

The Puzzle of Temporal Experience¹

Sean D. Kelly
Princeton University

To appear in *Philosophy and Neuroscience*, eds. Andy Brook and Kathleen Akins (Cambridge: Cambridge University Press, forthcoming).

Last modified: August 16, 2004

I. Introduction

There you are at the opera house. The soprano has just hit her high note – a glass-shattering high C that fills the hall – and she holds it. She holds it. She holds it. She holds it. She holds it. She holds the note for such a long time that after a while a funny thing happens: you no longer seem only to hear it, the note as it is currently sounding, that glass-shattering high C that is loud and high and pure. In addition, you also seem to hear something more. It is difficult to express precisely what this extra feature is. One is tempted to say, however, that the note now sounds like it has been going on for a very long time. Perhaps it even sounds like a note that has been going on for *too* long. In any event, what you hear no longer seems to be limited to the pitch, timbre, loudness, and other strictly *audible* qualities of the note. You seem in addition to experience, even to *hear*, something about its temporal extent.

This is a puzzling experience. For surely, it seems, you never actually hear anything but the note as it is currently sounding. You never actually hear anything, one might have thought, except for the *current audible* qualities of the note. But how long the note has been going on – its temporal extent – does not seem to be an audible feature of the note that is on a par with its pitch, timbre, and loudness. So how is it possible for one to experience the note as *having gone on for a very long time*? How is possible, indeed, for one *now* to hear something as *having been* a certain way? This is a particular example of the kind of problem I am interested in. It is a very specific kind of problem that concerns our capacity to experience time.

When I first began thinking about our experience of time I expected to find a vast quantity of literature devoted to the problem. I expected this literature to span the disciplines of philosophy, psychology, and neuroscience, and I therefore expected the area to be ripe for

¹ I have benefited, in writing this paper, from the feedback of many people. I would like to thank audiences at MIT, Wake Forest, Georgetown and Bryn Mawr, as well as what seemed like the entire Psychology Department at Princeton. All these audiences listened and responded, sometimes vociferously, to early drafts of the paper. I would like especially to thank Andy Brook and Kathleen Akins for the wonderful conference they organized on Philosophy and Neuroscience at Carleton University, at which I presented the original draft. In developing the ideas in the paper, and other related ideas that I hope will appear later, I remember especially interesting discussions with Adrian Bardon, Michael Berry, Dave Chalmers, Cheryl Kelly Chen, Bert Dreyfus, Adam Elga, Brian McLaughlin, Ann Treisman, Frank Tong, and Michael Tye. In addition to several very helpful discussions on this topic, I would like to thank my wife Cheryl Kelly Chen for everything else, not least of which includes all her help and support.

interdisciplinary research. In some sense I was not disappointed. There is, for instance, a vast literature in neuroscience on the neural basis of temporal illusions like the flash-lag illusion;² likewise, there is in psychology a considerable industry devoted to such issues as the temporality of gestalt switching, the temporal limits on attention, and the psychophysics of visual motion;³ and in philosophy, of course, there is a growing movement concerned with questions about the metaphysics of time.⁴

To my surprise, however, I discovered that there is in fact very little written in any of these areas nowadays about the particular problem I am interested in. There is very little written, in other words, about our experience of the passage of time, at least if the issue is construed in a certain way. This is partly because, I believe, the problem I am interested in is rather difficult to articulate. The goal of my paper, therefore, is quite modest: I hope to show you there is a problem here, one that is being overlooked in all of these fields. I will have succeeded if I manage to make you puzzled about something that you hadn't already known was puzzling.

2. Preliminary Statement of the Puzzle of Temporal Experience

For the purposes of this paper, I will approach the puzzle through the examination of a cluster of phenomena that share the same general form. These phenomena include: 1) the perception of objects as persisting through time; 2) the perception of objects as moving across space; 3) the perception of events as occurring in a unified but temporally ordered manner (as, for instance, in a melody); and generally speaking, 4) our experience of the passage of time.

Any careful and sensitive treatment of these phenomena would of course distinguish them in a variety of ways. In this paper, however, I will treat all of these phenomena as manifestations of the same general philosophical problem. At the general level I am interested in I do not believe the differences among these phenomena are as relevant as their similarities. As a result I will be cavalier about my use of these examples, blithely moving among them as if they were identical in all the relevant respects. This is clearly false, but I hope my strategy is not therefore illegitimate. For my main goal is to isolate a philosophical problem the various sides of which are seen more or less clearly through the various phenomena listed above.

The puzzle I am interested in can be stated, in preliminary form, as follows:

The Puzzle of Temporal Experience: 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?

² See, e.g., Nijhawan 2002, Whitney et. al. 2000, Krekelman 2000, Eagleman & Sejnowski 2000, 2001, 2002. See also footnote 35 for further references.

³ See, e.g., Pöppel 1988, Palmer 1999 and references therein.

⁴ See, e.g., Sider 2001, Le Poidevin 1998, Mellor 1998.

This philosophical puzzle of time and experience has largely been ignored in philosophy for the past fifty years. Even so, it is possible to isolate in the history of philosophy two traditional ways of resolving the puzzle.

The first approach centers on what has been called the theory of the Specious Present. According to this theory, we are wrong to think of our experience as providing us with static snapshots of the world. Rather, we are in direct perceptual contact with an ordered, temporally extended window on the world. The strategy of the Specious Present Theorist, therefore, is to reject the puzzle by rejecting the model of experience it presupposes. The principal advocates of the Specious Present Theory, in its various forms, are William James, C. D. Broad, and more recently, the British philosopher Barry Dainton.⁵

The second approach to the philosophical puzzle of temporal experience depends upon what I will call the Retention Theory. According to the Retention Theorist our experience is fundamentally a presentation of the world at a time. The Retention Theorist, in other words, accepts the presupposition that experience presents us with a snapshot of the world. But this is not *all* that experience gives us, according to the Retention Theory. Rather, in experience the snapshot that we get of the world is always supplemented with memories or “retentions” from the past and anticipations or “protentions” of the future. These retentions and protentions are the key to our experience of the passage of time. The principal advocates of the Retention Theory are Immanuel Kant, Edmund Husserl, and possibly Maurice Merleau-Ponty.⁶

Although these are the only going ways of resolving the puzzle, I do not believe that either of them works. The Specious Present Theory, I will argue, simply makes no sense. It is committed to claims about experience that have no sensible interpretation. The Retention Theory, by contrast, though it is correct as far as it goes, ends up raising more questions than it answers. Although one can make sense of the theory, it does not give an explanatory account of the puzzle. This is what makes the philosophical puzzle of temporal experience so puzzling: the only natural ways to resolve it are either senseless or insufficient. My strategy, therefore, is to make you puzzled about temporal experience by showing how difficult it is to deal with it adequately.

Before I get to this stage, however, I want to begin by distinguishing the puzzle I have proposed from two other, more familiar problems that also deal with time and experience. I call these problems, instead of puzzles, since they seem to admit of various more or less tractable solutions. The puzzle of temporal experience, by contrast, seems to require a reconsideration of some of its motivating assumptions.

⁵See James 1890, Broad 1923, 1938, and Dainton 2000.

⁶ See Kant 1781/1787, Husserl 1893-1917, and Merleau-Ponty 1945. It may seem odd to Husserl scholars that I distinguish the Specious Present Theory from the Retention Theory since, at least on a cursory reading of Husserl’s time consciousness lectures, he appears to advocate both. It is crucial, however, to recognize that these two views are opposed to one another. To the extent that Husserl seems to advocate the Specious Present Theory it is because, I believe, he understands this theory differently from the way I will be presenting it, and indeed differently from the way its proponents intend it.

