



Universiteit  
Leiden  
The Netherlands

## On the causal nature of time

Gijsbers, V.A.; Kleinberg, S.

### Citation

Gijsbers, V. A. (2019). On the causal nature of time. In S. Kleinberg (Ed.), *Time and causality across the sciences* (pp. 49-71).

Cambridge: Cambridge University Press.

doi:10.1017/9781108592703.004

Version: Publisher's Version

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3239295>

**Note:** To cite this publication please use the final published version (if applicable).

# 4

## On the Causal Nature of Time

Victor Gijssbers

### 4.1 Introduction: Time and Causation

What is the relation between time and causation? That there *is* a relation seems extremely likely, since, to give just one example, causes tend to happen before their effects rather than after them.<sup>1</sup> But this observation immediately raises an important question. Suppose that it is true that event *a* is a cause of event *b*; and that it is also true that *a* is earlier than *b*. What is the relation between these two facts? Can one of them perhaps be explained by the other?

David Hume, whose ideas continue to be a major influence on the philosophy of causation, answered this question in the affirmative. He gives the following definition of ‘cause’:

An object precedent and contiguous to another, and where all the objects resembling the former are placed in like relations of precedency and contiguity to those objects that resemble the latter. (Hume, 1739/1978, sec. 1.3.14)

For Hume, then, causation is nothing over and above one type of event always being followed by another type of event. By definition, the cause is precedent to (earlier than) its effect. If events of type *A* are always followed in time by events of type *B*, then by definition any particular event *a* is the cause of the event *b* that follows it. The temporal relation between events is thus more fundamental than the causal relation, since causation is defined in terms of – among other things – time.

If we look at Hume’s overall philosophy, it is clear why he wants and even needs to define causal relations in terms of temporal relations. Hume’s crucial presupposition is that all our ideas must be based on perception, since any

<sup>1</sup> One could also point at the important role of time lapses in the perception of causation, as described in the seminal psychological studies of Albert Michotte (1946); or at the fact that philosophical discussions about the possibility of time travel invariably lead to issues of backwards causation and closed causal loops (Smith, 2018, section 3).

concept that cannot be traced back to perceptions is merely an empty word. Hume also believes that we can perceive temporal relations, but that we have no (further) perception of causation. I can *see* that I put my hand in the fire just before I felt the pain, but there is nothing additional in my experience that corresponds with a causal relation between the two events. We cannot perceive causation. But then how is it possible, given Hume's presupposition, that we have the idea of causation? Hume's answer is that causation can be defined in terms of time.<sup>2</sup>

Now, for Hume's strategy to work, his definition of causation must be *reductive* – that is, it must define the concept of causation in terms that are not themselves causal. For instance, if it turns out that the term “precedent” is to be defined in terms of cause and effect, then Hume would have defined causation in terms of causation and he would not have achieved his philosophical aim. If we call *causally innocent* all concepts that are not based on the concept of causation, and all other concepts *causally culpable*, then it is crucial for Hume's project that temporal concepts such as precedence are causally innocent. Let us call this claim the *causal innocence thesis*.<sup>3</sup>

Contemporary philosophers tend not to share Hume's conviction that all our ideas are based on perceptions. But the form of his theory of causation, namely, giving a reductive definition of causation in terms that include temporal concepts, is still the form that contemporary theories of causation generally take. Almost all the major approaches to causation – including regularity theories, probabilistic theories, causal process theories, and counterfactual theories<sup>4</sup> – aim to tell us what causation is by defining it in non-causal terms. The general strategy of these approaches is to provide rules for deciding which events cause which other events, given some rich enough description of the world. And for all of them, temporal relations are an important part of this description.

Again, Hume's theory can be used to illustrate this. Suppose that I'm wondering whether putting one's hand in the fire causes the pain that follows it. Hume's theory allows me to answer this question, given that I am provided with a sufficiently rich description of the world. What I need to be able to do is consider every holding-one's-hand-in-the-fire event, and then check for each of them whether they are followed by a temporally and spatially contiguous pain

<sup>2</sup> Hume's total story about causation is more complicated than this brief exposition suggests, but we need not go deeper into it for present purposes. See Morris and Brown (2017, section 5 and 6), for an overview.

