

Memory and Perception: Remembering Snowflake

Jordi FERNÁNDEZ

BIBLID [0495-4548 (2006) 21: 56; pp. 147-164]

ABSTRACT: If I remember something, I tend to believe that I have perceived it. Similarly, if I remember something, I tend to believe that it happened in the past. My aim here is to propose a notion of mnemonic content-accounts for these facts. Certain proposals build perceptual experiences into the content of memories. I argue that they have trouble with the second belief. Other proposals build references to temporal locations into mnemonic content. I argue that they have trouble with the second one. I propose a notion of mnemonic Content that can account for the rationality of both beliefs.

Keywords: memory, perception, time, intentionality.

1. Introduction

Memories, like other mental states such as beliefs or perceptual experiences, are intentional states. Memories have content in the minimal sense that, if a subject remembers something, the world is thereby presented to her as having been in a certain way. The aim of this essay is to investigate how we should construe the content of memories.¹ For the sake of simplicity, I shall concentrate on those memories that we report as being about events in the outside world, such as those memories that we express by uttering sentences like ‘I remembered that there was a book on this table an hour ago’ or ‘I remember that you were at the party on Friday’. What sorts of things do memories like these put us in cognitive contact with? Are they, as our verbal reports suggest, events in the outside world? Are they mental events?

In section two, I will delimit the project. First, I will distinguish two varieties of memory, ‘episodic’ and ‘semantic’, and point out that our project only concerns the former variety of memory. Next, I will say more about what I mean by the content of a mental state by using the notion of truth-conditions. Then, I will specify the general form of a proposal about the content of memories. Finally, I will highlight two features of episodic memories that will constrain the task of determining what those memories are about. These features concern the relation between memory and belief. Basically, they amount to the fact that, when a subject episodically remembers something, she is inclined to believe two things. She is inclined to believe of the remembered event that it took place in the past. And she is inclined to believe that she has perceptually experienced it. I will take it to be an important virtue of any construal of the content of memories that it accounts for those two connections between memory and belief.

¹ I shall use the following locutions equivalently: ‘the content of A is that p ’, ‘ A is about p being the case’ and ‘ p being the case is the object of A ’ where ‘ A ’ stands for an intentional state and ‘ p ’ for an event.



Section three will be concerned with the belief that one has perceptually experienced the event that one apparently remembers. I will explore a certain kind of strategy for explaining the inclination to have this belief when one episodically remembers something. This type of strategy builds a reference to the subject's perceptual experiences into the content of her memory experiences. I will discuss two versions of this strategy and argue that one of them attributes too much information to memories whereas the other one attributes too little.

Section four will be concerned with the belief that the remembered event took place in the past. I will discuss a certain kind of account of the inclination to have this belief when one has an episodic memory. Within this account, a certain reference to the temporal position of the remembered event is built into the content of the corresponding memory. I will discuss two versions of this strategy and, once again, argue that one of them attributes too much information to memories whereas the other one attributes too little.

The upshot of section three will be that the content of a memory must involve a particular kind of reference to the subject's perceptual experiences. The upshot of section four will be that the content of a memory must involve a specific kind of reference to the temporal position of the remembered event as well. In section five, I will offer a positive proposal about the content of memories that incorporates both features. This proposal builds on some remarks by John Searle on the 'causal self-referentiality' of memory. And it is aimed at explaining both our inclination to believe of remembered events that they happened in the past as well as our inclination to believe that they were perceived by us.

A note on terminology. In what follows, I shall use expressions of the form '*S* remembers that *p*' and '*S* has a memory that *p*' to refer to a subject being in a state wherein she veridically remembers a certain fact. With the locutions '*S* apparently remembers that *p*', '*S* has an apparent memory that *p*' and '*S* has a memory experience that *p*', I will refer to a subject being in a state wherein she either veridically remembers or misremembers a certain event. (Likewise for perception talk.)

2. *Episodic memory and semantic memory*

There are two varieties of memory for events that it is important to distinguish, since our *explanandum* only concerns one of them. An example may be useful to tell them apart.

Suppose that, when I was a child, my parents took me to the zoo of Barcelona and I saw Snowflake (a charming albino gorilla who lived at that zoo for many years) playing there. There are two senses in which, later, I could remember that there was an albino gorilla playing at the local zoo. In the first sense, I remember it because I am having a certain quasi-perceptual experience whereby it appears to me as if there had been an albino gorilla playing there. Presumably, the experience I am having when I remember something in this sense involves some imagery of the event and it elicits some awareness of it in me. In this sense, I can remember that there was an albino gorilla playing at the zoo even if I did not form the belief that there was one at the time I

perceived him. Imagine, for instance, that I did not have the concept *gorilla* at the time I visited the zoo, or that I did not pay attention to him because I was fascinated by the elephant in the background. It is still possible for me to remember that there was an albino gorilla at the zoo in this sense.

There is a different sense in which I may remember that there was an albino gorilla playing at the zoo years after my visit there. In that different sense, I remember that there was an albino gorilla playing at the zoo if, in the past, I formed the belief that there is such an animal playing at the zoo, I now believe that he was playing there, and the reason why I believe it now is that I believed it then. In this sense, I can remember that there was an albino gorilla playing at the zoo even if, at the current time, I am not having any quasi-perceptual experience whereby it appears to me as if there had been an albino gorilla playing at the zoo and, therefore, I would not now be able to answer questions about any detail in the scene that I did not appreciate at the time I perceived it. In this sense, I cannot remember that there was an albino gorilla playing at the zoo if I did not form the belief that there was one at the time I perceived him.