3. Two Problems of Time and Experience: The Time Stamp Problem and the Simultaneity Problem.

3.1 The Time Stamp Problem

One problem that philosophers and psychologists have worried about is what I will call the Time Stamp Problem. The leading question behind the Time Stamp Problem is this:

Time Stamp Problem: How do we come to represent events as occurring at a particular time, and therefore to represent some events as occurring before others?

The Time Stamp Problem is familiar from recent work in philosophy, but it has a long history in psychology as well. Major contributors to the discussion of this problem include Daniel Dennett, Hugh Mellor, and the Gestalt psychologist Wolfgang Köhler.⁷ Although the Time Stamp problem lies at the intersection of issues concerning experience and temporality, it is not the problem I am interested in here. Let me say a word or two about this problem, however, in order to distinguish it from my own.

There are two main approaches to the Time Stamp Problem, each presented with typical clarity by Dennett in his book *Consciousness Explained*.⁸ According to the first approach, the time at which an event is represented to occur is determined by the time at which the relevant brain process occurs. The relevant brain process is the one that, so to speak, correlates with the experience. So, for example, if the brain process correlating with my experience of a green flash of light occurs before the brain process correlating with my experience of a red flash of light, then I will experience first a green flash and then a red. On such an account one can say that time is its own representation. The time at which an event is represented to occur is itself represented only by the already existing temporal order of the brain processes involved. There is an isomorphism, in other words, between the temporal order of the brain processes and the represented temporal order of the events in the world. Köhler was perhaps the main proponent of this isomorphism view, though it was typical of gestalt psychologists generally. As Köhler writes in his classic book *Gestalt Psychology*, “Experienced order in time is always structurally identical with a functional order in the sequence of correlated brain processes.”⁹

The second approach to the Time Stamp Problem rejects Köhler’s simple isomorphism view. According to this second strategy, the time at which the correlated brain process occurs is completely irrelevant, since representations can have time stamps – independent indicators of the time at which an event is represented to occur. On such a view, time is represented by something other than itself. This idea is very familiar to us from the way we date our letters. If we are interested in the time at which the events in a letter are represented to occur, we should look at

⁷ Dennett 1991, Mellor 1998, Köhler 1947.

⁸ Dennett 1991.

⁹ Köhler 1947, p. 62.

the date on the letter rather than pay attention to the time at which the letter arrives. (Think of the time of the arrival of the letter here as analogous to the time of the occurrence of the brain process.) Indeed, just because letter B arrives in my mailbox on Monday and letter A arrives on Tuesday, there is no guarantee at all that the events in letter B are also *represented to occur* before the events in letter A. If letter A has come from farther away, for example, and therefore took a longer time to arrive, or if it simply got lost in the mail, then it could very well happen that letter A tells me about events that are represented as occurring before the events in letter B, even though letter A arrives later. As Dennett writes, in discussing brain processes instead of letters:

What matters is that the brain can proceed to control events “under the assumption that A happened before B” whether or not the information that A has happened enters the relevant system of the brain and gets recognized as such before or after the information that B has happened.¹⁰

Dennett’s own Multiple Drafts theory of consciousness depends upon the plausibility of this second approach to the Time Stamp Problem. I certainly believe the problem is an interesting one, and that much work remains to be done in this area.¹¹ But I do not intend to say anything more about it here. My main goal at this juncture is simply to emphasize that the Time Stamp Problem is not the problem with which I am concerned.

3.2 The Simultaneity Problem

A second problem in the general area of time and experience is what I call the Simultaneity Problem. This problem is deceptively simple to state:

Simultaneity Problem: Which events do we experience as simultaneous?

The Simultaneity Problem is not much discussed in the literature today, but it was important for Russell during a certain period of his development,¹² and in any case it is quite easy to generate.

¹⁰ Dennett 1991, p. 149.

¹¹ One principal problem is to determine what a neural time stamp could look like and how we should go about finding one. It is one thing to look for correlations between the firing rate of a neuron and the presence within a receptive field of, for instance, a horizontal line. But if the stimulus is not a horizontal line but instead the time at which the horizontal line is represented to appear, then it seems the methodology for varying the stimulus will be trickier. The problem is even more complicated when you realize that in the simple cases, like neurons in V1, Köhler’s isomorphism story is likely to be right. More complicated cases of temporal illusion are clearly the place to focus here, and the burgeoning literature on the flash-lag illusion is perhaps a case in point. (See references in footnote 2.) But these models do not go unequivocally in Dennett’s direction, since they tend to rely upon a combination of low-level temporal isomorphism and higher-level non-isomorphic time stamping.

¹² See chapter 6 of Russell’s unpublished 1913 manuscript *Theory of Knowledge*. (Russell 1913.)

Let us assume a simple, non-relativistic notion of simultaneity for the purposes of discussion. Obviously, then, there are events in the world that are not as a matter of fact simultaneous, which we nevertheless experience as simultaneous. (This follows trivially from the fact that the temporal acuity of experience is not infinitely fine-grained.) The Simultaneity Problem, then, is to determine which events these are.

One might wonder why experienced simultaneity is a problem. After all, one might think, can't we simply say that two events are experienced to be simultaneous if and only if they seem to the subject to be simultaneous when they are experienced? That would certainly be the simplest approach. But a problem quickly arises if we proceed this way. For if we define experienced simultaneity in this obvious manner then the relation turns out to be non-transitive. Since it is meant to be an equivalence relation, this would be a devastating blow.¹³ Let us see in a bit more detail how this works.

Consider any two events A_1 and A_2 that seem to me to occur simultaneously. For example, imagine that A_1 is a flash of light and A_2 is a tone. Then take some third event A_3 , which actually occurs later than both A_1 and A_2 but is experienced as simultaneous with at least the second. Perhaps A_3 in this case is another flash. By this method we can generate a series of events A_1, \dots, A_n such that each is experienced as simultaneous with the next when considered pair-wise, but such that A_1 and A_n , when considered pair-wise, are experienced as non-simultaneous. In this way we can show that the relation of experienced simultaneity is non-transitive.

The Simultaneity Problem is formally equivalent to Nelson Goodman's problem about the non-transitivity of appearance properties.¹⁴ Recall that Goodman considers a series of color chips R_1, \dots, R_n such that R_j appears identical to R_{j+1} for all j from 1 to $n-1$, but it is not the case that R_1 appears identical to R_n . Just as with the Simultaneity Problem, if we define the "appears identical" relation in terms of the way the colors look to the subject in pair-wise comparisons, then we can show that this relation is non-transitive.

There is an obvious way to avoid this deficiency, and it is perhaps not surprising that Russell seems to have discovered the same solution to the problem as Goodman. The trick is to deny that two events are experienced to be simultaneous if and only if they seem to the subject to be simultaneous when they are experienced. Rather, two events A and B are experienced to be simultaneous if and only if they meet the following two conditions:

¹³ One possibility, of course, is simply to deny that experienced simultaneity is an equivalence relation. Few have been tempted by this approach, however, perhaps because it has the effect of divorcing the world so radically from our perceptual representation of it. This divide is more uncomfortable for some perceptual theorists than others. Disjunctivists about perception, for example, are typically committed to the idea that in perception we are presented with facts about the way the world is. (McDowell is a good example of a view like this.) On such a view it would be extremely awkward to let simultaneity in the world be an equivalence relation but to deny this for our perceptual representation of it.

¹⁴ See Goodman 1951.

1. A and B seem to be simultaneous.
2. There is no third event C such that one of A or B seems to be simultaneous with C and the other does not.

