

The Structure of Time in Autobiographical Memory

John Campbell

1. Time-Identification in Memory

Much of ordinary memory is autobiographical; memory of what one saw and did, where and when. It may derive from your own past experiences, or from what other people told you about your past life. It may be phenomenologically rich, redolent of that autumn afternoon so long ago, or a few austere reports of what happened. But all autobiographical memory is first-person memory, stateable using 'I'. It is a memory you would express by saying, 'I remember I . . . '.

My question is whether this autobiographical memory requires the conception of time as linear. To think of time as linear is to suppose that times are ordered by 'earlier than' in a way that is irreflexive, anti-symmetric, transitive and connected. If we identify times using the ordinary Western clock and calendar we do seem to be supposing that time is linear, and we are using a method of identifying times that could be said to be 'canonical' in the sense that the temporal relations between any two times can simply be read off from the ways in which they are identified: it is immediately apparent that, for example, Thursday 25th December 1963 is before Friday 26th December 1963. But autobiographical memory could hardly be thought to depend upon the use of the clock and calendar. Presumably humans had autobiographical memory even before the invention of the clock and calendar, children now still develop autobiographical memory before they learn how to use the clock and calendar, and people in contemporary societies which do not use the clock and calendar at all may still have autobiographical memory. So if autobiographical memory involves the identification of past times, it must use a way of identifying times that is more primitive than the clock and calendar. My question then is whether in using this more primitive way of identifying past times we are taking it that time is linear. Other structures that might be assigned to the time of autobiographical memory are, for example, those involving a tree structure, or disjoint though individually linear time-streams.

Psychologists who work on human memory are often hostile to the idea that autobiographical memory involves a conception of time as linear (Friedman 1990). One reason for this is that they want to resist the idea of an 'intrinsic temporal code', a kind of internal, innate analogue of the clock and calendar system, which would be canonical in the sense that the temporal relations between any two times identified using the code could simply be read off from the ways in which they were identified. There is no reason to suppose that there is such a

code, and the idea that there is runs into trouble over the banal fact that you may remember quite a lot about the temporal locations of two past events without being able to locate them temporally with respect to one another.

Important though it is, this point does not of itself settle the question of whether autobiographical memory involves the conception of time as linear. It simply leaves us with the problem of saying something positive about how past times are identified in memory. What primitive methods of time-identification do we have?

One possibility is that the time of a remembered event is specified in terms of just how long ago the event happened. That does seem to be how it works for some memories, for example memories of what happened yesterday. But in general we do not keep track of just how long ago remembered events occurred. There are two points here. In general we do not keep track of how long ago past events occurred with enough precision to let us locate them temporally with respect to one another, and anyhow the present does not seem to play a distinguished role in the way we identify times in memory (Friedman 1991). Nonetheless, as we shall see, the idea that these questions about just how long ago something happened do have determinate answers, whether we know them or not, does have a role to play in our having the conception we do of time in memory.

A better answer is that we identify remembered times by their relations to remembered events. So we might identify a past day as 'the day on which Thatcher resigned', or 'the day after Thatcher resigned'. We can use this to locate other events by their time-relations to Thatcher's resignation. But if that is the right way to think of the identifications of past times in memory, then the next question is what it comes to, that I remember Thatcher's resignation itself as having a temporal location. It is not just that Thatcher's resignation is itself temporally related to the other events I remember; the whole network of roughly contemporaneous events is itself temporally located. But what does this come to? If I remember Thatcher resigning, I seem to add nothing to that if I say that I remember it as having happened on the day on which Thatcher resigned.

The same question is raised by the idea that in memory past times may be identified using demonstratives such as 'that time', where the demonstrative has its reference fixed as the time from which most of the current memory causally derives. Again, if I say that I remember Thatcher's resignation, what do I add if I say I remember it as having happened on 'that day' – the day from which my information about the resignation causally derives?

The reason something is added by these time identifications is precisely that we do operate with some conception of the structure of time, so that we do think of the times we identify as standing in temporal relations to each other. What is added is what is given by the conception of time as linear, branching or composed of disjoint streams. So in thinking about how time is represented in memory we can hardly avoid the question what structure time is taken to have.