Let us call the first kind of cognitive achievement ‘episodically remembering’ and the second kind ‘semantically remembering’. Likewise, if a subject remembers a certain event episodically, we may call the experience she has an ‘episodic memory’ of the event and, if she remembers it semantically, we can call her belief that the event happened a ‘semantic memory’ of it. In this discussion, I shall concentrate on episodic memory. Thus, when I speak of remembering an event, I shall mean episodically remembering unless otherwise specified. (Similarly, talk of memories *simpliciter* should be understood as referring to episodic memories.) The issue that will concern us is, then, how we should construe the content of episodic memories.

Like perceptual experiences, episodic memories have the power to elicit certain beliefs in us. Thus, when we apparently remember a certain event, we are disposed to believe that some period of time has passed between the time at which it happened and the present moment. Similarly, when we apparently remember a given event, we are inclined to believe that it was perceived by us. We can give a quite strong formulation to these two ideas in the following ‘attribution of pastness’ and ‘attribution of experience’ principles:

(AP) *Attribution of Pastness*

For any subject *S* and proposition *p*:

If *S* remembers that *p*, then *S* believes that, in the past, it was the case that *p*.

(AE) *Attribution of Experience*

For any subject *S* and proposition *p*:

If *S* remembers that *p*, then *S* believes that she has perceptually experienced that *p*.

Why are AP and AE quite strong as formulated? The reason is that both AP and AE should really be read as including implicit *ceteris paribus* clauses, since they could certainly have exceptions. Suppose that a given subject has a certain memory experience, but she also has what she takes to be overwhelmingly strong evidence that her experience is not a memory experience but, say, an episode of imagination. Then, she may form neither the belief that the event that she seems to be remembering happened in the past nor the belief that she perceptually experienced it. Nevertheless, in normal circumstances, it seems that memory experiences do elicit in us both the belief that the remembered events happened in the past and the belief that we have perceptually experienced them.

Both AP and AE should play a constraining role in our theorizing about the intentionality of memory. Any hypothesis about the content of episodic memories should, at least, respect those two principles. Ideally, a hypothesis about the content of episodic memories should illuminate why those two principles hold true. For methodological purposes, this means that we are entitled to require from any proposal about the content of episodic memories that the facts described by AP and AE do not emerge as mysteries if the proposal is right. And we can arbitrate among different proposals about the content of episodic memories based on which have some prospects of explaining either AE or AP, and which don't.

Let me now be more precise about what I mean by a proposal about the content of episodic memories here. An almost trivial point about memory is that memory experiences can be evaluated as true or false. For each memory experience, there are conditions under which it is true and conditions under which it is false. Thus, it is natural to think that, if you want to know what the content of a given memory experience is, you should ask yourself what it would take for it to be true. For the purposes of this discussion, it will be convenient to represent the truth-conditions of episodic memories by means of certain abstract objects, namely, propositions. I will construe propositions as ordered pairs of properties and objects. (As far as I can see, nothing in this discussion hinges on that view about the nature of propositions.) Thus, I will construe the proposition that captures the truth-conditions of a certain memory experience M as an ordered pair of an object o and a property P , where o having P is meant to be what it takes for M to be true. I shall refer to such pairs with expressions of the form ' $\langle P, o \rangle$ '. The issue that will concern us is, in those terms, what sort of objects and properties are the constituents of those propositions that capture the truth-conditions of episodic memories.

Any given specification of the content of a memory experience should attribute just the right amount of information to it. By 'the right amount' I simply mean the following. If it is the case that, according to our pre-theoretic intuitions about memory, a given memory experience is true of a certain possible situation, then our theory should attribute to it such truth-conditions that the experience turns out to be true of that situation. Conversely, if it is the case that, according to our pre-theoretic intuitions about memory, a certain memory experience is false of a given possible situation, then our theory should attribute to it such truth-conditions that the experience turns out to

be false of that situation. We may call this the ‘right amount of information’ test, or RAI. To endorse the right amount of information test is partly to accept that, to the question ‘Would such-and-such possible situation make this memory experience true?’ one may not answer ‘It depends on what we take the content of that memory experience to be.’ Our pre-theoretic intuitions about veridicality come first. And content must be tailored according to those intuitions.

Thus, there are two types of constraints that we will need to respect when we try to find out what the constituents of those propositions that capture the contents of episodic memories are. On the one hand, it will be an important virtue of any hypothesis about those constituents that it accounts for the connections between memory and belief expressed by AE and AP. On the other hand, any such hypothesis will need to pass the right amount of information test.

3. Perception and Memory

Let us start our investigation at the attribution of experience principle. If one is interested in accounting for the truth of AE, then a natural way of construing mnemonic content is to build a reference to the subject’s perceptual experiences into the content of her corresponding episodic memories. That is, one can try to account for AE by making a perceptual experience of a certain event one of the constituents of the content of the memory experience whereby a subject apparently remembers the event in question. Let us call this approach the ‘reflexive approach’ to mnemonic content. The reflexive approach comes in many varieties. The most popular seem to be two theories that we may call the ‘causally self-referential view’ and the ‘neutral view’ of memory. Let us consider each of them in order and evaluate them with respect to AE, AP and RAI.

