

Provided by the author(s) and University College Dublin Library in accordance with publisher policies. Please cite the published version when available.

Title	The life of a process
Author(s)	Stout, Rowland
Publication Date	2003
Publication information	G. Debrock (ed.). Pragmatic Process : Essays on a Quiet Philosophical Revolution
Publisher	Rodopi
Link to publisher's version	http://www.ucd.ie/philosophy/staff/stout/the%20life%20of%20a%20process.pdf ; http://www.rodopi.nl/senj.asp?BookId=VIBS+137
This item's record/more information	http://hdl.handle.net/10197/4961

Downloaded 2016-02-29T16:45:42Z

Some rights reserved. For more information, please see the item record link above.



Penultimate draft of Rowland Stout, "THE LIFE OF A PROCESS", in Guy Debrock (ed.), *Pragmatic Process Philosophy*, Rodopi, 2003.

1. Introduction

When a child is growing, something is happening - the process of that child growing. We can identify and talk about this process. It started before birth and will last until the child reaches full size - one and the same process throughout. When that child has stopped growing, we can retrospectively identify the process as something that was happening. But we can also retrospectively identify something else. We can identify what happened - the event of that child growing. The event is not an on-going thing, nor was it ever. By its very nature it is a gone-over sort of thing.

The aim of this paper is to consider the relationship between such things. And I will approach this by considering the relationship between our ways of describing such things. So, although I am interested in developing the metaphysics of processes and events, I approach this by considering how we talk about processes and events. My conclusion will be that our talk of gone-over events depends on our talk of going-on processes.

This conclusion may suggest the metaphysical conclusion that processes are somehow ontologically more basic than events. But it will emerge a little paradoxically that because of the radical difference between the way we talk about processes and the way we talk about events, it is not easy to argue that there is a metaphysical distinction between these things - not at any rate at what people call the level of deep ontology. I shall be satisfied with a conclusion that makes no claim to be part of deep ontology. The conclusion is that it is not possible to talk of a world containing events and no processes.

I think that this has fundamental importance for understanding our talk of change and causation. The standard approach here is to try to work out what change and causation amount to in a world of events. This approach has faced familiar problems which stem I think from this assumption that the only kinds of happenings we can describe are gone-over happenings - events in my sense. How are we to make sense of the world going on just in terms of a world of gone-over happenings? Somehow the dynamic aspect of change - the changing - cannot be accommodated by a description of all the changes that have or will have happened.

Consider Whitehead's version of Zeno's arrow paradox.¹ Suppose happenings, like the motion of the arrow, are all temporally extended and infinitely divisible. Consider a motion that includes time t . It can be divided into that part that happens up to and including time t and that part that happens after t . But there is nothing that happens - no further happening or motion - that effects the transition from the first part to the second. There is no motion that carries the arrow across time t . This argument can be applied to any time; and so there is no motion that carries the arrow across any time. Hence there is at no time motion of the arrow.

This argument assumes that we require a causal notion of motion. It assumes that motion should not be understood in Russell's way as the arrow being in one state and then in another and in the meantime being in all the intervening states.² The arrow's motion is what gets it through this continuous series of states - it effects the transition. Temporally extended happenings by themselves cannot effect transitions. They are

transitions. They cannot take us from one side to the other, because with a temporally extended happening, there is not one and the same thing present on both sides of the transition. There are two distinct parts, one on either side. What we need are things which, by persisting themselves through the time of the transition, can carry the transition.

By having a category of dynamic continuants - persisting happenings or processes in my sense - we can satisfy this requirement. The very motion that is happening at time t was happening before and will happen after. What carries the arrow through t is simply this persisting motion. It cannot be divided temporally because it is not temporally extended. So we do not have the extra problem of what effects the transition from one part to the next. Of course there are plenty of other ways that one might respond to Zeno's paradox, but this way seems to provide the most natural solution.