The temporal relations in which remembered events stand to one another have implications for the possibility of causal relations between those events. If one

event came after another, it cannot have been the cause of the other event. Also, if there is a causal connection between two events, they must be temporally related; the cause must have come before the effect. These principles connect temporal relations to the possibility of causal relations between different events. These principles need not be explicitly grasped by someone who has autobiographical memory. Rather, they relate to the principles governing the construction of narratives in autobiographical memory, principles of plot construction which the subject may observe without ever having explicitly formulated them.

Theorists of memory often stress its fragmented character. Memories are interrupted. There are many gaps in the record, most strikingly the amnesia of early childhood, but again and again in later life the memories from this time or that simply peter out. And memory is plural, with many narratives, apparently disjoint, making up the total. And it is not just that there are many narratives: there seems to be some plasticity in the narratives themselves, in that how we will tell the story of past events depends to some extent on who we are telling it to. These aspects of memory narratives, together with the fact that they are structured as narratives, with evident background influences from the local culture in the forms that the narratives take, mean that it is fruitful to think of memory narratives as literary texts, and to try to illuminate them in the way one tries to illuminate a literary text (Albright 1994). But if we have this literary approach to memory, we may suppose that there is no more to the facts than is reported in the narratives, just as there is little more to the truth about a set of fictional characters than the text reveals to us. The various events recounted in a particular reminiscence may all be temporally related to one another, within the story. But the memory narratives themselves may provide no ground for relating the time of those events to other times which are remembered. So the proponent of a literary view of memory concludes that time in memory is very far from being linear: the various times remembered may be temporally related within individual memory narratives, but there is no supposition that all times are temporally related to one another. That supposition is imposed by our later grasp of the clock and calendar.

2. The Unity of the Remembered Self

All autobiographical memories have a first-person character, they are memories of what I saw and did. So, we tend to suppose, they are inferentially integrated. If I remember that I was on the beach and I remember that I saw a volcano, then it was one and the same person who was on the beach as saw a volcano. And the memories are inferentially integrated with my first-person knowledge of my current circumstances and current plans. If I once was lost but now I am found, it is one and the same person who once was lost but now is found. This datum, the logical connectedness of all the first-person knowledge, I will call 'the transparent unity of the self in memory'. In this section I will argue that the subject has to recognise the possibility of an empirical grounding for that transparent unity,

and that this grounding is provided only by the conception of the self as spatio-temporally continuous, which requires the conception of time as linear.

This transparent unity brings with it a demand that the subject recognise the need for it to be possible to find an empirical grounding for the identity that underpins all these logical connections. As Williams (1973) observed, the idea that you can have inferential integration of first-person statements without any empirical grounding is responsible for mistakes about the role of the first person in imagination. There is such a thing as imagining being Anthony Trollope, you can imagine sitting at the kitchen table on a cold morning, gripping the pencil and grimly turning out hundreds of words, making remarks like 'I shall have to be off to work in a moment'. Within the context of this imagination, you will use the first person to report what is being thought and done. And that use of the first person refers not to you, but to Trollope. You make a mistake if you describe this as 'imagining that I am Trollope', and suppose this to show that you could have been Trollope: that you could have had different parents, lived in a different century, and so on. The use of 'I' in the imaginative project, and the use of 'I' to refer to oneself, are not inferentially integrated.

The literary view of memory has implications for the unity of the remembered self. On the literary view of memory, there is no more to the facts than is reported in the narratives. And it is a fallacy to suppose that the mere use of 'I' in a number of narratives gives one the right to regard them as inferentially integrated. Inferential integration needs some empirical grounding, a point which is vivid when we think about people who seem to remember their past lives as emperors or as racehorses. Given the interrupted character of memory, there seems to be no grounding in the narratives for the assumption that it is the same person who is remembered as who now remembers. And given the plural character of memory, it seems to be a fallacy to suppose that it is the same person who figures in all the different memory narratives one has.