According to the causally self-referential view of memory (hereafter CSR), the intentional object of a memory experience is a complex causal relation, whose *relata* are the following. One of them (the cause) is the fact that the event reported as being remembered caused the subject to have a perceptual experience of it. The other one (the effect) is the event that consists in the subject having a certain memory experience, namely, the very experience whose content is constituted by the just specified causal relation. The view can be formulated thus:

CSR

For every subject *S*, memory experience *M* and proposition *p*:

If *S* has a memory experience *M* that she would express by saying that she remembers that *p*, then the content of *M* is the proposition

<Being caused by a perceptual experience of *p* in turn caused by *p* being the case, *M*>

John Searle seems to have had this view in mind when he claimed that memory is causally self-referential. Searle introduces causal self-referentiality as a feature of the

intentionality of perception. Commenting on the satisfaction conditions of a visual experience, he claims:

It is part of the conditions of satisfaction (in the sense of requirement) of the visual experience that the visual experience must itself be caused by the rest of the conditions of satisfaction (in the sense of things required) of that visual experience. Thus, for example, if I see the yellow station wagon, I have a certain visual experience. But the Intentional content of the visual experience, which requires that there be a yellow station wagon in front of me in order that it be satisfied, also requires that the fact that there is a yellow station wagon in front of me must be the cause of that very visual experience. Thus, the Intentional content of the visual experience requires as part of the conditions of satisfaction that the visual experience be caused by the rest of its conditions of satisfaction, that is, by the state of affairs perceived. [...] The intentional content of the visual experience therefore has to be made explicit in the following form:

I have a visual experience (that there is a yellow station wagon there and that there is a yellow station wagon there is causing this visual experience). (Searle 1983, p. 48.)

Searle claims that this is not only a feature of the intentionality of perception but it is a feature of other forms of intentionality as well, such as memory. Searle does not really elaborate on this idea, but the following remark strongly suggests CSR:

The memory of seeing the flower represents both the visual experience and the flower and is self-referential in the sense that, unless the memory was caused by the visual experience which in turn was caused by the presence of (and features of) the flower, I didn't really remember seeing the flower. (Searle 1983, p. 85.)

According to a view substantially different from CSR, the neutral view, the objects of our episodic memories are our perceptual experiences of past events. Thus, episodic memories turn out to be neutral on whether those events 'out there', in the world, that we claim to remember did happen or not. According to neutral theorists, all memory tells us is that we had perceptual experiences of those events. It does not really tell us that those events did happen. We may formulate the view as follows:

The neutral view

For every subject S , memory experience M and proposition p :

If S has a memory experience M that she would express by saying that she remembers that p , then the content of M is the proposition

<Having had a perceptual experience of p being the case, S >

This view seems to have been held, among others, by Wolfgang Von Leyden:

It is also true that our memories, veridical and non-veridical alike, always appear to be about some objective past event or fact, not about our past perceptions. [...] The best way to convince ourselves of the fact that *our recollections concern our perceptions of past events, not those events themselves*, is to take the case of a distorted or delusive perception of a certain past event and ask ourselves in what sense, if any, any subsequent memory purporting to have this event as its object is mistaken. (Von Leyden 1961, p. 61.)

How well do CSR and the neutral view fare *vis à vis* AE? It is easy to see that both of them can account for the attribution of experience principle being true. For both of them take perceptual experiences of past events to be constituents of those propositions that capture the contents of our current memory experiences of those events.

(Actually, CSR takes several other things to be constituents of mnemonic contents as well. But it does not need to appeal to them in order to explain AE.) According to both views, the fact that I seemed to perceive them is part of what I am in cognitive contact with when I have memory experiences of those events. Thus, the fact that I am inclined to believe of those events that I now remember that I have perceptually experienced them is, within either CSR or the neutral view, not at all surprising.

Unfortunately, the ways in which the two theories build perceptual experiences into the content of episodic memories raise important difficulties for their prospects of passing the right amount of information test. As it turns out, CSR happens to attribute too much information to episodic memories. The neutral view happens to attribute too little. I turn to those difficulties now.

Let us consider CSR first. Consider the above-mentioned situation where I am standing at a certain spot in the zoo, looking at an albino gorilla who is playing there and, years later, I have a memory experience that would I express by saying that there was an albino gorilla playing at the zoo. Let us call this situation W1. Let P1 be the perceptual experience I have in W1 and let M1 be the memory that, in W1, I have years later. Now, consider a possible situation W2 that is, let us stipulate, exactly like W1 except for the fact that, in W2, the perceptual experience that I had in the past was not caused by the presence of the gorilla. I did have an experience that is phenomenologically indistinguishable from P1 when I looked at the gorilla, and an experience that is phenomenologically indistinguishable from M1 years later. (We may call them P2 and M2 respectively.) Plus the gorilla was there, behaving exactly as he was in W1. But the experience I had in W2 when I was looking at the gorilla was not caused by his presence but by something else.

Let us concentrate on M1 now. The question that RAI raises about W2 is: Is M1 true of W2 or not? In other words, is W2 accurately represented by M1 or not? It is clear that, according to CSR, M1 is false of W2. What it takes for any situation to be accurately represented by my memory experience is that it has a certain causal history in that situation. And, in W2, my memory experience does not have the right kind of causal history. Thus, W2 turns out not to be accurately represented by M1. Our intuitions, though, seem to be that M1 is true of W2. W2 is, intuitively, one of those situations that M1 represents accurately. After all, the gorilla was, in W2, playing exactly as he was playing in W1. In W2, I was looking at him from the spot from which I was looking at him in W1. And he was looking to me just like he was looking to me in W1. So the only difference between the past scenes in W1 and W2 is that a causal relation between the gorilla and my perceptual experience is missing in W2. And it does not seem that this makes an intuitive difference to the accuracy of the respective memory experiences. So the relevant portion of my past life in W2 seems to be accurately represented by M1. Yet CSR entails that M1 is false of W2. In this sense, CSR seems to attribute too much information to episodic memories.