If experienced simultaneity is defined in this more restrictive way, then the problem of non-transitivity does not arise. That is because the second condition effectively excludes, by fiat as it were, all of the cases that could give rise to a problem. But the solution has a surprising consequence as well. It turns out, on this approach, that just because two events seem to me at the time to be simultaneous, I cannot thereby conclude that I am experiencing them to be simultaneous. I will only know that I have experienced them to be simultaneous once I can show in addition that there is no possible third event such that it seems to be simultaneous with one but not with the other. Experienced simultaneity, in other words, is more fine-grained than momentary reflection would indicate. As a phenomenologist I find this result congenial. That is because phenomenologists believe in general that experiences are richer and more complicated than one can know simply by having them. But again, my goal here is not to discuss the Simultaneity Problem in any great detail; it is only to emphasize that this is not the puzzle with which I am concerned.

4. Development of the Puzzle of Temporal Experience.

Both the Time Stamp Problem and the Simultaneity Problem are focused on *when* we experience events to occur. The Time Stamp Problem is concerned with whether we experience A to be before or after B. The Simultaneity Problem is concerned with how to define the equivalence class of events that are experienced to be simultaneous. But the question of time and experience in which I am interested is not the question *When* do we experience events to occur but rather the question *How* do we come to experience events as occurring *through* time at all? This is a question about experiencing the *passage* of time, not just a question about *at what time* we experience an event to occur.

The distinction between experiencing events to occur *at a time* and experiencing them to occur *over or through time* has been discussed since at least the writings of St. Augustine.¹⁵ By the period of the Renaissance this distinction had worked its way into the culture more generally. Consider, for instance, Shakespeare's Sonnet number 104, in which he writes of a gracefully aging friend:

Ah! yet doth beauty, like a dial-hand,
Steal from his figure and no pace perceived,
So your sweet hue, which methinks still doth stand,
Hath motion and mine eye may be deceived.

In this stanza Shakespeare compares the graceful decline of his friend's beauty with the slow movement of the hand on a clock: although one never sees any change occur ("no pace

¹⁵ See Augustine, *Confessions* Book 11.

perceived”) it is nevertheless clear after some time that there has been a change (“Hath motion and mine eye may be deceived”).¹⁶

In general, it is a common observation that some movements or changes happen too slowly for us to experience them *as movements or changes occurring through time*. For example, Locke considers in the *Essay* cases in which

The Body, though it really moves, ... [nevertheless] seems to stand still, as is evident in the Hands of Clocks, and Shadows of Sun-dials, and other constant, but slow Motions, where though after certain Intervals, we perceive by the change of distance, that it hath moved, yet the Motion it self we perceive not.¹⁷

By contrast, successive events sometimes occur too swiftly for us to perceive them as part of a movement or change through time. Locke makes this point in a colorful way:

Let a Cannon-Bullet pass through a Room, and in its way take with it any Limb, or fleshy Parts of a Man; 'tis as clear as any Demonstration can be, that it must strike successively the two sides of the Room: 'Tis also evident, that it must touch one part of the Flesh first, and another after; and so in Succession: And yet I believe, no Body, who ever felt the pain of such a shot, or heard the blow against the two distant Walls, could perceive any Succession, either in the pain, or sound of so swift a stroke.¹⁸

But there is a wholly distinct phenomenon that seems to lie somewhere between perceived precedence and perceived simultaneity, and that is the phenomenon of perceived movement or change. Perceived movement does not seem to be the kind of thing in principle that can be explained in terms of perceived precedence or perceived simultaneity. It seems, in other words, to be a basic kind of perception. As C. D. Broad writes:

It is a notorious fact that we do not merely notice that something *has* moved or otherwise changed; we also often see something *moving* or *changing*. This happens if we look at the second-hand of a watch or look at a flickering flame. These are experiences of a quite unique kind; we could no more describe what we sense in them to a man who had never had such experiences than we could describe a red color to a man born blind. It is also clear that to see a second-hand *moving* is a quite different thing from “seeing” that and hour-hand *has* moved.¹⁹

If perceived movement is a basic kind of perception, then a certain question naturally arises: What must perceptual experience be like if it is to allow for the possibility of the perception of motion or change? Since all motion takes place not only across space but also

¹⁶ Shakespeare does not specify which dial-hand he is talking about, but naturally the metaphor would not work if it were the second hand on a clock. Probably there were no second hands in Shakespeare’s day.

¹⁷ Locke 1690, Book 2, chapter 14, §11.

¹⁸ Locke 1690, Book 2, chapter 14, §10.

¹⁹ Broad 1923, p. 351.

through time, this question is closely related to the question with which we are principally concerned: What account of perceptual experience explains how it is possible for us to experience the passage of time? Two theories are prominent in historical accounts, but neither seems to do the job.

5. The Specious Present Theory

The Specious Present Theory found its first major proponent in William James, who developed the theory at length in *The Principles of Psychology*. The central idea behind the theory is that instead of giving us a snapshot of the world at a time, perception presents us with a temporally extended window of events. In particular, the Specious Present Theory proposes that a subject is in direct perceptual contact with an ordered, temporally extended, unified expanse. As James writes, in a famous passage from *Principles*:

The practically cognized present is no knife-edge, but a saddle-back, with a certain breadth of its own on which we sit perched, and from which we look in two directions into time. The unit of composition of our perception of time is a *duration*, with a bow and a stern, as it were – a rearward- and a forward-looking end ... We do not first feel one end and then feel the other after it, and from the perception of the succession infer an interval of time between, but we seem to feel the interval of time as a whole, with its two ends embedded in it.²⁰

According to this theory of the nature of perceptual experience, we are in direct perceptual contact not only with what is *now* occurring but also with what *has recently occurred* and indeed with what is *about to occur* as well.²¹

Although I believe there is something important about the perceptual phenomenon that this theory is intended to explain – and indeed I believe that no theory of perceptual experience would be complete without an explanation of the temporal effects in question – it nevertheless seems to me that there are at least three devastating objections to the Specious Present Theory itself. These objections coalesce around the following three questions:

1. How can I be directly aware of something that is no longer taking place?
2. How can I be directly aware of a duration?
3. How can I be directly aware of the future?

²⁰ James 1890, pp. 609-10.

²¹ James's version of the Specious Present Theory is one of the few I know of that advocates both a past and a future aspect to the specious present. Most accounts limit the window of the specious present to a short period of time that stretches back from now. James's proposal follows (as always) from his acute observations of the perceptual phenomena. For it is certainly true that our current experience is directed toward both the recent past and the near future; theories of the specious present that deny this fact are certainly less attuned to the phenomena they hope to explain. As we will see in a moment, though, the claim that we are in *direct perceptual contact* with the future is itself specious at best.

I will take these questions in order.

5.1 Problem 1: The experience of the recent past

The Specious Present Theory proposes that I am in direct perceptual contact with events that occurred in the recent past. This is at best an odd suggestion. After all, the events in the recent past are no longer occurring, and one might naturally wonder how I can be directly aware of something that is no longer taking place. I can think of two obvious ways to interpret the theory so that it gives an answer to this question. Unfortunately, neither of them does the work that the Specious Present Theory needs. I conclude that the theory has no obvious interpretation on which this central claim makes sense.

So, then, how can I be directly aware of something that is no longer taking place? One suggestion is to emphasize the time lag that always exists between the occurrence of an event and the subject's experience of the event's occurring. In extreme cases this time lag can be millions of years. When you look into the night sky and see the explosion of a star, it may be that the event you are now witnessing actually took place many millions of years ago (if the supernova occurred millions of light-years away). Of course, not every time lag between event and experience is this extreme. Sometimes the experience of the event occurs only a very short time after the event itself. Even so, the experience always occurs at least some time after the event, since it always takes time for the waves carrying information about the event to reach the sensory organs, and some further time for the brain to process this information in such a way that it gives rise to an experience. The result, therefore, is that I am *always* aware of events that are no longer taking place, since experience lags behind the event experienced. Perhaps, then, this time lag between an event and the experience of it is a way of making sense of the claim that I can experience what is past.