One way to get at the point is to consider a parallel between memory and knowledge of other minds. You may read articles by a right-wing columnist who has what strikes you as a curious world-view that you slowly get the hang of, by simulating his state of mind, throwing yourself into his shoes and finding how the world looks from there. Within the scope of the imaginative project, you may use 'I' in reporting his state of mind, and that use of 'I' will refer to him rather than to you (Gordon 1995, Williams 1973). You might also do this with your next-door neighbour, a tempestuous character whose reaction to a Jehovah's Witness you can readily imagine, and again, you may use 'I' to refer to him rather than to you, when imagining his reaction. Now it may be that unknown to you, the right-wing columnist is in fact your neighbour, but that does not of itself give you the right to regard the uses of 'I' across the imaginative projects as inferentially integrated; the assumption of sameness of person here would need some empirical grounding. Now the knowledge that you have in memory of your own past can be thought of as involving simulation of your past self; memory provides a kind of imaginative understanding of how it was for you then. If I remember my life as a teenage hooligan, and I remember my life as an enthusiastic gardener, I may

be simulating the same person twice over, but so far there may be nothing in the contents of the memories themselves to sustain that identity. And in that case, I have no right to regard the uses of 'I' in reporting the two sets of memories as inferentially integrated. If this line of thought is correct, then the literary view of the disunity of memory is vindicated: memory narratives are not in fact all inferentially integrated through the fact that they are all first-personal.

The basic problem with this is that the use of 'I' within a simulation depends on a more fundamental use of 'I' outside the simulation, the strict and literal use of 'I', and the involvement of that literal use of 'I' in memory does guarantee inferential integration. Simulation of a particular individual always needs some identification of the person being simulated, even if it is only as 'that man before me', so that there is somewhere for the simulative project to start. In the case of simulating another person, there are endlessly many different kinds of third-person identification to be had, such as 'the author of this article', or 'my next-door neighbour'. But in autobiographical memory the only way I identify the target of the so-called simulation is by the use of the first person: the target of the simulation is not identified by some description, for example, but simply as me. This use of the first person is a use outside, not inside the simulation, and its effect is to collapse together the uses of 'I' outside and inside the simulation; in effect it is questionable whether we have here a simulation properly so-called at all. There may be some place for simulation in memory, as when I relive the events of last year, once again going through the emotions of that momentous day. But this is not the ordinary case, as when I simply recount to you what happened, and the use of the first person in memory seems to have its most basic place in that recounting of what happened, which does not involve simulation at all. It is only against the background of that capacity to use the first person in simple memory narratives that we can have simulation in memory. The use of 'I' within a simulation depends upon the plain use of 'I' outside the simulation; and there is no way of resisting the inferential integration of those uses of 'I' in plain memory narrative, with each other and with present-tense first-person reports. But the challenge raised by the literary view is still there, namely, to identify what kind of empirical grounding, if any, memory must provide for this inferential integration, given its interrupted and plural character.

We will not find this empirical grounding in the contents of the memories themselves; we should look rather at the rules to be followed in constructing memory narratives. For anyone there are many principles of plot construction to be followed in elaborating reminiscences. There are literary aspects to plot construction, as when one uses narrative forms that are pervasive in the culture. Or one may tend to cast oneself as the hero, the heroine or the victim of the action. But there are also principles of plot construction which are essential to the claim memory has to be a form of knowledge. For example, we take it as axiomatic that no-one can be in two places at the same time; I might say, 'I seem to remember her being there, but that can't be right because she had left the country by then'. Or we have the point that to remember something on the basis of my own observation of it, I must have been there at the time. So I can say, 'I seem to remember

the resignation being announced, but that can't be right because I wasn't here when it was made'. One of the most basic principles of plot construction is that the remembered I traces a continuous spatio-temporal route through all the narratives of memory, a route continuous with the present and future location of the remembering subject. It may be that there are interruptions in what one remembers, but these can never be taken to imply violation of that principle, and, of course, the mere existence of gaps in memory does not imply violation of the principle. This principle imposes a kind of unity on all the narratives; there has to be a coherent story to be told about my movements which will fit with the contents of all my various memories.