The neutral view has trouble passing RAI as well. The reason is that, according to this view, the truth-conditions of our memory experiences do not have anything to do with whether the apparently perceived events that we report to remember happened

or not. Whether our memory experiences are true or not depends, strictly speaking, on whether we had the appropriate perceptual experiences in the past. It does not depend on whether those events that were apparently perceived actually took place. For this reason, it attributes too little information to memories. It is ironic that Von Leyden's thought experiment in the passage above illustrates this point rather nicely.

Imagine, with Von Leyden, a possible situation where I have a false perceptual experience of an albino gorilla playing at the zoo. Call this situation W3 and call the perceptual experience I have in it P3. Consider the subsequent memory experience that, in W3, I would express by saying that there was an albino gorilla playing at the zoo. Let us call this experience M3, and ask ourselves whether M3 is false of W3 or not. Von Leyden thinks that this sort of experience only 'purports', or 'appears', to have an objective event as its object, so his position would surely be that M3 is not mistaken. This is quite surprising. Suppose that my perceptual experience was false because, as it turns out, I was really looking at a polar bear. Then, the memory experience whereby it now seems to me that there was an albino gorilla playing at the zoo is surely just as wrong as my past perceptual experience.² Yet, according to the neutral view, all it takes for a given situation to make that memory experience true is that, in that situation, I seemed to perceive an albino gorilla playing at the zoo. So M3 turns out to be true of W3, which is quite counter-intuitive. In this sense, the neutral view attributes too little information to episodic memories.

The upshot is that, even though reflexive theories of memory can easily account for AE, they do not pass RAI. Thus, whether or not reflexive theories can account for AP, they do not seem to deliver good candidates for mnemonic content. However, there is a positive lesson that we may learn from our discussion of the reflexive approach, namely, that building past perceptual experiences into the contents of corresponding memory experiences is a good strategy to explain why AE holds true. We should keep the difficulties that RAI raised for reflexive views in mind, though. They illustrate that building perceptual experiences into the content of a memory experiences is risky business. On the one hand, the perceptual experience cannot appear in mnemonic content as a mere experience that may or may not correspond to the world. On the other hand, it cannot appear as part of a causal chain that leads from the remembered event to the memory experience either.

4. *Time and memory*

Let us turn to the attribution of pastness now. If one is interested in accounting for AP, then a straightforward way of specifying the content of episodic memories is by building a reference to the temporal position of a remembered event into the content of the corresponding memory experience. That is, one can try to account for AP by

² It is important not to collapse two issues here. One of them is whether M3 is false of W3 or not. The other one is whether memory malfunctioned in W3 or not. Clearly, memory did not malfunction in W3. But the faculty of memory need not be at fault for a certain memory to be false. The W3 scenario is, so to speak, a case of inheritance of a mistake, not a case of a mistake in inheritance.

making a certain moment in time associated to a given event one of the constituents of the proposition that captures the truth-conditions of the memory experience whereby one remembers that event. Let us call this approach the ‘temporal approach’ to mnemonic content. The temporal approach comes in many varieties, depending on how one thinks of the reference to the temporal location of a remembered event. The most natural versions of this approach seem to be two theories that we may call the ‘absolutist’ and ‘relativist’ views of memory. Let us consider them in order.

According to the absolutist view, when a subject seems to remember a given event, the event in question is presented to her as occupying a certain temporal location independently of the temporal position at which the subject has her memory experience.³ Descriptions of temporal positions of this sort will include expressions such as ‘in 1985’, ‘on Monday’ or ‘at noon’ (as opposed to ‘twenty years ago’, ‘yesterday’ or ‘two hours ago’). The view can be formulated as follows:

Absolutist view

For every subject S , memory experience M and proposition p :

If S has a memory experience M that she would express by saying that she remembers that p , then there is a non-indexically described time t such that t is earlier than the time of M 's occurrence and the content of M is the proposition
 <Occurring at t , the fact that p >.

This formulation of the view is meant to leave two questions open. The first of them is: Suppose that the fact that p took place at time t^* . How close should time t be to t^* in order for a memory experience M whose content is <occurring at t , the fact that p > be veridical? The reason why this question should be left open is that the answer to it seems to vary from memory experience to memory experience. It depends, for instance, on the amount of time that has passed between the occurrences of p and the subject's having M . The shorter that time is, the more inclined we are to require that t must be close to t^* in order for M to be considered veridical.

Compare the following two examples. Suppose that we are enquiring about the truth conditions of the memory experience that I would express by saying that I remember starting to write this paper. And suppose that I started to write it in late September 2005. That was only a month ago. (It was at the time I write this, anyhow.) That is why we are inclined to think that, if the memory experience in question represented the event as happening in early August and, so to speak, missed for some weeks, it would not be veridical. If our sympathies lie with the absolutist, we will take the content of that memory experience to be appropriately captured by something like <Occurring at the end of September 2005, the fact that I started to write ‘Memory, Perception and Time’> instead. By contrast, suppose that we are enquiring about the

³ This seemed to be the guiding idea behind the so-called ‘time-tagging theories’ of memory in psychology. See Glenberg and Swanson (1986), for instance.

truth conditions of the memory experience that I would express by saying that I remember riding a bike for the first time. Suppose it happened in, say, late 1978. Even if we granted to the absolutist that there is a temporal reference built into the content of memory experiences and that the reference in question is independent from our temporal location, we would not want to count as false a memory experience that represented the event as happening in early 1979 (and missed for a few months, as it were). This is because the event happened many years before my apparent memory of it takes place. To the extent that the idea of an absolute temporal reference is plausible, we must allow for considerable flexibility in the degree of accuracy that we require from that temporal reference. Surely the absolutist must claim that, the more recent the remembered event is, the more accurate that temporal reference in the content must be. Otherwise, the absolutist view is a non-starter.