<sup>3</sup> We cannot go deeply into methodological issues here, nor adequately discuss the relation between metaphysics and conceptual analysis. I will assume that we have temporal and causal concepts that adequately latch on to something real – time and causation. The causal innocence thesis then is to be understood as *both* the claim that temporal concepts can be possessed without the possession of causal concepts *and* the claim that time does not ontologically depend on causation. Hume certainly endorses both these propositions.

<sup>4</sup> See Beebe et al. (2009) for a good set of introductory essays.

event. If all of them are, if holding one's hand in the fire is always followed by pain, then I can conclude that there is a causal relation between them. (This is simply Hume's definition of cause.) Otherwise, there is no causal relation. Since during this process I need to determine the precedency and contiguity of events, the sufficiently rich description of the world has to contain temporal information.

Hume's theory wears its reliance on temporal information on its sleeve, while many recent theories do not. But the reliance is nonetheless there. For instance, it might initially seem that the counterfactual theory of Lewis (1973a) does not need spatial and temporal information, since it defines causation in terms of what would have happened in different circumstances. (Roughly, Lewis claims that *a* causes *b* if and only if *b* would not have happened if *a* had not happened.) But if we look at how Lewis evaluates claims about what would have happened in different circumstances, it turns out that temporal information is indeed crucial.<sup>5</sup>

So most contemporary theorists of causation are in the same boat as Hume: they can only achieve their philosophical aims if the causal innocence thesis is true. Otherwise, they are only providing us with ways of making causal judgments after we have already made them, which is of decidedly limited utility. Contemporary philosophical thinking about causation thus gives a high degree of urgency to the question whether temporal concepts are indeed causally innocent, or whether they are somehow based on or inextricably entwined with causal concepts.<sup>6</sup>

Unfortunately, this question is rarely asked. One may wonder why; and perhaps the answer is that philosophers, under the influence of a neo-Humean orthodoxy, have simply taken the causal innocence thesis for granted. If this is indeed the case, then the current chapter, which argues against the causal innocence thesis, may awaken us from our dogmatic slumber – or may at least ensure that our sleep is haunted by strange Kantian nightmares.

Now, there are two distinct ways in which temporal concepts could be causally culpable: temporal concepts could depend on causal concepts; or temporal and causal concepts could be mutually dependent on each other. These two options correspond to two important historical precedents: *causal theory of time*, which attempt to analyse time in terms of causation; and the argument developed by Kant, especially in his so-called 'second analogy of experience', to the effect that conceiving of the world as temporal and conceiving of it as causal are two sides of the same coin. I will argue that

<sup>5</sup> Lewis asks us to judge the closeness of a possible world in part by assessing the size of the spatiotemporal region where this world exactly matches ours.

<sup>6</sup> See Gijsbers (n.d.) for a more thorough analysis of the ways in which contemporary theories of causation require us to answer questions of causal innocence.

the second, Kantian, option – which we may call the *mutual dependence thesis* – is to be preferred.

The chapter is organised around these two historical precedents. In Section 4.2, we will discuss causal theories of time. An early example of such a theory is found in the work of Leibniz (see Futch (2008)), but they were especially popular around the middle of the twentieth century and then disappeared again in the 1970s. I will show that the arguments brought against them at the time did not warrant this disappearance. But I then go on to develop a new argument against causal theories of time: the *argument from synchronicity*. Somewhat surprisingly, a slightly adapted version of the argument also works against many forms of the causal innocence thesis.

In Section 4.3, we will go on to look at Kant's second analogy of experience. Kant argues that we cannot think of objects as placed in time unless we also think of them as falling under universal causal laws. It is a solid argument, but its reliance on Kant's philosophy of transcendental idealism makes it less suitable for convincing philosophers (and scientists) who are not committed to this metaphysical position. (And that would be nearly everyone.) For that reason, I will develop an argument of my own that is inspired by Kant, but approaches the topic from a less controversial basis. I show that insisting on the causal innocence thesis leads to an extreme form of scepticism about temporal claims, a form of scepticism that nobody wants to embrace.

If there is an a priori link between time and causation, but this is neither because temporal concepts are based on causal concepts, nor because causal concepts are based on temporal ones, then there are only two possibilities left: either there is some third concept, *X*, on which both causal and temporal concepts are based, or the mutual dependence thesis is true. But there are no plausible candidates for *X*.<sup>7</sup> We are thus led towards embracing the mutual dependence thesis, which should lead to a radical rethinking of our philosophical approaches to both causation and time.