The time lag suggestion, however, is irrelevant to the Specious Present Theory for at least two reasons. First, this suggestion implies, against the Specious Present Theory itself, that I am never aware of the present, never mind aware of the future. The Specious Present Theory proposes, you will recall, that I am aware of a temporal window of events that reaches into the past but includes at least the present and possibly also the future. But the time lag argument emphasizes that the events I am aware of are always events that occurred in the past. Second, the Specious Present Theory says not just that I am perceptually aware *of* the past, but that I am aware of it *as* the past. Without this the theory cannot make any sense at all of my experience of the passage of time. The time lag phenomenon, however, shows only that the events I am aware of occurred in the past, not that I am aware of them as occurring in the past. Indeed, these two phenomena usually come apart: it is always a surprise for children to learn, for example, that looking into the night sky is the same as looking back in time.

A second suggestion makes more headway, but in the end is insufficient as well. This suggestion derives from what Russell calls, in *The Analysis of Mind*, "akoluthic sensations". The akolouthoi, in ancient Greece, were the so-called camp followers who trailed behind the soldiers on their way to battle, providing services for them and collecting valuables they left behind. The

akoluthic sensations, in Russell's analysis, are the sensations one has of events as just having occurred. Take, for example, the acoustic case. When I clap my hands together, the primary sensation is as of a sharp sound that lasts for an instant. But after this primary sensation runs its course, it is natural to say that one continues to experience the clap of the hands, but now one experiences it as recently past. This is an akoluthic sensation of the clap. Russell's classically empiricist analysis of this phenomenon proposes that hearing a sound *as* past in this way is hearing the same sound but with less "force and vivacity". He writes:

Succession can occur within the specious present, of which we can distinguish some parts as earlier and others as later. It is to be supposed that the earliest parts are those that have faded most from their original force, while the latest parts are those that retain their full sensational character. ... Sensations while they are fading are called "akoluthic" sensations.²²

Perhaps Russell's notion of akoluthic sensations gives us a clue as to how we should develop our theory of perception so that it allows us to be directly aware of the recent past.

I am not sanguine about this approach. The problem is that Russell's account fails to make sense of one of the main problem cases we are hoping to solve – that is, the case of perceived motion. Consider the experience of watching Derek Jeter's throw travel through the air from shortstop to first. On Russell's account the earlier phases of the moving ball are now seen, but are seen less forcefully than the current phase. If that were right, then seeing a moving object would be a matter of seeing it vividly at its current position, but with a continually fading trail. Needless to say, this does not capture the experience of watching Jeter's throw, unless your television is significantly worse than mine.

5.2 Problem 2: The experience of duration

A second challenge for the Specious Present Theory is to explain how I can be directly aware of a duration. The default position seems to be that I am directly aware of what is presented to me now. As with awareness of the past, however, awareness of a duration of time requires that some part of what I am directly aware of is not occurring now. So how can I be perceptually aware of something that is not now occurring?

Once again, the most natural approach to this problem turns out to be flawed in a crucial way. The most natural approach emphasizes that we always experience a duration, but we experience it as a moment in time. Recall Locke's example of the Cannon-bullet. The point of that example was to show that our perceptual experience is relatively coarse-grained – it groups together events that happen successively into a single undifferentiated experience. Even though the bullet hits first one wall and then the other, for example, we are apt to experience only a single sound. In this sense, therefore, we always experience events that happen across a span or duration of time.

²² Russell 1921, p. 145.

It is certainly true that our experience is relatively coarse-grained in this way. It is, however, irrelevant to the challenge raised for the Specious Present Theory. That is because there is a difference between the claim made by the Specious Present Theory and the fact explained by the suggestion above. The suggestion proposes rightly that we experience temporally distinct events as simultaneous. But the Specious Present Theory claims more than this. It claims that we experience temporal extension itself. This strange claim is not in any way clarified by the correct but trivial observation that experience groups together successive events. Perhaps there is some other way to explain the claim that we experience temporal extension directly. If so, it will require a broader defense of the theory of perception that lies behind the story of the specious present.

5.3 Problem 3: The experience of the near future.

A third challenge for the Specious Present Theory is to explain how I can be directly aware of the near future. I have already mentioned in passing that not every version of the theory insists that the specious present has a future dimension. But perhaps the most famous version of it, manifest in the passage quoted from James above, has experience looking in two directions into time. It is odd enough to think that we could be directly aware of something that was occurring but no longer is. But it is odder still to think that we can be directly aware of something that has not yet occurred. This view of perception seems to smack of extra-sensory powers. What could the Specious Present Theorist mean?

As with the previous two problems for the theory, there is a natural way to answer this third question. But once again it turns out not to do the necessary work. The natural suggestion is to give examples in which it really does seem to make sense to say that I can be aware that something is *about to occur*. Here's one. Sometimes if you are standing in the batter's box watching the pitch come in it becomes immediately clear to you that you are about to be hit by the pitch. It is not ESP. It is just that you can see the pitch heading for your shoulder and you are immediately inclined to duck out of the way. This kind of thing seems to happen all of the time. You can see that the light is about to turn red, that the pedestrian is about to walk out into the street, that the professor is about to trip over the extension cord at the front of the lecture hall, and so on. Perhaps it sounds strange in the abstract to say that we are directly aware of these events before they occur, but after all isn't that the phenomenon?

I do not intend to deny the phenomenon. As with the experience of the recent past and the experience of duration, there is no doubt that there is a sense in which our experience has a futural dimension to it. The question is whether the best explanation of this phenomenon is to say that we have direct perceptual awareness of the future. Without a more robust and independently motivated theory of perception to back up a claim like this it seems a very large pill to swallow. After all, nobody can deny that there are some intentional states by means of which we can be directed toward the future. Anticipation, expectation, hope, desire, and so forth are all quite good ways of doing this. But that we have a direct perceptual awareness of the future is an unsubstantiated claim at best.

5.4 A final problem.

Let us suppose for the sake of argument that the Specious Present Theorist can get around the previous three objections. Let us suppose, in other words, that we can make sense of the central claims of the theory. There is still an important final problem for the theory – namely, that it fails to explain the central problem case it was invented to explain. That is the problem, as Shakespeare puts it, of pace perceived.

To see that the perception of motion remains unexplained, we need only to notice that the specious present, by nearly all accounts, lasts only a relatively limited time. Recent estimates generally agree that it is in the area of three seconds or so.²³ But we often experience things to be moving for periods that are longer than this. If you watch an airplane taking off from the runway you can follow its continuous motion for several minutes before it disappears. Even on the Specious Present Theory, therefore, we must *keep track* of the earlier phases of long movements in some way other than by perceiving them directly. That we have some relation to the past and the future other than direct perception of it, however, is the main point of the Retention Theory. I turn to this alternative theory now.

6. The Retention Theory.

The Retention Theory has its origins in the empiricism of Locke and Hume. Recall that one central tenet of their empiricism is that in the most basic cases our ideas of properties and objects in the world originate in our perceptions or impressions of them. We have the idea of square-ness for example, on such a view, because we have had a perception or impression of a square. The question naturally arises, then, what the origin is of our ideas of succession and duration, and generally of our idea of time. Locke denies outright that we have any perception of duration by means of which we can get our idea of it.²⁴ That is to say, Locke denies the central tenet of the Specious Present Theory. Hume expands upon this idea in the *Treatise*:

The idea of time is not deriv'd from a particular impression mix'd up with others, and plainly distinguishable from them; but arises altogether from the manner, in which impressions appear to the mind, without making one of the number. Five notes play'd on a flute give us the impression and idea of time; tho' time be not a sixth impression, which presents itself to the hearing or any other of the senses.²⁵

If we do not get our ideas of succession and duration from perceptions of time, according to the Empiricists, then what precisely is the origin of these ideas? Locke and Hume do not give

²³ These estimates are ascertained empirically by measuring the amount of time between, for instance, gestalt shifts in the perception of a Necker cube. See Pöppel 1988. William James is again a renegade here, since his estimates for the length of the specious present have it lasting somewhere between 3/4 of a second and twelve minutes. The principle of charity requires that we leave this aspect of his position behind.