2. People and their Lives

I will begin by introducing the crucial terms of my argument by making a different but analogous distinction – that between a person and their life. The point of this is that the distinction should be fairly obvious and so provide some unproblematic stage setting for the distinction I want to elaborate between processes and their events.

A person's life has temporal extension and temporal parts. A life is extended in time - its extent is some interval of time. My childhood is part of my life. It can be located in a smaller interval of time. Depending on its temporal extent, we can describe a life as long or short. We can also describe a life as eventful or uneventful. This suggests that events make up parts of lives. The events of my birth, marriage and possibly even death are all parts of my life, as is the event of my writing this paper. My whole life might be described as one event.

None of these things that can be said of my life can ordinarily be said of me. I am not in a temporal interval, my childhood and my marriage are not parts of me, except in the sense of contributing to making me what I am. And the sense in which I am long or short is quite different from the sense in which my life is. My length is spatial; my life's length is temporal.

When we say that lives have properties, we are attributing these properties timelessly. The "have" is not the "have" of the present moment; it is tenseless. Lives and parts of lives may of course have temporal properties - properties of temporal location and extension. But these temporal properties are had timelessly. For example, my childhood has timelessly the property of taking place in the 1960s and 70s. It is as true of my childhood in 2001 as it is of my childhood in 1969 that the period of its taking place is the 60s and 70s.

There are exceptions to the idea that properties are always attributed to lives timelessly. For example, I might say that my life is changing irrevocably now that I have had a baby. My life was free and easy; now it is not. One day my life will be nearly over. These are all tensed attributions. The problem here is one of ambiguity. When I am dead, my life in one sense will be non-existent. But in another sense it will have various properties incompatible with not existing, like that of spanning several decades. I

am interested here in a life in the second of these two senses - as a whole event rather than as a process of living. As a whole event, my life has properties timelessly.

This ambiguity is an instance of the wider difficulty that the same form of words often picks out both a process and an event. I am thinking of lives here as events, but we can also think of a life as the on-going process of living a life. While this makes it more difficult for me to illustrate the distinction between processes and events, it also helps to explain why this distinction has proved so elusive to philosophers.

By contrast, when we say that I have properties, we are attributing these properties at a time. The “have” is the “have” of the present moment. This is what makes room for the idea of change. For example, suppose that I was a child in 1969 and an adult in 2001. We want to say that I have changed in this respect. We can make room for this by relativizing the instantiation relation and saying that I have-in-1969 the property of being a child and I have-in-2001 the property of being an adult. This is precisely Johnston’s solution to the supposed problem of identity through intrinsic change raised by Lewis developing Quine.³

Incompatible properties are attributed to one and the same thing. But Leibniz’s Law is not infringed since these properties are attributed at different times.

Given this we must understand Leibniz’s Law in a way which incorporates relativized instantiation relations as follows:

If $a=b$ then if a has at t some property F then b has at t F .

The fact that b has at t^* not- F does not contradict b having at t F ; it merely captures the idea of one and the same thing changing.

We cannot make room for this idea of change by saying that I-in-1969 is timelessly a child, and that I-in-2001 is timelessly an adult. For, I-in-1969 would not be the same subject as I-in-2001, and it would not be the case that one and the same thing was both child and adult. Nor can we make sense of it by saying that I timelessly have the property of being-a-child-in-1969, and I timelessly have the property of being-an-adult-in-2001. For, while at least in this case we have the same subject having two properties, there is no incompatibility between the two properties. The subject timelessly has them together, and does not change from having one to having the other. We might also try to relativize the whole proposition to a time and say that it-was-true-in-1969 that I be timelessly a child, and it-was-true-in-2001 that I be timelessly an adult. But this does not seem to have any advantages over just relativizing the instantiation relation.