It may be true also that we can make some sense of the idea of a subject who does not have any grasp of that principle, who simply does not use it in constructing memory narratives. But, on the present approach, that subject cannot be said to have grasped the use of 'I' as a referring term, because he has no grasp of the empirical grounding needed for the inferential integration of uses of 'I' across narratives. Such a subject may construct narratives, and they may be more or less related to what actually occurred, but he will not have autobiographical memory.

If this is correct, then the remembering subject cannot coherently say that he remembers what he saw and did at a time temporally unrelated to the present. So we can abandon the idea of discontinuous epochs being remembered as discontinuous. For the same reason, we have to abandon the idea of remembered times being represented as branching back into the past. The demand for spatio-temporal continuity of the self means that the remembered times must all be temporally related to the present, but it also means that the times one remembers can all be linearly ordered, as having been experienced earlier or later by one and the same self (Quinton 1993).

The really difficult question is what we say about the time-structure of memory if, in constructing memory narratives, we take a functionalist, in contrast to a substantialist, view of the identity of the self. For a functionalist, the identity of the self is simply constituted by the existence of causal relations between the ways the self is at various times. For a substantialist, in contrast, the identity of the self at one time with the self at another time is the ground for there being causal relations between the self at one time and the self at another; it is what explains the existence of these causal relations. A functionalist view may, but need not take the form of a reductionist view of the self. For the reductionist, individual mental states, such as particular memories, are taken to be more fundamental objects than the self; the self is a kind of construction out of those more basic things and the causal connections between them (Hume 1888, Parfit 1984). But a functionalist view need not take this form. The functionalist may think of psychological states as properties of objects, and the object will be constituted by the causal potentialities it has by having those properties (Shoemaker 1984). So long as a change in the properties at one time internally causally affects what properties are had later, we have one and the same person.

The key point is that whether or not the functionalism takes a reductionist form, there may be no more to the causal relations constituting a person than the

causal relations which are described and constituted by the memory narratives of that person. There are two types of causal relation here. There are those which the person remembers as holding among the various things and events remembered. And there are those which are constituted by the memory narratives themselves. The point is that so far as the subject is concerned, the causal relations – even the potential causal relations – among the things and events remembered, including the subject, may peter out there. There is no ground for supposing there to be any potential for causal relations beyond those actually detailed, or constituted by the acts of remembering themselves. On the substantialist view, the ground for supposing there to be such potential causal links would be supplied by the conception of the remembering self as a persisting substance, with a linear temporal organisation, so that all the remembered events could be temporally ordered with respect to it, and derivatively with respect to each other. But this substantial ground is precisely what the functionalist view rejects. To that extent, the literary view of memory will be vindicated by functionalism. The subject will not have any ground for presuming there to be potential or actual causal relations other than those which are explicitly part of the memory narratives. Correlatively, on this view, the subject has no ground for thinking that in principle, there must be discernible temporal relations among all the events remembered.

It may be said that if the functionalist takes the memories themselves to be part of the causal connectedness that constitutes a person, then memory of times disjoint from the present can be ruled out, since all remembered events must be thought of as causing, and hence preceding, the memories of them. The functionalist is not, however, in any position to rule out the possibility that the subject may assign a branching structure to the times remembered, a structure rooted in the present and branching backwards, so that all remembered times are earlier than the present, 'earlier than' is recognised to be irreflexive and transitive, but not all remembered times are temporally related to one another: there are remembered times, assigned to events in different memory narratives, which bear no temporal relations to one another.

To sum up: if, as our principle of plot construction governing the use of 'I', we use a functionalist conception of the self, then it would be quite consistent with that, in fact the natural procedure would be, to assign a backwards branching structure to the times remembered. The various narratives of memory will each have their own, individually linear temporal structure, and each will be temporally related to the current remembering subject, but the subject need not suppose that there is, even in principle, any answer to the question what temporal relations there are between the events remembered across different narratives. In contrast, if our principle of plot construction for the use of 'I' is provided by the conception of the self as a spatio-temporally continuous substance, the subject has the basis for the supposition that all the times remembered are temporally related not just to the present but to each other, because each is temporally related to the past self. So if the principle of plot construction is given by the conception of the self as a substance, the subject is bound to assign a linear structure to the times remembered.