## 4.2 Causal Theory of Time

### 4.2.1 Introduction

The most ambitious attempt to prove the causal culpability of temporal concepts, and the one that would deliver the most dramatic result if it were

<sup>7</sup> There are attempts to explain the *direction* of causation and the *direction* of time in terms of some third *X*, say, the direction of increasing entropy. But these attempts presuppose the concepts of time and causation. Note, for instance, that 'the direction of increasing entropy' means 'the direction *of time* in which entropy increases'.

to be carried out successfully, is the construction of a *causal theory of time*. A causal theory of time is a theory that reduces our temporal concepts to causal concepts. Here is a toy example:

*a* is earlier than *b* if and only if *a* causes *b*.

This theory is false, since events often precede other events without being their cause: for instance, the death of Moctezuma II happened about six months before the excommunication of Martin Luther, but there is no causal relation between them. However, at least earlier events *can* causally influence later events, while the opposite is often held to be impossible. A historian *might* convince us that the death of Moctezuma did cause Luther's excommunication, through some extraordinary chain of events, but no historian could convince us that Luther's excommunication caused Moctezuma's death six months earlier. So a more defensible causal theory of time might be the following:

*a* is earlier than *b* if and only if *a* could causally influence *b*.

Developing and defending this theory would not be easy, in particular because we would have to show that the idea of possible causal influence can be understood without a prior understanding of temporal relations. But the idea that a theory like this might be correct was quite popular among scientific-minded philosophers of time around the middle of the twentieth century, especially because of the way it fits Einstein's special theory of relativity.

In special relativity, whether one event *a* is earlier or later than another event *b* often depends on the velocity of the observer. But not always: some events are in the *absolute future* (or the *absolute past*) of *a*. For such events, all possible observers agree that they happen later (earlier) than *a*. The theory identifies the absolute future of *a* with those parts of space-time that could be reached by a signal from *a*; in other words, with all parts of space-time that could be causally influenced by *a*. The absolute past, by contrast, is identified with all parts of space-time that could causally influence *a*. So the idea that the structure of causation determines the structure of space and time is, if not quite inescapable, at least one that naturally arises within the confines of special relativity.

And thus it was that although the idea of a causal theory of time can already be found in the work of Leibniz,<sup>8</sup> it rose to modern prominence in the

<sup>8</sup> Van Fraassen 1970 chapter II.3 traces the history of causal theories of time through the writings of Leibniz, Kant, and Lechallas, while Mehlberg (1980) presents a more thorough treatment of the same topic. Futch (2008), chapter 5, gives a full discussion of Leibniz's causal theory of time. We will come to Kant in §4.3.2, but can remark here that his aim is not to construct a causal theory of time in the strict sense: he does not believe that we can reduce time to causation.

work of Hans Reichenbach and Adolf Grünbaum, who were deeply involved in philosophical reflection on Einstein's theories of relativity. Reichenbach developed two versions of his causal theory of time (1928; 1956), and so did Grünbaum (1963, 1968). In each case, the latter version was supposed to fix some of the problems in the former (van Fraassen (1970), chapter 6). Van Fraassen (1970) then developed Grünbaum's theory yet further. It is van Fraassen's version that I want to look into, since it is both relatively simple and a representative example of the movement as a whole.

The plan is as follows. First (§4.2.2) we discuss van Fraassen's theory. We will see that he and his predecessors make their own job very hard by limiting themselves to a symmetric concept of causation. If we take an asymmetric concept of causation, we can formulate a much simpler causal theory of time (§4.2.3). We then (§4.2.4) delve into four arguments that were developed against causal theories of time: the argument from the relative unclarity of causation; the argument from modality; the argument from general relativity; and the argument from causally isolated events. The first three all depend upon a scientific conception of metaphysics. The fourth forces the causal theorist to embrace the Eleatic Principle, according to which causation is a criterion of reality – but she might be completely happy to do so. Finally, I develop a new argument against the causal theories of time, an argument based on synchronicity (§4.2.5). We will see that this argument suggests that we cannot reduce temporal concepts to causal concepts, but that temporal concepts and causal concepts depend upon each other.

### 4.2.2 Van Fraassen's Causal Theory of Time

Van Fraassen (1970, section 6.4) bases his causal theory of time on the idea of 'genidentity', where two events are genidentical just in case they involve the same object. The events of me being born, me writing this article and me dying are thus genidentical, since they all involve the same object: me. For any object  $O$ , the set of all such events is called 'the world line of  $O$ ', which we can abbreviate as  $W_O$ .