The second question that the formulation of the absolutist view is meant to leave open is: Must the moment in time that is partly constitutive of the content of a memory experience be mentioned in a date-like manner, as the examples above suggest, or can it be mentioned by reference to other remembered events? The reason why this should be left open is that the answer to that question seems to vary from memory experience to memory experience as well. Thus, suppose that we are enquiring about the truth conditions of the memory experience that I would express by saying that I remember starting to write this paragraph. And suppose that it's 2:48 pm and I started writing it a minute ago. We are inclined to think that a memory experience that represented the event accurately (and absolutely, that is, without making reference to my temporal position) would not need to mention the temporal position of the event as 2:47 pm in order to be veridical. Intuitively, it would be sufficient to refer to it as happening before some other remembered event, such as my checking my email. Thus, it seems that the absolutist will need to allow that, in some cases, the temporal locations of seemingly remembered events are presented as relative to those of other seemingly remembered events whereas, in other cases, they are presented in a date-like manner.

The two questions above raise many interesting issues about our conception of the structure of time and the relations between memory and that conception which, unfortunately, I cannot pursue here. For the purposes of this discussion on mnemonic content, the important feature of the absolutist view to keep in mind is that, according to it, a reference to the temporal location of an event is one of the constituents of the content of a memory experience of that event. And that reference does not mention the temporal position of the remembering subject, even though it may very well make reference to other remembered events.

The absolutist seems to have reasonable prospects of accounting for AP, even though this would be far from an easy task. Recall that there are two kinds of temporal references that the absolutist accepts: temporal references such as 'at 2:47 pm' and temporal references such as 'right before I checked my email.' In either case, the absolutist needs to claim that, ultimately, a subject who remembers, say, starting to write a certain paragraph is inclined to believe that this happened in the past because she has certain beliefs about her own temporal location at the time of remembrance. I take it

that this is obvious in the former case. Why would a memory that presents my starting to write this paragraph as happening at 2:47 lead me to believe that I started it in the past if I did not think that it is after 2:47 pm now? Things are more complicated in the latter case, though.

Suppose that my starting to write the paragraph is presented to me as happening right before I checked my email and I believe that I checked my email in the past. Then, it is not surprising that I am inclined to believe that I started to write the paragraph in the past when I remember doing it. But the question of how I came to hold the belief that I checked my email in the past arises at that point. If this is a memory-based belief, then I may have arrived at it because I remember the fact that I checked my email as happening before some other event, which I believe to have happened in the past. But we may then ask how I came to believe that the event in question happened in the past, and so on. Eventually, this sort of chain must end at some event whose temporal location is represented in a date-like manner (if it is not to be represented by reference to my own temporal location, that is). And this takes us back to the former case where, as we have just seen, the absolutist must claim that AP holds because the subject has beliefs about her own temporal position. To the extent that the absolutist can make a case for the view that a subject has beliefs about her temporal position when she believes of those events that she apparently remembers that they happened in the past, the absolutist can account for AP through the just-sketched line of thought.⁴

Unfortunately, the absolutist view does not have an easier time passing RAI than reflexive theories did. The absolutist view seems to attribute too little information to memory experiences. There seem to be memory experiences and possible situations such that, intuitively, we would not count those situations as being accurately represented by those experiences and, yet, the proponent of the absolutist view must claim that they are. Recall situation W1 where I have a perceptual experience P1 whereby I seem to perceive an albino gorilla playing at the zoo. And, years later, I have memory experience M1, which I would express by saying that there was an albino gorilla playing at the zoo. Let us stipulate that, in W1, P1 occurs on 7/2/1980 and M1 occurs on 11/7/ 2005. Now, consider a possible situation W4 where I do not exist but the gorilla is playing at the zoo on 7/2/1980, exactly as he was playing in W1.

Since the advocate of the absolutist view claims that the content of M1 makes no reference to my temporal location, she is committed to the claim that M1 is true of W4 if it is true of W1. Whatever time the absolutist builds into the content of M1 in order to make it true of W1, the resulting content will make M1 true of W4 as well. Why is that? The reason is that whether or not I exist in a given possible situation is,

⁴ Contrary to what it may seem at first glance, the task is not hopeless. There is considerable evidence of internal processes that keep track of time through different cycles. (On internal timers, see Church (1984) as well as Friedman (1990).) Perhaps an argument could be constructed to the effect that such internal processes make beliefs about our own temporal position available to us. Thus, in Tzeng (1976), it is hypothesized that the outputs of some organic pacemaker might be associated with perceived events, which could encode temporal information for later retrieval in memory.

according to the absolutist view, irrelevant to whether that situation is accurately represented by a given memory experience of mine. Thus, whether or not I exist in W4 is irrelevant to whether M1 is true of W4 or not. All that matters is that the gorilla is playing at the time that he is represented as doing it in the content of M1. And, given that the gorilla is playing at the same time in both W1 and W4, there is no way of construing the content of M1 so that it can be true of W1 but false of W4 (no way available to the absolutist anyway). So the absolutist must conclude that M1 is true of W4 if it is true of W1.