I said that a functionalist principle of plot construction of 'I' will assign a backwards branching structure to the times remembered. But further reflection suggests that this undermines the idea that we have here something that could serve as a principle of plot construction for the first person at all. If we have a remembered self in two disjoint narratives, what makes it the case that we have one and the same person here? On the functionalist view, all that makes it the case that it is the same person across narratives is that the doings of both are being remembered by the same current rememberer. The doings of the two remembered selves are temporally disjoint: neither is temporally related to the other. But this temporally delayed grounding for the identity between the remembered selves does not seem to be recognisable as a way of grounding an identity at all. If the identity between the two holds, it should hold whatever comes later, whether or not there ever is a remembering self to reflect later on what happened then. But that is precisely what we do not have in this case. So it does not seem that the functionalist gives a way of grounding the use of the first person as a referring term at all. For the functionalist, the transparent unity of the self in memory, the datum from which I began, really would seem to be a sham. My own conclusion from this would be that we should reject functionalism, and move to the conception of the self as a persisting substance.

Suppose we agree that the autobiographical character of the memories demands that the time of the self be thought of as linearly organised. The point now is that the self is not some magical entity that can somehow confer linearity on the time of memory. The problem is to understand how it can be that we have a conception of time within which we can make sense of the idea of the self as spatio-temporally continuous. What makes the problem sharp is the force of the point urged by the proponent of a literary view of memory: that memories are characteristically interrupted and plural. How can we have come by the conception of time as linear?

3. Rhythms and Cycles

When thinking about how we represent time, it is an inescapable thought that one's apprehension of natural periods, such as the light-dark sequence, or the passage of the seasons, must have a fundamental role to play. But in pinpointing the role, we have to make a distinction, which I will mark by the differentiation between 'rhythm' and 'cycle'. This is a difference between ways of representing time. A 'rhythm', as I will use the term, is a way that a creature has of keeping synchronized with the natural periods in its environment, so that some of its activities can be timed to fit with the times of events around it. 'Biological clocks' keep creatures synchronized with the environment in this way (Gallistel 1990, Winfree 1987). For example, most animals can regulate their activity by the light-dark cycle, either so that they are resting during the day when their predators are around, or so that they are active only at times when vision can be put to use. Animals need to synchronize their feeding times with the times at which food is available, and they

need to synchronize their mating periods with the seasons of the year, so that their young are born at favourable times. A biological clock can be dedicated or general-purpose. That is, if we consider the clocks governing the life of a plant, each may be specific to a particular task, such as getting it to flower at the right time of year. The clock registering the light-dark cycle has, for many animals including humans, the purpose of governing the onset and offset of activity, the offset and onset of sleep. But for humans at any rate, the clock seems to be general-purpose, in that the phase of the clock can govern endlessly many different activities. We can learn at which phase of the day to expect this or that new phenomenon, such as the presence of food at a particular location, or the presence of students expecting to be taught. This is one reason for drawing a distinction between the purely biological time of a plant, and the cognitive time of a creature with a general-purpose clock.

In contrast, a 'cycle', as I will use the term, is a spatial or quasi-spatial image of a recurrent period (Friedman 1986). So, for example, a timetable for a twenty-four hour period, stating what happens during morning, in the afternoon, late afternoon, evening, and so on, represents the period. It of itself is not enough to enable one to keep in time with the period as it passes. It simply gives one a representation at a time, of what happens during the whole period. We could put the contrast by saying that cycles are synchronic representations, available to the creature in their entirety at a single moment, which represent a wide span of time. Rhythms, in contrast, are diachronic, and are not available to the creature at a single time, but can be put to work only over a period, to keep the creature in time with the various phases of the period.