²⁴ “we have no perception of duration...” Locke 1690, Book 2, chapter 14, section 4.

²⁵ Hume 1739, Book 1, chapter 2, section 3, pp. 34-5.

a particularly detailed story about this. They do suggest, however, without much further explanation, that the succession and duration of our ideas themselves is central in the development of our ideas of succession and duration. As Hume says, “From the succession of ideas and impressions we form the idea of time”.²⁶

If the suggestion here is that merely by *having* a succession of ideas we can thereby get an idea of one thing’s following another, then the suggestion is clearly wrong. For imagine a creature who has a succession of experiences, but at each moment forgets all the previous ones. Such a creature, though it has a succession of experiences, does not seem thereby to have the resources for forming the idea that each experience follows the previous one. For without any memory of the previous experience, such an idea is impossible to form. That is why William James (following Kant) says, rightly, that “A succession of feelings, in and of itself, is not a feeling of succession”.²⁷

If a mere succession of experiences is not sufficient for the idea or experience of succession, what must be added? Kant famously develops a story about this in the A-Deduction of the first *Critique*. It is the story of the so-called “threefold synthesis”. The key to Kant’s account is that at every moment I must not only have an experience of the thing before me now, but also I must “reproduce in imagination” the things I experienced in the recent past. Notice that I don’t *perceive* these past events, according to Kant, but rather *reproduce* them in *imagination*. As he writes:

When I seek to draw a line in thought, or to think of the time from one noon to another, or even to represent to myself some particular number, obviously the various manifold representations that are involved must be apprehended by me in thought one after the other. But if I were always to drop out of thought the preceding representations (the first parts of the line, the antecedent parts of the time period, or the units in the order represented) and did not reproduce them while advancing to those that follow, a complete representation would never be obtained.²⁸

But what kind of intentional attitude, we might naturally ask, is the attitude of reproduction in imagination? There are lots of views among Kant interpreters, and I certainly do not intend to wade into this thorny literature here. But one prominent view, which is suggested at least implicitly by Robert Paul Wolff, is that *reproducing* an event in imagination is like *remembering* it. As Wolff explains:

What I must do ... as I proceed from one moment to the next, is to reproduce the representation which has just been apprehended, carrying it along in memory while I apprehend the next. In looking at a forest, I must say to myself, “There is a birch; and there is an elm, plus the birch which I remember, etc.”²⁹

²⁶ Hume 1739, Book 1, chapter 2, section 3, p. 35.

²⁷ James 1890, p. 628.

²⁸ Kant 1781, p. A102.

²⁹ Wolff 1963, p. 128.

Whatever the virtues of this as an interpretive claim, I do not believe that the subject's intentional relation to recently past events is anything like a memory of them. In this I follow Husserl, whose notion of retention is meant to do the work that Wolff believes something like memory can do. Before I present Husserl's approach, therefore, let me say why memory is not sufficient.

Memory comes in a variety of different forms, so it is important to specify which kind we have in mind. I will consider two. First, there is the kind of memory in which one is reminded of something in a flash. It occasionally happens to me, for example, that as I walk out the front door of my apartment, and hear it closing locked behind me, I am reminded instantly that my house keys are on the kitchen table inside. This kind of memory brings something to mind in a flash and apparently without any effort on the subject's part at all. It has the feeling of a kind of instantaneous revelation, in which all of a sudden something completely new is brought to mind.

This kind of instantaneous remembering cannot be the relation we have to the recent past. This becomes clear when we think again of Hume's example of hearing a melody. Hearing the fifth note of the melody does not involve being instantaneously reminded of the earlier ones at all. It is not as if when I hear the fifth note the earlier notes jump immediately to mind like the image of the keys when the door clicks shut. As Husserl says, "A present tone can indeed "remind" one of a past tone, exemplify it, pictorialize it; but ... the intuition of the past cannot itself be a pictorialization."³⁰

A second kind of memory falls naturally under the heading of *entertaining* a memory. One entertains a memory when one goes over in one's head, often from a first person point of view, events in which one was previously involved. I might, for example, entertain the happy memory of my wedding ceremony. When I do so, I go back over the events of the ceremony in my mind, experiencing them again as if they were occurring now. Some people are extremely good at this kind of remembering, and can even produce autonomic responses such as increased heart rate by entertaining memories in this way.

The problem with this kind of remembering is that in putting myself again in the situation I experience the events as occurring *now*. If we model the experience of the melody this way, then the earlier notes in the melody are experienced as present together with the current one. It would be as if, while hearing the fifth note, I simultaneously entertain the memory of the notes before. But this would give me the experience of a chord instead of the experience of an extended event. It would be a very odd type of chord, of course, since one of the notes would now be *heard* while the others would simultaneously be *entertained in memory*, but it would be a type of chord nevertheless.

Husserl concludes from considerations like these that our relation to recently past events cannot be any kind of memory. He coins the term "retention" to characterize our actual relation to the recent past, and claims that retention is a unique kind of intentional act that is unlike any kind of reproduction or memory. Retention is, by definition, a way of being directed towards objects and events *as just-having-been*. About perceived motion, for example, Husserl says:

³⁰ Husserl 1893-1917, §12.

During the time that a motion is being perceived, a grasping-as-now takes place moment by moment; and in this grasping, the actually present phase of the motion itself becomes constituted. But this now-apprehension is, as it were, the head attached to the comet's tail of retentions relating to the earlier now-points of the motion.³¹

The problem with Husserl's account, as I see it, is that it seems only to name the phenomenon instead of to explain it. We have no interesting account of what it is now to experience something *as just-having-been*, except to say that it is the phenomenon involved in the experience of the passage of time. But this is the phenomenon we are trying to *explain*. It does no good just to give a name to its various parts. What, after all, is it now to experience something *as just-having-been*? I know what it is now to think of George W. Bush *as the President of the United States*. Our whole theory of intentionality is built around cases of linguistic predication such as this. But what kind of perceptual phenomena can we point to that will help us unravel the *temporal aspect* of perceptual experience? Husserl does not even attempt to say.

I conclude from this that Husserl's theory is some kind of advance, since at least it does not have the disadvantages of Kant's view or of the specious present. But Husserl's account seems to raise more questions than it answers. What we would like is a standard set of examples that give us the feel for what it is to experience something now *as just-having-been*. It would be even nicer to formulate these examples in such a way that we could pose empirical hypotheses about their features. In short, our project should be to give examples like the kind that Russell uses to expound his theory of akoluthic sensations, but to give examples that are phenomenologically apt instead of obtuse. I believe there are good examples of this sort, and I would like to think more about how to discuss them in the future. But for the moment I will be content if I have managed to convey how puzzling the phenomenon of temporal experience really is.

7. Appendix: Temporal Experience and the Science of the Brain

The goal of this paper was to make you puzzled. I have discovered, however, that it is very difficult to make psychologists and neuroscientists (as well as some philosophers) puzzled about this topic in the appropriate way. Usually these unflappable types have one or the other of two different responses. The first is to assert that we are already busy filling in the neuroscientific details that, when complete, will constitute a satisfying answer to the problem I have posed. There is no particularly gnawing puzzle of temporal experience, on this account, just an empirical problem that we are busy solving. The second is to admit that there is a puzzle, but to identify it as a version of the so-called hard problem of consciousness. (The hard problem is really a puzzle in my terminology, since we don't have the conceptual resources yet to understand what the form of a satisfying answer would be.) On this account there is a genuine puzzle about how the brain could give rise to anything like experience at all, but the question how it could give rise to temporal experience is just a special version of that familiar puzzle.