However, our intuitions are that M1 does not accurately represent W4 even though it does represent W1 accurately. Why doesn't it represent W4 accurately? The reason is that no fact in W4 corresponds to a certain aspect of the information provided by M1. The relevant piece of information is that the apparently remembered event happened *in the past*. Granted, in W4, the albino gorilla was playing at the moment at which he was playing in W1. But why should that day count as being in the past in W4? Clearly, no fact about 7/2/1980 makes that date, as opposed to a date forty years later, a past date in a possible world where I do not exist. Any date, in any possible situation, counts as a past date (or a future date, for that matter) only relative to the time that counts as the current time at that situation, that is, the time at which I exist there. If I do not exist in W4, then there seems to be no point of reference relative to which 7/2/1980 can count as a past date. Thus, it seems that nothing in W4 corresponds to a very salient piece of information that M1 is carrying, namely, that the event that I seem to remember happened in the past. The upshot is that W4 is not a truth-maker of M1, which suggests that the absolutist view attributes too little content to memory experiences.⁵

According to a different version of the temporal approach, the relativist view of memory, when a subject seems to remember a given event, the event in question is presented to her as occupying a certain temporal location relative to her own. The view can be formulated as follows.

Relativist view

For every subject *S*, memory experience *M* and proposition *p*:

If *S* has a memory experience *M* that she would express by saying that she remembers that *p*, then there is a period of time *T* such that the content of *M* is the proposition

<Occurring *T*-earlier than now, the fact that *p*>.

⁵ In (2006), I used a scenario like W4 against CSR. There, I basically argued that CSR made M1 false of W4 even though, intuitively, M1 is true of W4. I was not concerned with issues about time in that essay. And I now think that, when we consider the fact that memories present events as being in the past, it is not intuitive at all that W4 is accurately represented by M1. Thus, I no longer think that CSR is vulnerable to the objection I raised in (2006).

The intuitive idea in the relativist view is that a memory must represent a certain event as happening a certain amount of minutes, hours, days or years ago. In either case, it must always represent it as being separated from the instant of remembrance by a certain period of time. Basically, it represents it as being this much or that much *earlier than now*.

There are some intimations of this view in Edmund Husserl's writings on memory. Thus, he writes (emphasis is mine):

I remember the lighted theatre of yesterday [...]. Accordingly, the theatre hovers before me in the representation as something actually present. I mean this, but at the same time I apprehend this present as lying back *in reference to the actual present* of perceptions now extant. [...] What is remembered appears as having been present, that is, immediately and intuitively. And it appears in such a way that a present intuitively appears *which is at an interval from the present of the actual now*. (Husserl 1964, p. 82.)

The question about veridicality and the exact length of that period of time arises here as well. Suppose that p took place two months before the time at which it is apparently remembered. Would a memory experience whose content is <occurring three months ago, the fact that p > be veridical? Would a memory experience whose content is <occurring two months and a half ago, the fact that p >? Our intuitions suggest that the answer to this kind of question varies from memory experience to memory experience. And, once again, the more distant in time the seemingly remembered event is, the more inclined we seem to be to relax the degree of accuracy that we require from the period of time built into the content of the relevant memory experience. Thus, it does not seem that any reasonable relativist theory could offer, for each event p , an exact degree of accuracy that a period of time T must satisfy in order for the content of a memory of p to be plausibly construed as <Occurring T -earlier than now, the fact that p >. Here, like in our discussion of the absolutist theory, we need to concede, for the sake of the argument, that if references to periods of time are constituents of the contents of memories, then there must be significant variation as to how precise such references are from memory experience to memory experience. For our purposes, the important point to remember about the relativist view is that, according to the view, those temporal references are relative to the temporal positions of the remembering subjects.

The relativist seems to have reasonable prospects of accounting for AP. If a memory experience of mine presents a certain event as happening earlier than now (no matter how much earlier), then I will be disposed to believe that the event happened earlier than now. And if I am disposed to believe that it happened earlier than now, then it is no wonder that I am disposed to believe that it happened in the past as well. There may be certain beliefs about the structure of time such that, if I held those beliefs, then they could prevent me from having the belief that the event happened in the past when I believe it to have happened earlier than now. But the fact of the matter is that, ordinarily, we have no such beliefs. We take it for granted that anything that happened earlier than now must have happened in the past. So AP emerges as a quite natural phenomenon if the content of memories is the content that the relativist attributes to them.

However, it seems that the relativist view, like the rest of views considered so far, does not pass RAI. The relativist view seems to attribute too much information to memory experiences: There are memory experiences and possible situations such that, intuitively, those memory experiences are true of those situations and, yet, according to the theory, they turn out to be false of them.

Consider the following case. In a certain possible situation W_5 , an albino gorilla is playing at the zoo at a certain moment in time, call it t , and you perceive him to be. Some time later, you travel back in time. In fact, you travel to some period in time before t . And, at some point during that period, you have a memory experience that you would express by saying that you remember that an albino gorilla was playing at the zoo. Call the experience M_5 . Now, is M_5 true of W_5 or is it false of W_5 ? Is W_5 one of the truth-makers of M_5 ? Intuitively, M_5 is true of W_5 . After all, your perception of the gorilla is something that constitutes part of your past life. It has certainly not become a figment of your imagination just because you traveled back in time. It represents an event that belongs to your past. Your experience of that event is part of your life. Intuitively, then, M_5 accurately represents W_5 . So whatever proposition expresses the content of M_5 , it should give M_5 truth-conditions according to which it is true of W_5 . However, the content posited by the relativist makes M_5 false of W_5 .⁶ The reason is that, in W_5 , the gorilla is not at the zoo before the time that counts as ‘now’, that is, the time at which you have M_5 . The gorilla is at the zoo at t and, therefore, after you have M_5 . Thus, we are forced to conclude, quite counter-intuitively, that M_5 is not representing the situation accurately. The upshot is that the relativist view delivers truth-conditions that are too demanding for memories.