There is no reason why the cycles in terms of which one thinks of time should be confined to representations of the patterns which one uses in one's rhythms, to keep synchronized with one's surroundings. For example, one could easily think in terms of an artificial 'week' of nine days, and get used to thinking in terms of such a week, without this having anything to do with any supposed underlying biological clock functioning on a nine-day cycle. But for humans the diurnal cycle is basic, and the other cycles in terms of which we think tend to be defined in terms of it, either as multiples, such as the seven-day week, or as fractions, such as the hour. So would it be right to say that the cycles in terms of which one thinks must initially be defined in terms of cycles which recapitulate the basic rhythms which keep one in step with one's surroundings? I think the most we can say is that there will be obvious reason to highlight a cycle which recapitulates the basic rhythms which keep one in step with one's surroundings. The kind of time-periods in which one is likely to be interested, and which will have significance for one, will very likely be the very periods for which evolution has provided one with a timer. So the patterns of events in one's surroundings to which one attaches significance, and the history of one's own interactions with one's surroundings, and indeed with other people, will be most naturally describable using cycles which recapitulate the rhythms on which one relies in practice.

More generally, we have to give causal significance to the assignation of times to events. One primitive way of doing this is through direct connections with action; alternatively, we can do it at a reflective level, at which we have a narrative

that explicitly details the causal relations between events. In the use of a biological clock we have a way of giving causal significance to the times of events directly through action: if I know where the food is at noon, I know how to act at noon. When I reflect on my actions and construct the narrative of what I have done, the assignation of times to remembered events by using the same set of cycles gives a way of connecting the assignation of causal significance at the primitive level to the assignation of causal significance at the level of reflective narrative.

Once we have recapitulated the rhythm of the clock, to form the conception of the cycle involved, we have the basis for the ability to arithmetize time, by counting the cycles backwards and forwards. And that gives one a way in principle of identifying the time of any event, a single set of cycles, counted backwards and forwards from the present, in terms of which one can assign a time to any event whatever. If the clock which is being recapitulated really is general-purpose, and used in the governance of all one's activities, then at the reflective level at which we think in terms of cycles we will be able in principle to log in the times of all remembered activities. And this, it seems to me, is one foundation of the conception of time as linear. Having this conception means that all the memory narratives one constructs are using a single set of cycles.

This picture has to be complicated by the fact that we give events temporal locations by reference to quite different scales, quite different cycles. I may identify the times of the one set of events by reference to the seasons of the year at which they all occurred, and the times of another set of events by reference to the times of day at which they all occurred. What is required for unity here is that one can count any one type of cycle back and forth endlessly, and think in terms of there being determinate ratios between the lengths of the different types of cycle one encounters.

You might object that our grasp of the possibility of arithmetizing cycles has to do only with our ways of measuring time, whereas in discussing whether time is linear, and in particular whether it is connected, we are looking at topology. This seems to be a question at a more fundamental level than anything that could be addressed by appealing to some point about the nature of the cycles that we use to measure time. But I think that conception of what is fundamental has to be challenged. It is true that for a theorist looking at a system of temporal reference it is possible to separate out the topological component from the metrical. But it by no means follows that the purely topological is somehow conceptually more fundamental than the metrical. In fact, the ability to think of time in memory in purely topological terms may be something that can be accomplished only by the theorist. There are not two separable stages in our most primitive understanding of time, one at which we impose the topology and a later stage at which we impose the metric; rather, we impose the topology by imposing the metric.

4. Causal and Temporal Structure

This does not end the argument, though, because we can still ask in what way our construction of memory narratives actually puts to work this linear character of

time in memory. Is it really used in autobiographical memory? It helps here to think of the position of, say, a Newtonian physicist who takes time to be linear. This physicist is not claiming to be able to construct a finely detailed, unified history of the world since the origin of matter. It may well be that the only histories that this physicist can in fact construct are partial, fragmentary, and difficult to integrate with one another; they say only what was happening in one or another region at some time or other. Nonetheless, what the physicist claims in saying that time is linear is that all these regional histories could in principle be integrated with one another.

One way in which the assumption of general causal connectedness is implicit in ordinary memory is through the fact that we think in terms of a single range of concrete things which form the background to all our activities, past and present. While some of these objects may not in fact act on one another, there is in general no dividing the whole totality into subgroups which are causally isolated from each other. We think of concrete objects as causally structured. One way to put this is in terms of the truth of a set of conditionals about how interactions at one time would affect the later condition of the object, the potentiality of the object for propagating causal influence over time. You make a mark on the thing at one time, and there the mark is later (Salmon 1984). More generally, if you were to mark this object now, it would bear that mark later. Marks made at one time are propagated to later times.