So a person belongs to a different metaphysical category from their life. The sorts of properties that are attributed to people and to their lives are quite different, and they are instantiated in quite different ways - at a time and timelessly. People do not have temporal extension or temporal parts. It makes no sense to say that me yesterday is part of me; whereas it makes perfect sense to say that my life yesterday is part of my life. When you refer to a person's life at a certain time you are referring to a time-slice of that life; when you refer to a person at a time you are referring to the whole person - one and the same thing that exists at other times. People persist in time - they can last for years, despite constantly changing. They are continuants. People's lives are extended in time, and are not continuants.

One way to challenge this sort of distinction would be to question the necessity of quantifying over both lives and people in order to be able to say everything we want to say about the world. Given some principle of ontological parsimony like Quine's principle of ontological commitment in "On What There Is", if we do not need both metaphysical categories to make sense of our experience we should discard one.⁴ It is usually taken to be enduring entities - like people - rather than temporally extended entities - like people's lives - that come off worse given this principle.

One kind of argument in favour of the ontological priority of lives over people might be based on an attempt like Perry's to construct a reductive account of the identity of people in terms of the identities of lives and life stages.⁵ According to such an account, we identify and re-identify people by identifying other things - life stages - and seeing that the right sort of relation holds between these stages. If the identity conditions for continuants must be given in terms of the properties of non-continuants, then the non-continuants have some sort of conceptual priority, which can be translated into an ontological priority given a principle of ontological commitment.

But this sort of account assumes that we can identify a stage in a person's life in advance of identifying that person, and that does not seem very plausible. Also it does not seem well motivated. Identification of one sort of thing cannot always be explained in terms of identification of another sort of thing - on pain of infinite regress. Things in the most basic categories are not identified reductively using identity conditions, but must be identified more directly.

Given this, it is more natural to think that 3-dimensional continuants rather than 4-dimensional lives of continuants belong to the basic category as far as identification is concerned. Like Strawson and Evans we may reject the Humean tradition in which the fundamental skill that enables one to identify continuants is to see similarities and continuities between different objects.⁶ So we may argue against the idea that we identify objects by observing relationships between different object-stages. Instead the ability to keep track of one and the same object through space may be the basic skill which enables one to identify continuants. If the perceptual skill of tracking is what underlies our ability to identify people over time, then we do not have to be able to identify life-stages in advance of identifying people.

On the face of it, there is a conceptual priority of people over their lives rather than the other way around. Normally our entitlement to talk about people's lives depends on our entitlement to talk about people and not vice-versa. Normally we learn how to identify and re-identify people and how to attribute properties to them first and then develop the more abstract talk of their lives, parts of lives and properties of their lives later. Normally we would not know how to refer to a life or a life stage without having in mind the person whose life or life stage it was of.

There are some standard ways in which our talk of people's lives may be abstracted from our talk of people. For example, if, during 1990, it is true of a person, P, that they are alive, then it is (timelessly) true that there is a life stage of that person that spans 1990. A person's life is the sum of all their life stages. If it is true at time t that a person, P, is born and it is true at t+30 years that they die, then P's life lasts 30 years. If it is true during 1990 that a person lived in San Francisco then it is true (timelessly) that the 1990 stage of their life is in San Francisco, and so on. In general, I suspect that any claim

we want to make about a life or life-stage can be derived from some claims about the person whose life it is.

3. Processes and their Events

The distinction I am making between temporally enduring and temporally extended entities is often taken in the literature to correspond to the distinction between objects and processes.⁷ But just as there are two different logical categories, one for people and one for their lives, so, I shall argue, there are two different logical categories, one for processes and one for their lives. Happenings may belong to either of these categories depending on whether they are enduring or temporally extended. Enduring happenings are dynamic continuants - I shall call them processes. Temporally extended happenings I shall call events.⁸

We can get a handle on this distinction between the two kinds of happening by seeing how it is marked at the grammatical level by the aspect of the verb "to happen". There are those things which will have been happening and there are those things which will have happened. I am stipulating that processes are things that will have been happening and events are things that will have happened. When I am travelling to work, I want to say that what is happening is a process of my travelling to work. But when I have traveled to work, I want to say that what has happened is an event of my travelling to work. This captures the idea that processes are on-going things and events are gone-over things.