As yet,  $W_O$  is an unordered set: there are some events, all pertaining to the same object, but we don't know which events come before others. If we want to make temporal claims and say that one event was earlier or later than another, we need to introduce an ordering on  $W_O$ . Now, it may seem as if introducing an ordering is the simplest thing in the world: we just use the asymmetrical relation of causation. My state right now is certainly one of the causes of my state in the near future, while my state in the near future is not

a cause of my state right now. So the ordering of the events in  $W_O$  is just the causal ordering, and this perfectly correlates with the temporal ordering of the events.

This would indeed be both possible and simple, but, for reasons that we will discuss later, Van Fraassen – like Grünbaum and the later Reichenbach – does not want to make use of an *asymmetrical* notion of causation. So he has to create an ordering in some other way.

He starts by defining the notion of a ‘continuous part’ of a world line. If we have a criterion for continuous parts, we can then say that  $b$  is between  $a$  and  $c$  just in case  $b$  lies on every continuous part of the world line that also contains  $a$  and  $c$ . To see this, just think about the real numbers. A number is between 3 and 5 just in case it is part of every continuous interval that includes 3 and 5. (Van Fraassen, like Grünbaum, actually introduces a more complicated scheme in order to be able to deal with circular time structures that are possible in general relativity, but we will ignore this.)

In order to define the notion of a continuous part, Van Fraassen uses the *symmetrical* concept of ‘causal connectibility’, as well as its opposite, ‘topological simultaneity’. Two events are topologically simultaneous just in case they are not causally connectible. (In special relativity, this would mean that they are not in each other’s past or future light cones, and thus that they will be simultaneous for some possible observers – hence the term.) Now, take an object  $O$  and some event  $e$  that is not part of  $W_O$ . Then, van Fraassen tells us, the set of events that lie on  $W_O$  and that are topologically simultaneous with  $e$  form a continuous part of  $W_O$ . The underlying intuition is that  $W_O$  has three parts: one that lies in the absolute past of  $e$ , one that lies in its absolute future, and then one in the middle that is topologically simultaneous with  $e$ . This van Fraassen defines as continuous, and we can see why, given that there are supposed to be no causal gaps in world lines. With this criterion for being a continuous part, and some further logical and mathematical apparatus, van Fraassen is capable of linking the events in  $W_O$  to the real numbers and proclaiming this link to be a time-coordinate assignment of the kind that physicists know and love.

The ingenuity of the theory is impressive; but it is a curious that almost all of it is needed *only* because van Fraassen starts from a *symmetrical* conception of causation. So why did Reichenbach (in the second version of his theory), Grünbaum, and van Fraassen all choose to base their causal theories of time on such a symmetric conception? We can start to understand this from the criticism levelled at Reichenbach’s original 1928 formulation of the theory. Van Fraassen’s objection is typical:



The major criticisms of this theory center on Reichenbach's use of the notion of *cause*. Since Hume, no philosopher can afford an uncritical use of this notion. But even if one takes the view that the notion of causal connection is prephilosophical and that the question is not whether there are causal connections but how they are correctly described, Reichenbach faces a problem. For he relies explicitly on the asymmetry of such connections, on the distinction between cause and effect. If he wishes to say, like Leibniz, that by definition the 'earlier' one of a causally connected pair is the cause, then he must provide a criterion for distinguishing the cause from the effect. (Van Fraassen, 1970, 174)

To this, the appropriate response would seem to be puzzlement. If one takes the causal relation as basic and undefinable, why not take its asymmetry as basic and undefinable too? Historically, the answer must lie in the empiricist commitments that Reichenbach, Grünbaum, and van Fraassen all share with Hume, as well as their idea that the difference between a cause and an effect cannot be observed. But one need not be an empiricist; and even if one is, one need not hold that we cannot experientially distinguish between causes and effects. To adapt an example from Zwart (1967, 25), suppose that I am frightened by an evil clown who suddenly jumps from the bushes. Surely, it is not the case that I experience both the clown and the fright but am left to speculate about the existence and nature of causal relations between them. I experience *being frightened by* the clown, and this experience leaves no doubt as to which is cause and which is effect. Asymmetry is part of the common sense concept of causation; and there seems little reason not to use it in a causal theory of time. Doing so allows us to simplify things considerably.