There is another aspect to the causal structure of concrete objects, which is that a concrete thing is capable of figuring in many interactions, and so being the cause of a range of correlated effects. For example, one person may show up in many ways in your life, a distinctive figure who haunts your various narratives, casting a distinctive shadow for better or worse across all of them.

Both of these aspects of causal structure are temporally asymmetric. The capacity of an object to propagate causal influence within itself over time runs from earlier to later. And the common cause of a range of correlated effects must be earlier than those effects. And since the same range of concrete things figures in all of one's memory narratives, causal influence can be propagated from the events of one narrative to another. The fact that it can be one and the same object that figures in interactions in different memory narratives, causing correlated effects in those different narratives, plays a role in putting to work the temporal connectedness of the events we remember. And again, an object entering into the events of one narrative may bear upon it the marks of its earlier interactions in other narratives.

These points apply also to the remembered self, which is one concrete object among many. The striking causal connection in a person is the potentiality for memory of past perceptions and actions. Memories of the events of one memory narrative are carried by the person into the events of another memory narrative. And the marks of one interaction can affect how the person behaves in many subsequent interactions, in many types of narrative stream. So the unity of the subject in many narratives itself puts to work the linearity of remembered time.

The only way in which this pervasive possibility of causal connection across

narrative streams might be blocked is if you thought that the properties of the various objects you encounter could somehow be partitioned, so that the events of one narrative could affect the object's possession only of one set of properties, and be affected only by one set of behaviours of the object, whereas the events of another narrative could affect the object's possession only of a quite different set of properties, and be affected only by a different set of behaviours of the thing. Let me set out this possibility a bit more fully. In general, the properties of a thing do not affect its behaviour on a one-by-one basis. The way in which the object tends to behave is the joint upshot of its possession of a family of properties (Shoemaker 1984). So, for example, the fact that a spade can be used for digging is owed to its being rigid, of the right shape and size, having a flat blade, and so on. What we will find, as we consider which properties are the bases of the various causal powers that the object has, is that we are appealing to an overlapping family of properties that the thing has, the same background pool of properties being drawn upon again and again. The hypothesis we are considering is that there might be a partitioning in the causal powers and underlying properties of an object, so that on one side of the partition we have one set of causal powers and one pool of underlying properties, and on the other side of the partition we have a disjoint set of causal powers and a disjoint pool of underlying properties. So one set of causal powers, and one pool of underlying properties might relate to the way in which the table figures in one narrative, and another, disjoint set of causal powers and another pool of underlying properties might relate to the way in which the table figures in the other narrative. So one and the same object could figure in the two narratives without there being any possibility of causal connectedness between the events of the two narratives. The events in the one narrative could affect only those properties of the object which are causally idle from the standpoint of the other narrative.

This is a very radical proposal, and obviously we do not ordinarily find this kind of partitioning in the causal powers and properties of objects. And on further reflection, it is not just that it does not often happen in ordinary life that we have this kind of partitioning of the causal powers and properties of the object. Rather, the suspicion that we have this kind of partitioning threatens the very idea that we have just one object here. It suggests that we have two objects, one having one family of properties and the other having the other family of properties. To put it more generally, we tend to think that a collection of properties comprises properties of a single object when the causal powers of that object are the consequence of the possession by the thing of various overlapping families of properties drawn from that underlying collection. Properties which could be deleted from the collection without loss to the account of which causal powers the object has are not, strictly speaking, properties of that object: they are 'Cambridge properties' – that is, they are properties such that we would not count a change in whether the object has the property to be a change in the object itself – or properties of some other object. If this is right, then the various narratives constituting the autobiographical memory of a single individual will be thought of as organized around a single linear time so long as there is some overlap in the persisting things which figure in the various narratives. These objects

will ensure the temporal connectedness of all the times remembered by ensuring the potential causal connectedness of all the events remembered. You might say that this could in principle be achieved, if we think of a subject whose various narrative streams simply involve different concrete things. But, of course, the self is another concrete thing, and the unity of the self in all these narratives will ensure the possibility of causal connectedness across the narratives.