It may seem at first glance that there could be no genuine metaphysical distinction here, just a difference of tense. But, first of all, the difference between saying that something was happening and saying that something happened is not strictly speaking one of tense; it is a difference between the imperfective and the perfective aspect. And crucially, as I shall argue, this difference corresponds to the difference between identifying something without temporal parts that endures through time and identifying something with temporal parts that is extended in time. In this respect, it mirrors the difference between a person and their life.

Just as people can be identified and re-identified by being kept track of through time, so processes can be. Suppose that I am travelling to work now and that I was travelling to work three hours ago. Is what is happening now the very same thing as what was happening three hours ago? If what is being asked is about the type of thing that is happening, then the question is trivial. What is happening now is that I am travelling to work and what was happening three hours ago was that I was travelling to work. So, obviously what is happening now is the same type of thing as what was happening three hours ago.

But there is a non-trivial question here too. Is what is happening now the same token of that type of process as what was happening three hours ago? If I got to work, went straight home and went to work again, then it is clearly not the very same token process of my travelling to work. But if it simply takes a long time for me to get to work, what is happening now may indeed be the very same token process of travelling to work as what was happening three hours ago.

You can tell which of these is the case in the same way that you can tell whether there are two people present at different times or the same person continuously present between these times. The basic skill is that of keeping track of the thing for the duration. If you keep your eye on the process, you either will or will not observe a discontinuity where one process ends and another one begins later. As with the skill of keeping track of a person, this can be done through direct observation or indirectly by inference. If you telephone me at work between your two observations of my travelling to work you can establish that these were two token processes. In fact keeping track of processes is not always a simple task, since processes may be happening even when not very much else is going on. For example, I am still travelling to work even when I am sitting quite still in the station waiting for the train.

Treating processes as identifiable and re-identifiable particulars is what distinguishes my notion of processes from the currently most influential notion, which is mainly associated with Mourelatos.⁹ According to this latter notion, processes are distinguished from events by being stuffs, or at any rate stuff-like, whereas events are particulars. According to this idea, the process of my travelling to work is the stuff of the event that is my journey to work. You can have more or less travelling to work, but you cannot have more or less journey. “Journey” is a count noun phrase; “travelling to work” is a mass noun phrase.

Mourelatos has since acknowledged that using the word “process” to describe the stuff of events was unfortunate.¹⁰ “Process”, unlike “stuff” is not itself a mass noun, so we cannot say grammatically that process is the stuff of events. For this reason alone, “activity” seems a better word to describe what these philosophers are after.

Distinguishing the stuff of events from particular events is associated with taking seriously the distinction between perfect and imperfect nominals. “John’s singing the Marseillaise” is an imperfect nominal. “John’s singing of the Marseillaise” is a perfect nominal.¹¹ Bennett makes much of this distinction as a way of distinguishing between particular token events on the one hand and what he describes as event facts on the other.¹² The verb is “alive and kicking” in an imperfect nominal. This is shown by the fact that we can use adverbs to modify imperfect nominals. For example, we can talk of John’s loudly singing the Marseillaise. We cannot talk of John’s loudly singing of the Marseillaise, but only of John’s loud singing of the Marseillaise. Perfect nominals refer to identifiable particulars. Imperfect nominals refer to less particular things - activities.

My distinction between processes and events is not this distinction. I use perfect nominals to describe both processes and events in my sense by making sure I talk of the process of something happening or the event of something happening. Such things are particular tokens. John’s singing the Marseillaise, thought of as the stuff of the event that is John’s singing of the Marseillaise, does not admit of the same type/token distinction that processes and events as I am understanding them do. John’s singing the Marseillaise today is the same activity as John’s singing the Marseillaise tomorrow. But the process of John’s singing the Marseillaise today is most likely distinct from the process of John’s singing the Marseillaise tomorrow.