### 4.2.3 A Simple Causal Theory of Time

Think of all events in the universe. Now order them along a line in such a way that causes always appear to the left of their effects. Given that there are no causal loops, this is possible. (There will in fact be infinitely many such arrangements – simply choose one.) Imagine that the line is labelled with the real numbers, just like the x-axis of a Cartesian coordinate system. Then we can simply look where each event falls on the line and assign the corresponding number to it. In this way, every event will get a number, and indeed a number that is less than the number given to any of its effects. We can call the line *time* and the numbers *temporal coordinates*. And that is all that time is: a numbering system set up to reflect the causal structure of the world. This, in a nutshell, is our simple causal theory of time.

**Simple Causal Theory of Time:** Temporal coordinates are nothing more than a pragmatically useful assignment of real numbers to events in such a

way that causes are always assigned lower numbers than their effects. This is all there is to say about the nature time.

We could leave it at that, but it may be useful to work out the idea in a little more technical detail. Anyone uninterested in such technicalities is invited to skip to the next section.

First of all, we must choose a specific causal relation as the basis of our theory. For instance, we could start from the relation of *direct cause*. Suppose that I order you to throw a stone at a window; you do so; and the window breaks. Then my order is a direct cause of your throwing; your throwing is a direct cause of the breaking; but my order is an indirect rather than a direct cause of the breaking, since your throwing lies ‘in between’. If we further assume that there are no causal loops, we can model the structure induced on the set of events by the relation of direct causation as a *directed acyclical graph*. This is the standard approach in the causal modelling literature (see, e.g., Pearl (2000)).

But let us instead start from the relation of *causal ancestry*, where  $a$  is a causal ancestor of  $b$  just in case there is a chain of direct causation from  $a$  to  $b$ . My order to you is a causal ancestor of the breaking of the window, since it directly causes your throwing and your throwing directly causes the breaking.<sup>9</sup> We will assume that the relation of causal ancestry is irreflexive (nothing is a causal ancestor of itself), asymmetric (if  $a$  is a causal ancestor of  $b$ , then  $b$  is not a causal ancestor of  $a$ ), and transitive (if  $a$  is a causal ancestor of  $b$  and  $b$  is a causal ancestor of  $c$ , then  $a$  is a causal ancestor of  $c$ ). In other words, the causal ancestry relation is a *strict partial order*.

Now let  $t(x)$  be a function from events to real numbers such that if  $a$  is a causal ancestor of  $b$ , then  $t(a) < t(b)$ . That there is such a function  $t$  is guaranteed by the fact that the causal ancestry relation is a strict partial order, and by the further assumption that the cardinality of events in the universe is not greater than the cardinality of the real numbers. Indeed, infinitely many such assignments are possible. All of these are *possible time-coordinate assignments*.

We can just choose one of these assignments and call it *time*. In practice, however, some assignments are much more useful than others. So let us designate a certain class of periodic processes as ‘standard clocks’. Choose an

<sup>9</sup> The advantage of the causal ancestry relation is that whether it holds does not depend on how we model the causal sequence. If we decide to add the event of the stone flying towards the window to our model, your throwing of the rock can no longer be said to directly cause the shattering; it now does so indirectly by directly causing the flying. But the relation of causal ancestry is independent of such choices.

arbitrary number  $\Delta t$  and impose the restraint on time-coordinate assignments that when they assign the number  $t$  to a standard clock event, they should then assign  $t + \Delta t$  to the equivalent event of that standard clock's next period. (Current science has chosen the transition between the two hyperfine ground states of the caesium 133 atom as the standard clock and has set  $\Delta t$  to 1/9192631770.) Then let us choose one of the remaining possible time-coordinate assignments and dub it the *actual time-coordinate assignment*. We can identify time with this particular assignment. And that is all that time is: a way of assigning numbers to the events in the universe that is informed by their causal structure.

#### 4.2.4 Arguments Against Causal Theories of Time

The idea that the temporal order simply is the causal order, and that we can introduce time-coordinates by setting up a correspondence between the natural numbers and the partial strict ordering induced on events by the relation of causation, has much to recommend itself. It seems eminently possible and it promises the kind of reductive analysis that philosophers (especially analytic ones) tend to love. So why did causal theories of time disappear from the scene? To understand this – and to decide whether this disappearance is to be applauded or decried – we need to consider the criticisms levelled against the causal theories. Three critical papers that together appear to have spelled the causal theory's doom are Lacey (1968), Smart (1969), and Earman (1972). Lacey's paper is very dependent on the details of Grünbaum's theory, and hence is of limited use to us. But it will be helpful to study the ideas of Smart and Earman.