The bottom line about the temporal approach is that temporal views seem to be capable of accounting for AP. However, they do not offer good candidates for mnemonic content, whether or not they may be able to account for AE. The reason is that relativist views turn out to be too strict about the content that they attribute to memories whereas (ironically) absolutist views turn out not to be strict enough. A more constructive moral that we may draw from our discussion of the temporal approach is that building a temporal location into the content of a memory experience is a promising strategy for the purposes of explaining AP. That being said, the troubles that temporal views had with RAI suggest that one must proceed with extreme care while pursuing this strategy. The reference to the temporal location of the remembered event must be relative to the remembering subject in order to constitute a reference to the past. But it cannot amount to the property of happening at a certain location relative to the location of the relevant memory experience in objective time.

⁶ There are a number of contents that the relativist may attribute to M_5 , of course. The period of time that, according to the content that a relativist may attribute to M_5 , separates the gorilla from you may be longer or shorter. The objection I am raising here is independent of the precise length of the period of time that the relativist may build into the content of M_5 .

5. *A positive proposal*

Let us take stock and reflect on what our discussions of the temporal and reflexive approaches have taught us.

The reflexive approach tells us that perceptual experiences can be constitutive elements of the contents of memories. The difficulties that the neutral view faced, though, suggest the following trade-off: On the one hand, construing mnemonic content in such a way that the content of a memory experience of an event makes no reference to our having perceptually experienced it makes accounting for AE a challenging task. On the other hand, simply building a perceptual experience into the content of a memory experience does not let us take the intuition that most of our memory experiences are about facts ‘out there’, in the world, at face value. We do not want to posit a kind of mnemonic content that forces us to conclude that memory cannot, as it were, reach beyond our skin. Making perceptual experiences part of the content of memories helps with AE, but we need to keep an eye on the fact that memory experiences are, so to speak, opinionated. Once we seem to remember a certain event, the question of whether it seems to us that the event really happened or not is no longer open to us. Memory experiences are assertive in that, when we seem to remember an event, that event is presented to us as having been the case, not only as having been perceptually experienced to be the case.

This leaves us with a first dilemma, which we may call the ‘world dilemma’. Either we build perceptual experiences into mnemonic content or we do not. In the former case, it is easy to account for AE but, if we do it as the neutral theorist does, then we will end up concluding that memories are never about the outside world. In the latter case, we may preserve the commonsensical idea that memory can put us in cognitive contact with events that happened in the external world, but we will have a rather difficult time accounting for AE.

The temporal approach tells us that temporal positions can be constitutive elements of the contents of memories. The problems of the absolute view, though, suggest that those temporal positions must be positions relative to the remembering subjects. Otherwise, it is hard to make sense of the idea that, in memory, events are presented to us as being in the past. This is the reason why the relativist view is appealing. Suppose that events are indeed presented to us as occupying certain temporal locations relative to ours when we apparently remember them. This would vindicate the idea that a robust feature of mnemonic content is that remembered events are presented to us as being in the past. However, we have seen that building earlier-than temporal positions into mnemonic content makes the truth-conditions of memories too strict. The reason is that references to temporal positions of that sort seem to carry with them information about the relative positions of the remembered event and the remembering subject in objective time. And this seems too much information to attribute to memories.

This leaves us with a second dilemma, which we may call the ‘time dilemma’. Either we do not build temporal locations into mnemonic content or we do. In the former case, it is going to be pretty hard to account for AP. In the latter case, it is possi-

ble to account for AP. But the most promising strategy for this purpose seems to involve building temporal positions relative to the subject's position into the content of her memory experiences. Basically, these constituents of mnemonic content come down to the properties of happening this much or that much earlier than the memory occurs. And, as it has just been mentioned, this raises significant difficulties for the view regarding RAI.

CSR enjoys certain attractive features that may allow us to get out of both the world and time dilemmas. Unfortunately, we saw that CSR had trouble with RAI and veridical hallucination scenarios. To get around that kind of trouble, let us consider, instead, the following 'veridically self-referential view', or VSR. We can formulate it as follows.

VSR

For every subject S , memory experience M and proposition p :

If S has a memory experience M that she would express by saying that she remembers that p , then the content of M is the proposition

<Being caused by a veridical perceptual experience of p being the case, M >.

VSR does not have the difficulty that CSR had passing the right amount of information test. Since, according to VSR, the content of a memory experience does not require that there is a causal link between the event that we claim to remember and our perceptual experience of it, VSR squares with the intuition that, in the kind of veridical hallucination scenario discussed in section 3, M1 represents the situation accurately. All VSR requires from a possible situation in order to make M1 true is that, in that situation, the memory experience is caused by a correct perceptual experience of an albino gorilla playing at the zoo. And, in W2, my perceptual experience of the albino gorilla is true and it is causing my memory experience. So VSR can accommodate the intuition that M1 is true of W2.