So I distinguish three levels. Firstly there is the activity - John’s singing the Marseillaise. This is referred to using an imperfect nominal. Then there is the process of this activity. Finally there is the event of this process. The perfect nominal, “John’s singing of the Marseillaise”, may refer to either the process or the event. And it is with

the clarification of the metaphysical distinction between these levels that I am concerned here.

Things that will have happened (events) have temporal extension and temporal parts. There are several stages to what will have happened in the course of the event of my travelling to work. The beginning and end of the event and the various stages of it are all part of what will have happened. Each can be located in a temporal interval which is a sub-interval of that of the whole journey. The event that will have happened may have been lengthy or short, eventful or uneventful. And it has these properties timelessly.

By contrast, the properties that processes have are had at a time, not timelessly. If I say of the process of my travelling to work that it is making me bored, I am attributing a property to this process at a time. The same process may not have this property in ten minutes time. Processes change their properties as they persist. What is happening now, while being the very same token as what was happening ten minutes ago is different from what was happening ten minutes ago. It has changed. But what happened ten minutes ago has exactly the same properties as any event it is identical with. There is no room for temporal modification of the instantiation relation when attributing properties to events; but there is for processes.

In general if what happened has a certain property, then it did and always will have that property. If a particular journey (the thing that happened) was tedious, then it was and always will be tedious. Of course, a journey may have tedious stages and exciting stages. But these stages are quite timelessly tedious or exciting. If I say that the event of my travelling to work this morning was tedious at half past eight, I can only mean that the half past eight stage of the journey was timelessly tedious. (I am assuming that the word "journey" picks out an event - a thing that will have happened. But, like the word "life", it may also be used to pick out a process. For example, it is perfectly proper to say that a journey is becoming tedious. This is a tensed attribution of a property to a process.)

But if what was happening has a property, then it may not have had that property before and it may stop having that property. The process of my travelling to work may have been tedious then stopped being tedious. What happened is extended in time. What was happening is a continuant; it endures, alters and might last for ages. Temporal extension is a property that things have timelessly, and so not the sort of property that processes have.

In these respects processes are like people and events are like lives. Indeed the relationship between the two pairs is more than just an analogy. A person's living is something that will have been happening; so it is a process. The event of the person's living is something that will have happened, and can plausibly be identified with the person's life. So the life of a person is the event of the process of a person living.

This means that there are two stages to the relationship between a person and their life: the relationship between a person and the process of that person living; and the relationship between the process of living and the event of that process - the life. Although a person and the process of that person living - i.e. what is happening with that person through their life - are clearly very different things, they both belong to the metaphysical category of continuants. (It is not my task here to explain the distinction between processes - dynamic continuants - and physical objects - static continuants, though clearly my approach owes such an explanation.)

4. The mereology and identity of processes and events

The earlier observations about the ways in which our talk of lives may be abstracted from our talk of people can be carried over to talk of events and processes. In general, I think it is possible to derive what we should say about the events of processes that will have been happening from what we say about the processes themselves. So we can try to construct a semantics for the “event of” operator. This is by contrast with attempts by Galton and Taylor to give semantics for our talk of things happening in terms of our talk of things that have happened – a semantics for the “process of” operator.¹³

I suppose that a thorough-going semantics for the “event of” operator would give an account of how all the timelessly attributed properties of events of processes may be derived from the properties attributed at a time to processes. I am not attempting this here, but will limit my task to making a tentative suggestion about one way that the mereology and identity of the events of processes derives from the mereology and identity of the processes themselves.

Since processes do not have temporal extension, they do not have temporal parts. For a part to be a temporal part it must occupy a temporal interval which is a sub-interval of the temporal interval occupied by the thing as a whole. But processes, unlike events, do not occupy temporal intervals. Something should be taken to have temporal parts only if it can notionally be divided temporally to yield these parts. But only something that is extended in time can be divided temporally.