³¹ Husserl 1893-1917, §11.

I believe that neither of these responses is to the point. The puzzle of temporal experience will not be resolved by empirical research of the type now being done. Even so, the issue is not hopeless: temporal experience is a much more tractable puzzle than that posed by the hard problem of consciousness. Indeed, the really nice thing about the puzzle of temporal experience is that it seems to be just the kind of puzzle that would benefit most directly from the interaction of conceptual, philosophical work and empirical psychological and neuroscientific progress. The goal of this appendix is to say how this interaction might work. I will pursue this goal by showing precisely why each of the two usual responses to the puzzle misses the point.

7.1 Temporal experience really is a puzzle

I begin by explaining why current neuroscientific research, though interesting and important, is not the right kind of thing to resolve the puzzle I have identified. I don't know about all the empirical research there is, of course, and it may be that some neuroscientist somewhere really is making progress. But I doubt it. All of the research I am familiar with is either directed at some version of the Time Stamp and/or Simultaneity problems, or else depends implicitly on either the Specious Present or the Retention theory. For this reason the current empirical research seems either to miss the puzzling issue of temporal experience, or to fail to appreciate its puzzling nature. Let me give one example in a bit of detail.

Some of the most interesting neuroscientific work on time and experience focuses on temporal illusions like the flash-lag illusion.³² In the flash-lag illusion, a flash occurs at a certain location precisely when a moving object reaches that location. The illusion is that the moving object and the flash appear to be offset. On one version of the illusion, for example, the moving object is a bar of constant velocity traveling in a straight line from left to right. When the bar reaches a certain point a second bar flashes directly above it; the first bar continues moving in the same direction. After being presented with this stimulus, the subject is asked to say where the flash seemed to be relative to the moving bar at the instant the flash occurred. In the case just described subjects systematically report that the flash seemed to be behind (to the left of) the moving bar. The flash, in other words, seemed to lag behind the bar; hence the name of the illusion. Interestingly, however, if the bar stops moving at the moment the flash occurs, then subjects report the flash and the bar to be in the same location; and if the bar reverses direction when the flash occurs then subjects report the bar to lag behind the flash instead of the other way around.³³

The principal neuroscientific problem here is to provide a computational model that accounts for the psycho-physical data. Since the judgment about the relative locations of the flash and the bar is affected by what happens after the flash occurs, one model proposes that the brain is taking these later events into account in determining what the percept represents to be

³² See footnote 2 for references.

³³ See Eagleman and Sejnowski (2000) for details. These authors used a moving ring and a flash that occurs inside the ring, but illusion is the same.

happening at the instant of the flash. This is the “Postdiction” account preferred by Eagleman et. al.³⁴ There are also prediction and online accounts of some of these phenomena.³⁵

The question what model the brain is using in order to give rise to this range of psycho-physical results is an enormously interesting and important question. No doubt the final answer will be established not just by the computational neuroscientists but also by their brethren in the cell recording industry. What I wish to emphasize here, however, is that this problem, interesting though it is, has no bearing on the puzzle of temporal experience that I have described. The reason is that the flash-lag illusion is a straightforward hybrid of (admittedly empirical versions of) the Time Stamp and Simultaneity Problems. Insofar as it combines elements of these two problems it is more complicated than either on its own; but it is not for all that a problem of a radically different kind.

To see this recall that the Time Stamp Problem is the problem of determining how we come to represent events as occurring at a particular time, and therefore to represent some events as occurring before others. One issue in the neuroscientific accounts of the flash-lag illusion is just that: what model is the brain using to represent the time at which the flash occurred and to represent the position of the moving bar at that time? The models currently under discussion do not so far take a stand on whether the brain uses time as its own representation in implementing these models, but they have turned this into an empirical question that they can begin to address.

In addition to the question how we represent events occurring at a time there is the question which events we experience as happening together. This is the Simultaneity Problem, and it is the second problem under discussion in the flash-lag illusion. The psycho-physical data that define the illusion just are data about which events the subject experiences to be simultaneous. Russell’s discussion of the Simultaneity Problem, of course, should make us wonder whether we really can take the subject’s report about perceived simultaneity as a good indicator of what the experience actually represents. To the extent that the psychologists assume the subject’s report to be a precise indicator of his percept they have failed to grapple with the philosophical issues behind the Simultaneity Problem. But the question which events we experience as simultaneous is nevertheless a central question for these experimenters.

Insofar as the Puzzle of Temporal Experience is different from either the Time Stamp or the Simultaneity Problem, as I have argued it is, it should be clear that the neuroscientific accounts of the flash-lag illusion fail to impinge on the puzzle I have tried to describe. They do not have any bearing at all, in other words, on the question how we are to conceptualize perceptual experience from the point of view of the subject so as to allow for the possibility of experiences as of continuous, dynamic, temporally structured, unified events. The flash-lag illusion, of course, is not the only temporal phenomenon now being studied. There are many other neuroscientific projects that cover the general territory of time and experience. These include, for example, discussions of the echoic buffer, of apparent motion and the motion

³⁴ *ibid.*

³⁵ For the prediction account see Nijhawan 1994. For the online account see Baldo and Klein 1995; Purushothamen *et al.* 1998; Whitney and Murakami 1998; expanded upon in Whitney, Murakami, Cavanaugh 2000.

detection neurons in area MT, of the neural basis of circadian rhythms, of the temporal structure of gestalt switching phenomena, and so on. Many of these are fascinating problems, and in the case of some of them important neuroscientific research is being done. But in no case that I know of is the work of the sort to address the Puzzle of Temporal Experience that I have characterized.

There is a simple reason why no empirical work is being done on this puzzle. The reason is that we're not yet clear enough on what kind of work would make progress. The puzzle, in other words, has not yet been turned into a problem. We would be clear enough about how to approach the issue if either the Specious Present Theory or the Retention Theory provided the right kind of background model for temporal experience. Some empirical work presupposes one or the other of these theories, and thus gives the appearance of making progress.³⁶ But if my criticisms of these theories have been on target, then what we need is a new background model on the basis of which we can begin to pursue empirical research. It is not my goal to provide such a model here – that will have to wait for another day. But what I can do is say something about why I think the prospects for such a model are better in the case of temporal experience than they are in the more general case of consciousness. I turn to this final topic now.

7.2 The puzzle of temporal experience is not a version of the hard problem (puzzle) of consciousness

The hard problem of consciousness, as David Chalmers calls it, is the problem of determining how the brain – a mass of physical stuff – can be the kind of thing that causes conscious experience. The hard problem is different from various easy problems of consciousness. These concern the question which part of the brain causes this or that kind of experience. The easy problems are the kind that are well-studied using, for instance, fMRI. There are many details and niceties to this kind of work, and it can be very challenging empirically, but the basic idea is very easy to understand. Induce a certain kind of experience in the subject and see which part of the brain lights up when he is having it. Having discovered the relevant kind of correlation, we have some evidence about which part of the brain causes the kind of experience in question.

The various easy problems of consciousness are vulnerable to just the kind of empirical attack that science is very good at making, and we do seem to be making progress on some of them. The hard problem, however, the question precisely how activity in the relevant part of the brain causes anything like experience, is left completely untouched by this technique. This is unfortunate, of course, but it is no great embarrassment for science. Scientists generally have accepted that the hard problem is, for the time being anyhow, outside their domain.

