual and Motor Skills* 82, 707–730.
Johnston, A. and Nishida, S. (2001). Time Perception: Brain Time or Event Time? *Current Biology* 11(11), 427–430.
Kelly, S.D. (2005). The Puzzle of Temporal Experience. In A. Brook and K. Akins (eds.). *Cognition and the Brain: The Philosophy and Neuroscience Movement* (218–248). Cambridge: Cambridge University Press.
Kiverstein, J. (2012) What is Heideggerian Cognitive Science? In J. Kiverstein and M. Wheeler (eds.). *Heidegger and Cognitive Science* (1–61). Basingstoke: Palgrave MacMillan.
Kiverstein, J. (2010) Making Sense of Phenomenal Unity: An Intentionalist Account of Temporal Experience. In P. Basile, J. Kiverstein, and P. Phemister (eds.). *The Metaphysics of Consciousness*. Royal Institute of Philosophy Supplement 67, 155–181.
Le Poidevin, R. (2007). *The Images of Time: An Essay on Temporal Representation*. Oxford: Oxford University Press.
Nijhawan, R. (1994). Motion Extrapolation in Catching. *Nature* 370, 256–257.
Noë, A. (2012). *Varieties of Presence*. Cambridge, MA: Harvard University Press.
Paul, L.A. (2010). Temporal Experience. *Journal of Philosophy* 107(7), 333–359.
Phillips, I. (2010). Perceiving temporal properties. *European Journal of Philosophy* 18(2), 172–202.
Phillips, I. (2011). Perception and Iconic Memory: What Sperling Doesn't Show. *Mind and Language* 26, 381–411.
Phillips, I. (forthcoming) The Temporal Structure of Experience. In V. Arstila and D. Lloyd (eds.). *Subjective Time: The Philosophy, Psychology, and Neuroscience of Temporality*. Cambridge, MA: MIT Press.
Pöppel, E. (1988). *Mindworks: Time and Conscious Experience*. Trans. T. Artin. Boston, MA: Harcourt Brace Jovanovich.
Rao, R, Eagleman, D., and Sejnowski, T. (2001). Optimal Smoothing in Visual Motion Perception. *Neural Computation* 13, 1243–1253.
Reid, T. (1850). *Essays on the Intellectual Powers of Man. Cambridge*, MA: *John Bartlett*.
Siegel, S. (2006). How Can We Discover the Contents of Experience? *Southern Journal of Philosophy* 45, 127–142.
Spence, C. and Parise, C. (2010). Prior-Entry: A Review. *Consciousness and Cognition* 19(1), 364–379. doi:10.1016/j.concog.2009.12.001.
Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.
Vibell, J., Klinge, C., Zampini, M., Spence, C., and Nobre, A.C. (2007). Temporal Order is Coded Temporally in the Brain: Early Event-Related Potential Latency Shifts Underlying Prior Entry in a Cross-Modal Temporal Order Judgment Task. *Journal of Cognitive Neuroscience* 19(1), 109–120.
Wertheimer, M. (1912). Experimentelle Studien über das Sehen von Bewegung. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane* 61, 161–265.
Wittmann, M. (1999). Time Perception and Temporal Processing Levels of the Brain. *Chronobiology International* 16(1), 17–32.

Zahavi, D. (2007). Perception of Duration Presupposes Duration of Perception or Does it? Husserl and Dainton on Time Consciousness. *International Journal of Philosophical Studies* 15.3, 453–471.

Further Reading

Husserl, E. (1991) *On the Phenomenology of the Consciousness of Internal Time* (1893–1917). Trans. J. Brough. Dordrecht: Kluwer.

Koehler, W. (1947). *Gestalt Psychology: An Introduction to New Concepts in Modern Psychology*. New York: Liverlight.