But there is a sense in which processes have parts. Consider the process that was my travelling to work this morning. At one stage while this process was happening I was walking to the station. The process of my walking to the station was part of the overall process of my travelling to work. The question of what makes one process part of another is beyond the scope of this paper, although I acknowledge that my claims about how the metaphysics of events depend on the metaphysics of processes only provide a very incomplete picture since I am refusing to commit myself to any claim about the metaphysics of processes here.

Both processes were happening at the same time though the overall process was going on for longer. Indeed someone might try to claim that the process of my walking to the station was a temporal part of the overall process of my travelling to work since it was happening during a sub-interval of the time during which the overall process was happening. But as I will now try to show, this would be the wrong way to describe the relationship between the two processes.

It helps if we start off by considering a slightly different sub-process - not the process of my walking to the station but the process of my walking to my reserved seat at the back of the train. I walk to the station, on to the platform, on to the train, and continue walking even when the train starts moving until I get to my seat at the back of the train. All this time the same thing is happening - the process of my walking to my seat on the train. When I am walking to the station, the process of my walking to my seat on the train is part of the overall process of my travelling to work. But when I am walking to the back of the train, even though this is exactly the same process of walking to my seat that has been happening since I left my front door, it is not at this time part of

the overall process of my travelling to work. It is something that is happening alongside the overall process, but forms no part of it.

What this shows is that the process of my walking to my seat is only a temporary part of the overall process. It is not timelessly part of the overall process; it is at a time part of the overall process. It is like a hair being at one time part of my body and being at another time, after I have shed it, not part of my body. Hairs are temporary parts of bodies. Temporary parts should be contrasted with temporal parts. Temporal parts are timelessly temporal parts whereas temporary parts are at a time parts. The hair exists as part of the body during a sub-interval of the interval within which the body exists; but this does not make it a temporal part.¹⁴

In this example, the process of my walking to my seat is continuing even after it is no longer part of the process of my travelling to work. This makes it clear that it can only be part of the process of my travelling to work at a time and not timelessly. The case is only slightly less obvious with the process of my walking to the station. In this case, the sub-process is not happening except when it is part of the overall process. Nevertheless, it is still only a temporary property of the overall process of my travelling to work that the process of my walking to the station is part of it. When I am sitting on the train, the process of my walking to the station is not currently part of the process of travelling to work. When I was walking to the station that was part of the process of my travelling to work; afterwards it was not. It is like a freckle, which can only exist as part of my body, but is nevertheless a temporary part. It is at a time part of my body just as the process of my walking to the station is at a time part of the process of my travelling to work.

I now want to make a tentative suggestion as to one principle that might govern our talk of the mereology of the events of processes. It is this:

P1 The event of process A is (timelessly) a temporal part of the event of process B if and only if process A is for its whole duration part of process B.

While the process of my travelling to the station is happening, that process is part of the process of my travelling to work. So, according to this principle, the event of my travelling to the station is part of the event of my travelling to work. On the other hand, the process of my walking to my reserved seat at the back of the train is not for its whole duration part of the process of my travelling to work; so the event of my walking to my seat at the back of the train is not part of the event of my travelling to work.

We can take as a limiting case of this suggestion concerning event mereology the following suggestion concerning event identity:

P2 The event of process A is numerically identical to the event of process B if and only if process A is numerically identical to process B.

Consider first the right to left implication - i.e. if process A = process B then the event of process A = the event of process B. This is to say that the “event of” operator is a mapping. I have already smuggled this in as an assumption by using the word “the” in talking about the event of a process. But the assumption that each process has only one event should be examined.

On the face of it the assumption seems obviously false. Consider the process of the earth orbiting the sun. What is the unique event of this process? It seems that we can identify the event of the first million orbits of the process of the earth orbiting the sun, or the event of the first million and one orbits. These are distinct events; but they correspond to the very same process. Different things happened even though only one thing was happening.