Sometimes when psychologists and neuroscientists, and even some philosophers, have heard me present the ideas in this paper they have accused me of trying to saddle them with a

³⁶ Work on the temporal structure of gestalt switching phenomena, for example, is sometimes taken to be telling us about the temporal extent of the specious present. See Pöppel 1988.

particular version of the hard problem. Their reasoning is something like this. You say that the experience of motion is one of the phenomena you're trying to explain. This is, after all, one of the central examples under discussion, from Locke through Broad and up to your case of Derek Jeter. (Often they pause at this point either to disparage or to applaud the Yankees more generally.) But the problem of motion perception is a good old empirical problem that we are making progress on. After all, we have known since Zeki's work in the early 1970's that a good number of the neurons in cortical area MT are direction-selective: they respond selectively to the movement of objects or fields of dots when they are traveling in one direction but are silent when they are traveling in a different direction.³⁷ There are lots of further details to be learned, of course, but in this work we have the foundation for an account of the neural basis of motion perception. Any further question about how the neurons in MT cause experiences of motion is nothing but another version of the hard problem. We admit that we have nothing to say about that, but after all that was not our task anyhow. The claim that there is a further issue about temporal experience that we should be addressing but are not is simply unfair. So they say.

I think the scientists are being too modest here. There are further issues in the area, and if properly conceptualized they can be pursued empirically. Let me try to say something about what these issues are.

An important clue comes from the so-called miracle cures. These are cases in which vision is restored to a patient who has been congenitally blind or blind from a very early age. The first and most well-known patient of this sort is S. B., who was initially studied by Gregory and Wallace in the early 1960's.³⁸ S. B. had lost effective sight in both eyes as an infant due to corneal disease. At the age of 52 he received a corneal graft that restored the optical mechanisms in both eyes. There was naturally some startup time in learning to deal with his new kind of sensory input, but S. B. learned in remarkably short order to recognize various visual forms such as uppercase letters and the face of a clock. (His success in these areas, by the way, may shed some light on the famous Molyneux Problem, at least if the usual claim about S. B. is correct. The usual claim is that he learns to recognize shapes as quickly as he does because he is transferring his tactile understanding of shapes into the new visual domain.) Perhaps even more surprising than his ability to recognize shapes, S. B. learned to identify colors by their names without much difficulty, though it should be noted that even when he was blind he had the ability to detect the difference between light and dark.

Despite these great successes, however, it is the limitations of S. B.'s visual understanding that are important for us here. The central limitation was his difficulty, even after a long period of recovery, in dealing with visually presented moving objects.³⁹ He had difficulty both when the object was moving and he was standing still, and when the object remained stationary as he moved around it. As to the first of these, crossing the street became a terrifying affair for him since he could not track the motion of the cars. As to the second, it was very

³⁷ See, e.g., Dubner and Zeki 1971 and Zeki 1974.

³⁸ See Gregory and Wallace 1963. The material from this monograph, plus some additional material, is reprinted in Gregory 1974. An overview discussion of this case and others like it can be found in Wandell 1995, pp. 388-390.

³⁹ Also notable, and perhaps related, was his difficulty in recognizing depth.

difficult for S. B., as he walked around an object, to integrate the various perspectival presentations of it into an experience of a unified thing. As Gregory writes:

Quite recently he had been struck by how objects changed their shape when he walked round them. He would look at a lamp post, walk round it, stand studying it from a different aspect, and wonder why it looked different and yet the same.⁴⁰

The problem of seeing an object as unified throughout various perspectival presentations of it is precisely the kind of problem Kant was concerned with in his discussion of the threefold synthesis. As we saw, Kant believed that a certain kind of memory was needed in order to have this unified experience. Another miracle cure patient speaks to just this point. Valvo's patient H.S. described his difficulty in learning to read after his vision was restored:

My first attempts at reading were painful. I could make out single letters, but it was impossible for me to make out whole words; I managed to do so only after weeks of exhausting attempts. In fact, it was impossible for me to remember all the letters together, after having read them one by one. Nor was it possible for me, during the first weeks to count my own five fingers: I had the feeling that they were all there, but ... it was not possible for me to pass from one to the other while counting.⁴¹

Difficulties like these have led some researchers to posit a kind of short-term visual memory on the basis of which normal subjects are able to have experiences of moving objects of the sort that is extremely difficult after the miracle cure. This short-term visual memory is reminiscent of the kind of memory proposed by Kant and the kind of retention proposed by Husserl. As Wandell writes:

To perceive motion, the visual system must be able to integrate information over space and time. To perform this integration, one needs a means of short-term visual storage that can be used to represent recent information and visual inferences. If this visual storage fails, perhaps because it did not develop normally during early blindness, motion perception will be particularly vulnerable.⁴²

The kind of visual experience that requires this type of short-term visual storage seems to me to be quite distinct from the pure visual experience of motion. This is an empirical claim, but one that bears strongly on the puzzle of temporal experience. I do not know of any empirical work that has been done in this area, so it seems to me quite likely that the conceptual distinction between the pure visual experience of motion, on the one hand, and the kind of visual experience of a moving object that requires short-term visual storage, on the other, has not been made clear enough in the neuroscientific literature. Let me say a bit more about this distinction.

⁴⁰ Gregory 1974, p. 111, quoted in Wandell 1995, p. 389. The question in what sense it *looked* the same for S. B. is an important one. The lamp post looks *to be* a single, unified object to us as we walk around it. But presumably this is just what is lacking in S. B.'s experience. It is very possible that the claim that it looked the same to S. B. is misleading.

⁴¹ Valvo 1971. Quoted in Wandell 1995, p. 390.

⁴² Wandell 1995, p. 390.

There is a kind of pure visual experience of motion that can be isolated in the so-called motion aftereffect (sometimes known more popularly as the waterfall illusion). To induce the motion aftereffect, the subject habituates to a constantly moving scene, like a waterfall, and then focuses his gaze on a stationary object. One has the very strong impression of motion in the direction opposite to habituation, although it does not look as though the stationary object is moving. We might think of this aftereffect as the pure visual experience of motion. The aftereffect shows that it is possible to have the pure visual experience of motion in the absence of the experience of an object as moving. Recent fMRI studies show that the neural basis of the aftereffect is principally to be found in area MT.⁴³

The pure visual experience of motion seems to me not to depend on anything like short-term visual memory. There is no question of needing to remember which unified object has the property of moving, since the experience is not as of an object moving at all. It would be very interesting to know, therefore, whether miracle cure patients are susceptible to the motion aftereffect. As long as there is no cortical damage to MT, it seems reasonable to think they will be. If that's right, then their inability to perceive objects as moving, or to perceive an object as identical throughout perspectival variations, would not be due to an inability to experience motion *per se*. Rather, it would be due to an inability to keep track, in some kind of short-term visual storage, of the thing that is moving and thereby to experience it as identical over time.

If this hypothesis is right, then the miracle cure patients are likely to be very different from patients who suffer from akinetopsia. Akinetopsia is the deficit, sometimes called motion blindness, in which a subject is unable to perceive objects as moving despite stationary objects remaining more or less visible. In the standard example, a patient suffering from akinetopsia is unable to pour a cup of coffee because the level appears frozen even while the cup is filling up. Similarly, a car moving toward such a patient will appear to be first far away and then all of a sudden much closer, without any experience of its moving in between. Interestingly, akinetopsia seems to be tied closely to a deficit in cortical area V5 (MT). It can be induced temporarily, for instance, by disabling V5 through transcranial magnetic stimulation (TMS).⁴⁴

Such a deficit seems quite distinct from what is described by the miracle cure patients. The akinetopsic patient experiences a series of snapshots in each of which a unified object appears to be located successively in different places; the miracle cure patients, by contrast, have difficulty keeping track of the moving object as the same throughout variations. These two deficits seem to be dissociable: the akinetopsic patient has no trouble counting his fingers, for instance. If in addition the miracle cure patients are susceptible to the motion aftereffect, then we will be on the way to a double dissociation.