Let us begin with Smart. He levels three main criticisms. The first is that temporal concepts are clearer than causal concepts, so that it makes little sense to try to elucidate the former by the latter. Carl Hoefer, in his very brief retrospective comments about the causal theories of time, echoes this sentiment (Hoefer, 2009, 697, n. 6), as does Earman (1972, 83). But this surely begs the question against the causal theorist. If causal concepts are unclear, and temporal concepts are causally culpable, then any simplicity that seems to attach to temporal concepts is illusory. In addition, anyone who has studied the philosophy of time will be surprised to hear that temporal concepts are supposed to be clear, rather than, say, among the most mysterious there are.

I suspect that this idea that temporal concepts are clear, while causal concepts are unclear, is based on the idea that a concept is bound to be clear if it has a well-defined use in an established scientific theory. Within the

mathematical theories of physics, time is clearly defined. Causation, on the other hand, does not feature in such theories. But of course this criterion for clarity works only if one is willing to hold that the physicist's story about time captures all of time's essential features. But the deepest controversies in the philosophy of time, such as that between A-theorists and B-theorists,<sup>10</sup> hinge precisely on whether this claim is true.

Smart's second criticism is that causal theories of time have to use modal concepts, that is, concepts that involve possibility and necessity. *Causal connectibility* – a concept used by both Grünbaum and van Fraassen – is a modal concept since it describes the world not in terms of what is actually the case, but in terms of what *could have been* the case. In fact, Smart goes on, even the concept of causal connection itself might well be disguisedly modal. After all, we are inclined to say that if the cause hadn't happened, the effect *would not* have happened either.

Let us grant this. Why would the use of modal concepts be a problem for causal theories of time? Smart tells us that "a physical theory should be based on a purely extensional language, and the predicates '... is necessary' and '... is necessary for ...' should not occur in it" (Smart, 1969, 389). This argument seems puzzling. Surely the causal theory of time is not a physical theory, but a philosophical theory. So how could Smart's strictures on physical theories be relevant to it? The answer is that Smart is thinking of the causal theory not so much a philosophical theory about time, but as a philosophical analysis of temporal terms used in scientific theories. Again, what seems to be at work here is the idea that the essence of time is captured by physics.

Before moving on to Smart's third – and most interesting – criticism, it will be useful to discuss another argument, one that Smart mentions only briefly (Smart, 1969, 391–392) but that is central to Earman's critique of the causal theories of time: the argument, namely, that such theories cannot be made to fit the theory of general relativity. Earman's discussion is quite technical, but we can illustrate it with the following example. General relativity allows for universes in which time loops back on itself, so that by going far enough into the future (or the past), we will come back to the present. In such a universe, all events are causally connectible. Thus, constructions like the one given by van Fraassen will no longer work; and the simple causal theory presented in the previous section will also run into trouble, since the causal relation no longer induces a strict partial order.

<sup>10</sup> Roughly, A-theorists hold that the present is importantly different from the past and the future, while B-theorists hold that all moments of time have the same metaphysical status.

It might well be possible to save the causal theories of time from such problems. Van Fraassen (1972) claims so; and I suspect that the simple theory too could be made to work in cyclical time. (This would involve using the relation of 'direct cause' and the assumption that this relation is locally non-cyclical.) But suppose that Earman is right and that there are space-times allowed by general relativity that cannot be fitted into a causal theory of time. Why would the causal theorist be worried by this? Surely, she can simply claim that these models of the general theory of relativity do not describe anything we should call time. This in no way involves rejecting the theory of general relativity as a good scientific theory; it does not even involve rejecting a realist interpretation of that theory. What it *does* involve is rejecting the idea that the general theory of relativity *defines* what time is. But the causal theorist should be happy to deny that, since she holds that her own causal theory of time defines what time is.