But the principle for the identity of events of processes can be held onto despite this. We should not require that every event is an event of a process. Corresponding to each process is an event that is the unique event of that process. Then further events can be derived from that event. Corresponding to the process of the earth orbiting the sun is the event whose temporal extension is exactly that time during which the process of the earth orbiting the sun is happening. The beginning of the event occurs when the process starts happening and the end of the event occurs when the process stops happening. Neither the event of the first million orbits nor the event of the first million and one orbits is directly the event of a process. They are both events in a derivative way by being parts or temporal segments of the event of the process of the earth orbiting the sun.

Since events are temporally extended it is convenient to have some kind of temporal mereology for them. It is useful to break events into parts even when these different parts do not correspond to different processes. These second-order events depend on processes at one remove and may be found a place in my account without too much difficulty. For example, we might have the following principle for constructing temporal segments of events. Unlike P1 this is not limited to the events of processes:

P3 If E is an event that lasts from t_1 to t_4 and t_2 to t_3 is a segment of t_1 to t_4 then there is an event E/ that lasts from t_2 to t_3 and is a temporal part of E.

So far we have the fairly weak suggestion that there is a mapping from processes to events. Not every event is in the range of this mapping. Some events may be identified derivatively from events in the range of this mapping as a result of temporal division of such events. Perhaps there are some events that do not depend on processes in any way at all. Bare Cambridge changes might be thought to be such events. If an object has some property at one time and does not have it at a later time, then we can say that a change has happened. The existence of this change does not seem to depend on the existence of any process but simply on the fact of an object having and then not having a property. It seems here that nothing can be said to have been happening although something has happened.

It is not uncontroversial that Cambridge changes are substantial enough to constitute particular events. But rather than getting into that issue here, I want to duck it by limiting myself to considering the qualified claim that a certain important class of events are to be identified in terms of processes. Put rather feebly the claim is that those events which are the events of processes depend for their identity on their processes.

The principle for event identity that I am considering is a biconditional. According to this principle, not only is there a mapping from processes to events of processes, but this mapping is a 1-1 mapping. No two processes yield the same event. This is equivalent to the reasonable thought that what happened cannot be the same if what was happening was not the same.

But my suggested principles for the mereology and the identity of events of processes become much less obvious if they are applied across possible worlds. Applied across all possible worlds they yield a novel essentialist thesis about events, one which conflicts with the standard kinds of time and place essentialism. The thesis is that the only thing that is essential to an event of a process is the particular process it is an event of.

The principle may be made to look implausible if we consider the possible world in which I trip and break my leg while leaving the house and end my journey to work at my front door. It does not seem entirely convincing to claim that this is the very same journey as the actual one in which I traveled all the way to work. Would the event of my writing this paper have been the very same event if it had taken me twenty years to do? I am not sure that there would be any bad metaphysical consequences of saying that it would have been the same event, but it is far from clear that our ordinary talk of events is determinate with respect to this sort of question.

I think we do not have sufficient intuitive grasp on the modal properties of event tokens to feel absolutely sure about this question of event identity. Indeed, at this stage there may be no fully determinate metaphysics of events. This rather anti-climactic result does not discourage me. It seems to count as further evidence that events do not figure as centrally in our metaphysical thinking as processes do.

If my suggested principle for event identity were accepted even across possible worlds, there arises the question of why we do not simply identify events and processes. After all, we do say things like: "What is now happening will one day have happened." And this looks like we are identifying processes and events in ordinary language. But we have to be very careful in interpreting this, since processes and events may be referred to by the same phrases. For example, my travelling to work is both a process that will have been happening and an event that will have happened. The fact that the event and the process are both referred to as my travelling to work does not mean that they are identical. So when we say that what is now happening will one day have happened we might be saying that the process that is now happening is of the same type as the event that one day will have happened. In the same sort of way we might say that some book will one day be a blockbusting movie, without identifying books with movies.