Admittedly what I have offered is a set of empirical hypotheses. These are the kinds of hypotheses one can make from the armchair but can prove only with good empirical science. But they are hypotheses that are made on the basis of a conceptual distinction, between the pure visual experience of motion and the visual experience of an object as moving, that seems not to

⁴³ Niedeggen and Wist 1998.

⁴⁴ Beckers and Homberg 1992.

have been emphasized in the empirical literature. To be sure, the motion aftereffect itself highlights this distinction. But until one thinks about the miracle cure cases in addition, it is all too easy to think that activity in MT is all one needs to experience Derek Jeter's throw as moving from shortstop to first. If that's what you think, then the puzzle of temporal experience really will seem like a version of the hard problem of consciousness. But it's much more interesting than that.

If it is right, as Kant and Husserl and Wandell all think, that some kind of memory or retention or storage is required for the experience as of a unified, temporally persisting object, then we can ask all sorts of further empirical questions about the psycho-physics of this kind of short-term visual memory. Because of the considerations adduced in §6, it seems to me that this short-term visual storage is unlikely to be assimilable to any familiar kind of memory. But what kind of experience is it, then? Is it, for instance, a version of iconic memory?⁴⁵ It seems to me very different from this as well. In iconic memory a subject can retain for short amounts of time a tachistoscopically presented visual image, and can read off some of its details after the fact. But the phenomenology of short-term visual storage seems not to be like this at all. Rather, what the miracle cure patients seem to be missing is tied much more directly to our visual concept of an object as a unified entity persisting through time. Can we isolate this kind of short-term visual memory and distinguish it from other kinds of visual phenomena? That is an important empirical question, and one which, if answered in the affirmative, would do an enormous amount to turn the puzzle of temporal experience into a genuine scientific problem.

⁴⁵ Sperling 1960.

References

St. Augustine, *Confessions*. See esp. Book 11.

Baldo, M.V. and Klein, S.A. (1995) Extrapolation or attention shift? *Nature* 378: 565-6.

Beckers, G. and Homberg, V. (1992) Cerebral visual motion blindness: transitory akinetopsia induced by transcranial magnetic stimulation of human area V5. *Proc R Soc Lond B Biol Sci.*, 249(1325): 173-8.

Broad, C. D. (1923) *Scientific Thought*, Bristol: Thoemmes Press, 1993. Originally published in 1923. See esp. chapter 10, "Date and Duration."

Broad, C. D. (1938) *Examination of McTaggart's Philosophy, Vol. II Part I*, London: Cambridge University Press. See esp. chapter 35, "Ostensible Temporality."

Dainton, Barry (2000) *Stream of Consciousness: unity and continuity in conscious experience*, London: Routledge.

Dennett, Daniel (1991) *Consciousness Explained*, Boston: Little, Brown and Co. See esp. the chapter on the perception of time.

Dubner, R. and Zeki, S. (1971) Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, 35: 528-532.

Eagleman, D.M. & Sejnowski, T.J. (2000) Motion integration and postdiction in visual awareness. *Science*, 287(5460): 2036-8.

Eagleman, D.M. & Sejnowski, T.J. (2001) The flash-lag illusion: distinguishing a spatial from a temporal effect, and why that matters for interpreting visual physiology. *Journal of Vision*, 1(3), 16a.

Eagleman, D.M. & Sejnowski, T.J. (2002) Untangling spatial from temporal illusions. *Trends in Neurosciences*, 25(6): 293.

Goodman, Nelson (1951) *The structure of appearance*, Cambridge, MA: Harvard University Press.

Gregory, Richard L. and Wallace, Jean G. (1963) *Recover from Early Blindness: a case study*, Experimental Psychology Society Monographs, no. 2.

Gregory, Richard L. (1974) *Concepts and Mechanisms of Perception*, London: Duckworth.

Hume, David (1739) *A Treatise of Human Nature*, L. A. Selby-Bigge (ed.) Oxford: Clarendon Press, 1978. Originally published in 1739. See esp. Book I, section III.

Husserl, Edmund (1893-1917) *On the phenomenology of the consciousness of internal time*, John Barnett Brough (tr.), The Netherlands: Kluwer Academic Publishers, 1980. Lectures given between 1893 and 1917.

James, William (1890) *Principles of Psychology, Volume I*, New York: Dover Books, 1950. Originally published in 1890. See esp. chapter 15, "The perception of time."

Kant, Immanuel (1781/1787) *Critique of Pure Reason*, Paul Guyer and Allen Wood (trs.), Cambridge: Cambridge University Press, 1998. Originally published in 1781/1787. See especially the A-deduction.

Köhler, Wolfgang (1947) *Gestalt Psychology: an introduction to new concepts in modern psychology*, New York: Liveright Pub. Corp.

Krekelberg, B. and Lappe, M. (2000) A model of the perceived relative positions of moving objects based upon a slow averaging process, *Vision Research*, 40: 201–215.

Le Poidevin, Robin (ed.) (1998) *Questions of time and tense*, Oxford: Clarendon Press.

Locke, John (1690) *An Essay Concerning Human Understanding*, Peter H. Nidditch (ed.), Oxford: Oxford University Press, 1975. Originally published in 1690. See esp. Book 2, chapter 14.

Mellor, D. H. (1998) *Real Time II*, London: Routledge.

Merleau-Ponty, Maurice (1945) *Phénoménologie de la Perception*, Paris: Gallimard.

Niedeggen, Martin and Wist, Eugene R. (1998) The Physiologic Substrate of Motion Aftereffects. In *The Motion Aftereffect: A Modern Perspective*, George Mather, Frans Verstraten, Stuart Anstis (eds.), Cambridge, MA: The MIT Press, pp. 125-155.

Nijhawan, Romi (2002) Neural delays, visual motion and the flash-lag effect, *Trends in Cognitive Sciences*. 6 (9): 387-393.

Nijhawan, Romi (1994) Motion extrapolation in catching, *Nature*, 370, 256-257.

Palmer, Stephen E. (1999) *Vision: from photons to phenomenology*, Cambridge, MA: MIT Press.

Pöppel, Ernst (1988) *Mindworks: time and conscious experience*, Tom Artin (tr.), Boston: Harcourt, Brace, Jovanovich.

Purushothamen, G. *et al.* (1998) Moving ahead through differential visual latency, *Nature*, 396: 424.

Russell, Bertrand (1913) *Theory of Knowledge: 1913 manuscript*, New York: Routledge, 1999. Written in 1913. See esp. chapter 6, “On the experience of time.”

Russell, Bertrand (1921) *The Analysis of Mind*, London: G. Allen & Unwin. See esp. chapter 9, “Memory.”

Sider, Theodore (2001) *Four-dimensionalism: an ontology of persistence and time*, Oxford: Clarendon Press.

Sperling, G. (1960) The information available in brief visual presentations, *Psychological Monographs*, 74: 1-29.

Valvo, A. (1971) *Sight Restoration after Long-Term Blindness: the problems and behavior patterns of visual rehabilitation*, New York: American Foundation for the Blind.

Wandell, Brian A. (1995) *Foundations of Vision*, Massachusetts: Sinauer Associates, Inc.

Whitney, David and Murakami, I. (1998) Latency difference not spatial extrapolation, *Nature Neuroscience*, 1, 656.

Whitney, D., Murakami, I. and Cavanaugh, P. (2000) Illusory spatial offset of a flash relative to a moving stimulus is caused by differential latencies for moving and flashed stimuli, *Vision Research*, 40: 137–149.

Wolff, Robert Paul (1963) *Kant's Theory of Mental Activity*, Cambridge, MA: Harvard University Press.

Zeki, S. (1974) Functional organisation of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey, *Journal of Physiology*, 236: 549-573.