In addition, VSR borrows from CSR two features that allows the view to dissolve the world and time dilemmas. According to VSR, *veridical* perceptual experiences are constitutive elements of mnemonic content. This means that the content of a memory experience can refer to events out there, in the world, while referring to perceptual experiences of them at the same time. If an event is, as VSR claims, presented to us in memory as having been perceived, then it must be presented to us as having been perceptually experienced and it must be presented to us as having been the case. The former aspect of mnemonic content allows us to account for AE. And the latter aspect of it allows us to overcome the difficulties of the neutral view. To appreciate the very last point, let us go back to possible situation W3. In that situation, I misperceive an albino gorilla playing at the zoo and, years later, I have a memory experience M3 that I would express by saying that I remember that an albino gorilla was playing at the zoo. Intuitively, M3 is false of W3. And VSR can account for this intuition: The

perceptual experience that, in the past, I had at the zoo was not veridical. Therefore, M3 turns out to be, according to VSR, false of W3. Which is what we wanted.

In addition, according to VSR, the content of a memory experience refers to the experience's own causal history. There is one leg of the causal history of memory experiences that CSR builds into their contents that is missing from the kind of content that VSR attributes to memory experiences. This is the leg from the perceptual experience of the event that we report to remember to the event in question. But VSR does build part of the causal history of a memory experience into its content, namely, the causal history of that experience up to the perceptual experience of the event that is reported as being remembered. This provides us with a sense in which that experience can refer to the subject's past, or her past life, even in time travel scenarios. Thus, consider W5 again. In W5, you have a memory experience M5 that you would report by claiming that an albino gorilla was playing at the zoo. But something happened at some point between your having a perceptual experience of the gorilla, which we may label P5, and your having M5: Unbeknownst to you, you traveled back in time to a moment before the gorilla is playing at the zoo. Now, in W5, your perceptual experience of the gorilla was correct, and it caused that experience that you are having when you report to remember the gorilla. Thus, according to VSR, your memory experience is representing the situation accurately. This squares with our intuition behind the idea that time-travel does not falsify the traveler's memory experiences. This is the intuition that memory provides us with information about *one's own* past rather than information about the past. The upshot of this is that VSR can explain AP as well as AE. Suppose that an event is, in memory, presented to us as matching a certain perceptual experience that caused the very memory experience that we are having. Then, it is presented as occupying a certain position in our personal history. That history goes (perhaps nomically necessarily) hand in hand with objective time. However, it is not logically necessary that the order of events in our personal history matches the order of events in objective time, as time-travel scenarios illustrate. The order of personal time is the order of causation and, to the extent that backwards causation is conceivable, the order of personal time does not logically need to match the order of objective time.⁷ What VSR tells us is that memory represents events as being in our own personal past, that is, as being earlier in our history. Since this corresponds to earlier in objective time, as a matter of natural law, it is not surprising that we are inclined to believe of those events that they happened in the past. But relative positions in objective time are not built into mnemonic content according to VSR. This is how VSR can avoid the difficulties of the relativist view regarding time-travel. The bottom line is that VSR explains AE by building perceptual experiences into the contents of memories and it explains AP by building causal relations into them. The perceptual experiences at issue are veridical, which allows VSR to overcome the difficulties of the neutral view. And the causal relations do not include a causal relation between perceptual experiences and events in the world, which allows it to over-

⁷ For more on these issues, see Lewis (1976).

come the difficulties of CSR. In addition, those causal relations make reference to the remembering subject by making reference to the very memory experiences whose contents they partly constitute. This allows VSR to overcome the difficulties of the absolutist view. And, finally, those causal relations provide us with a sense in which events can be represented, in memory, as being earlier in the order of events that corresponds to the subject's history without being, necessarily, earlier in the order of objective time. This allows VSR to overcome the difficulties of the relativist view. Thus, it seems that not only does VSR account for AE and AP, but it also seems to pass the right amount of information test with a pretty decent grade.

REFERENCES

- Church, R.M. (1984). 'Properties of the internal clock', in J. Gibbon and L. Allen (eds.), *Timing and Time Perception*. New York: New York Academy of Sciences.
- Fernández, J. (forthcoming). 'The intentionality of memory', in the *Australasian Journal of Philosophy*.
- Friedman, W. (1990). *About Time: Inventing the Fourth Dimension*. Cambridge: MIT Press.
- Glenberg, A.M., and Swanson, N.G. (1986). 'A temporal distinctiveness theory of recency and modality effects', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12, 3-15.
- Husserl, E. (1964). *The Phenomenology of Internal Time-Consciousness*. (Ed. M. Heidegger). Bloomington: Indiana University Press.
- Lewis, D. (1980). 'Veridical Hallucination and prosthetic vision', *Australasian Journal of Philosophy* 58/3, 239-49.
- (1976). 'The paradoxes of time travel', *American Philosophical Quarterly* 13, 145-152.
- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Tzeng, O.J.L. (1976). 'A precedence effect in the processing of verbal information', *American Journal of Psychology* 89, 577-599.
- Von Leyden, W. (1961). *Remembering: A philosophical problem*. London: Duckworth.

Jordi FERNÁNDEZ is a research fellow at Macquarie University. He works mainly in epistemology and the philosophy of mind, and he has interests in continental philosophy as well. His main projects include writing a series of articles on self-knowledge that pursue the connections between privileged access and the transparency of mind and a series of articles that explore the relation between intentionality and phenomenology in memory.

ADDRESS: Philosophy Department, Macquarie University, Sydney NSW 2109, Australia. E-mail: jfernandez@scmp.mq.edu.au