So it is not obvious that events and processes are identical. But equally it is not obvious that they are distinct. A possible argument against identifying events and processes would be that since events have temporal parts and processes do not, then, by Leibniz's Law, they must be distinct. How can one and the same thing both have and not have temporal parts? But I argued earlier that Leibniz's Law must be relativized to the instantiation relation. Events timelessly have temporal parts. Processes do not have temporal parts because the properties that processes have are not had timelessly. So we cannot say that events timelessly have temporal parts and processes timelessly have no temporal parts. Nor can we say that processes at a time endure while events at a time do not endure. Processes and events belong to different categories as far as the temporal mode of their instantiation relation is concerned; so Leibniz's Law cannot be applied to them both at once.

So it is a possible view that there is really no ontological distinction between processes and events. We just have two different ways of talking about the same things - two different aspects. We can talk about them as on-going or as gone-over. Since these

different ways of talking - the process-way and the event-way - involve different sorts of instantiation relation we will not get any contradiction by identifying the objects of these different ways of talking. If ontological parsimony were motivating us then this would be quite a reasonable position. The various claims I am making about the distinctions and relations between processes and their events would then have to be read as claims about the distinctions and relations between process-talk and event-talk. The key claim would be that event-talk depends on process-talk.

NOTES

¹ Alfred North Whitehead, *Process and Reality* (Cambridge: Cambridge University Press, 1929), pp. 94-5. I paraphrase to avoid Whitehead's rather confusing language.

² Bertrand Russell, *The Principles of Mathematics* (Cambridge: Cambridge University Press, 1903), section 442. The approach has been defended more recently by Wesley Salmon, *Scientific Explanation and the Causal Structure of the World* (Princeton, NJ.: Princeton University Press, 1984).

³ Mark Johnston, "Is there a problem about persistence?" *Proceedings of the Aristotelian Society* Supp vol. 61 (1987): 107-35. David Lewis, *On the Plurality of Worlds* (Oxford: Blackwell, 1986). Willard van Orman Quine, "Identity, Ostension and Hypostasis," *Journal of Philosophy* 47 (1950): 621-33.

⁴ Willard van Orman Quine, *From a Logical Point of View* (Cambridge Mass.: MIT Press, 1953), ch. 1.

⁵ John Perry, "Can the self divide?" *Journal of Philosophy* 69 (1972): 467 ff.

⁶ Peter Strawson, *Individuals* (London: Routledge, 1959) ch. 1, and Gareth Evans, *The Varieties of Reference* (Oxford: Oxford University Press, 1982), pp. 174-5.

⁷ For one of countless examples, see Peter Simons, *Parts: A Study in Ontology* (Oxford: Oxford University Press, 1987), pp. 129f.

⁸ See Stout, "Processes," *Philosophy* 72 (1997): 19-27. See also Hugh McCann, "Nominals, facts and two conceptions of events," *Philosophical Studies* 35 (1979): 129-49. McCann makes a very clear case for the distinction between temporally persistent and temporally extended events.

⁹ A. P. D. Mourelatos, "Events, processes and states," *Linguistics and Philosophy* 2 (1978): 415-34. See also Barry Taylor, *Modes of Occurrence* (Oxford: Blackwell, 1985), and Helen Steward, *The Ontology of Mind: Events, States and Processes* (Oxford: Oxford University press, 1997).

¹⁰ A. P. D. Mourelatos, "Aristotle's kinesis/energeia distinction," *Canadian Journal of Philosophy* 23 (1993): 385-88.

¹¹ The example is taken from McCann, "Nominals, facts and two conceptions of events."

¹² Jonathan Bennett, *The Act Itself* (Oxford: Oxford University Press, 1985), pp. 31f. Here he is following Zeno Vendler, *Linguistics and Philosophy* (Ithaca: Cornell University Press, 1967), ch. 5, and consolidating his own earlier position in *Events and their Names* (Oxford: Oxford University Press, 1988).

¹³ Anthony Galton, *The Logic of Aspect* (Oxford: Oxford University Press, 1984), and Taylor, *Modes of Occurrence*.

¹⁴ See Simons, *Parts*, for a useful exposition of this distinction between temporary and temporal parts.