5. Other Cultures

It is sometimes said that there are other cultures which do not think of remembered time as linear (Friedman 1990). The data here are often difficult to interpret. But the fundamental question that has always to be addressed is: suppose we do find that we have a people who do not think in terms of time as linear, do we then have a people who genuinely have autobiographical memory? Even if what they have counts as a form of memory, in the sense that it is a kind of non-inferential knowledge of the past, it is hard to see how it could count as autobiographical, given that the remembered self is not being conceived as spatio-temporally continuous.

If these people have a form of words that they use to make reports about the past, but they do not have the conception of the self as spatio-temporally continuous, then they are in no position to insist on a requirement that we often impose on memory: namely, that to remember something – an observed scene – you had to be there at the time. Of course, we accept that it may be that your knowledge of the scene does not derive from your own observation of it, but from what you were told about the event by other people. But still we demand that those other people had to be there at the time, even if you were not; at some point, the memory had to derive from past observation. When we apply the demand, we use the idea of spatio-temporal continuity to determine which interactions between people could have occurred. But we cannot impose this demand if we do not have the conception of time as linear, and so do not have the conception of the self as spatio-temporally continuous.

It may be that what we have here is a way of knowing about the past, but that it is not memory: it rather involves the use of non-memory-based evidence, such as the remains of a fire or the dead body of a man. On that basis our people could make statements about the past without having had to be there. But there is another possibility, which is that these past-tense statements do not involve the use of non-memory-based evidence, but do not have any requirement of the sort 'you had to be there'. There may nonetheless be such a thing as getting it right or wrong in the use of such statements. In that case, what we are dealing with may not be a form of memory at all: we may be dealing with a people who have no distinction between memory and myth. For such a people, the literary view of memory would seem to be correct.¹

John Campbell
New College
Oxford, England

NOTES

¹ Earlier versions were presented to UCB Philosophy Dept., the Meno group at UCB Psychology, the Society for Philosophy and Psychology, seminars and discussion groups in Oxford, and a round table on the first person organized by Charles Travis at Stirling University. Thanks also to Philippa Foot and an anonymous referee for the *EJP*. This paper was prepared during tenure of a British Academy Research Readership, and I am grateful to the Academy for its support.

REFERENCES

- Albright, Daniel (1994), 'Literary and Psychological Models of the Self', in Ulric Neisser and Robyn Fivush, eds., *The Remembering Self*. Cambridge: Cambridge University Press.
- Friedman, William J. (1986), 'The Development of Children's Knowledge of Temporal Structure', in *Child Development* 57, 1386–1400.
- Friedman, William J. (1990), *About Time*. Cambridge: MIT Press.
- Friedman, William J. (1991), 'The Development of Children's Memory for the Time of Past Events', in *Child Development* 62, 139–155.
- Gallistel, Charles R. (1990), *The Organisation of Learning*. Cambridge: MIT Press.
- Gordon, Robert (1995), 'Simulation Without Introspection or Inference from Me to You', in Martin Davies and Tony Stone, eds., *Mental Simulation*. Oxford: Blackwell.
- Hume, David (1888), *A Treatise of Human Nature*, ed. by L. A. Selby-Bigge. Oxford: Oxford University Press.
- Parfit, Derek (1984), *Reasons and Persons*. Oxford: Oxford University Press.
- Quinton, Anthony (1993), 'Spaces and Times', in Robin Le Poidevin and Murray MacBeath, eds., *The Philosophy of Time*. Oxford: Oxford University Press.
- Salmon, Wesley C. (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Shoemaker, Sydney (1984), 'Identity, Properties and Causality', in his *Identity, Cause and Mind*. Cambridge: Cambridge University Press.
- Williams, Bernard (1973), 'Imagination and the Self', in his *Problems of the Self*. Cambridge: Cambridge University Press.
- Winfrey, Arthur T. (1987), *The Timing of Biological Clocks*. New York: Scientific American Library.