Again and again, we see that these arguments against the causal theory of time are based on the same conviction: that physics tells us what time is and that a philosophical theory of time should thus be nothing over and above an elaboration or clarification of the physical theory. This makes sense historically. In the 1960s and 1970s, there was still a widespread distrust of metaphysics within analytic philosophy and hence an idea that if we could talk about seemingly metaphysical topics at all, it was only by linking them to established scientific theories. It thus also makes sense that these arguments made an impression and persuaded philosophers that causal theories of time might not be worth taking up. But the distrust in metaphysics has, if not vanished, at least dwindled; and it seems clear that none of the arguments discussed so far need dissuade a contemporary thinker from taking causal theories of time seriously.

This brings us to Smart's third argument, which appears to be the only one with clear contemporary relevance. He wonders how a causal theory of time can deal with events that are neither causes nor effects of other events:

It at least *seems* to me that I can consistently envisage a universe of purely random events spread out through space-time. The objection that only if such events were causes or effects of other events could they in fact be located ought, I think, to be dismissed as too verificationist. (Smart, 1969, 394)

Let a *closed causal set* be any set of events that is closed under the causal relation; that is, a set such that if event  $e$  is in the set, then so is every cause and effect of  $e$ . Then Smart is asking us to consider a world in which there are two or more distinct closed causal sets. (In the example he gives there would even be  $2^N$  such sets, with  $N$  the number of events in the world.) A

causal theory of time will allow us to establish temporal relations between events within one set, but it also commits us to the implausible, even absurd claim that there are no temporal relations between events in different sets. How could there be pairs of events in the world that are not temporally related? Doesn't this make a mockery of the very idea that they are part of the same world?

Indeed, that is precisely what the causal theorist must stress if she wishes to avoid the absurdity. The causal theorist must claim that it is impossible for a world to contain two or more distinct closed causal sets (ignoring the empty set). Why would that be impossible? Because a world just *is* a closed causal set. To be real, to be part of our world, just *is* to be (directly or indirectly) causally connected to us: the causal theorist should embrace what has become known in the current literature as the Eleatic Principle (e.g. Colyvan (1998); Cowling (2014)). Hence, Smart is mistaken when he thinks that he can consistently envisage a universe of events that are not causally connected. (This mistake is easily made, since one imagines oneself overlooking a world of random events without realising that this establishes oneself as a causal nexus that binds all the events together as a world. Take away the very idea of an observer. Is it still clear that one is envisaging anything at all, let alone a world?)

#### 4.2.5 The Argument from Synchronicity

So far, causal theories of time seem to be very much defensible – none of the arguments discussed in the previous section need worry the causal theorist. In this section, however, I wish to develop an argument against causal theories of time that is at the same time an argument in favour of the mutual dependence thesis.

The argument starts from the phenomenon of *synchronicity*. Carl Jung coined this word to designate an 'acausal connecting principle' (Jung, 1973) that could explain otherwise unexplainable coincidences. Thus far, my usage of the term agrees with his. But where Jung is thinking of rare, potentially life-changing events containing symbolic meaning, I am instead thinking of the most common of everyday phenomena – of the phenomenon, namely, that identical processes happen at identical 'speeds'.

Take two identical clocks and put them next to each other on a table. Let them both start out at twelve o'clock. We will witness the minute hands move clockwise until they are back at '12', while the hour hands will meanwhile move forward even more slowly to reach '1'. For each clock, there is a perfectly good causal explanation – involving their internal cogs and gears or

electrical circuits – for why these particular states follow each other in this particular order. But here is the crucial question: why do the two clocks reach one o'clock at the same time? Or, to ask the question without presupposing the notion of time: why does an observer looking at both clocks always see them in the same state?

Is there a *causal* explanation of this fact? That seems extremely unlikely. To give a causal explanation of a coincidence is to give some common cause that determined that the coincidence would take place. But for all intents and purposes, the two clocks in our example are causally isolated systems. (They are not truly isolated, and they have, for instance, been in gravitational contact. But it is hard to see how such interactions could provide an explanation for their synchronicity.) We have one causally explainable process and another causally explainable process, but there seems to be no causal explanation for the fact that these two processes run in tandem.

Wait a minute!, I hear you cry. The problem is that I am assuming that causal processes are to be defined as nothing but patterns of succession. But if I describe causal processes as successions of states over certain time intervals, then the problem immediately disappears. For instance, if the basis for my causal story about a quartz clock is only that the quartz crystal keeps alternating between two states, then there is indeed no way to explain why two quartz clocks run in tandem. But if I start by stating that the quartz crystal alternates between its two states 32,768 times per second, then it's a different matter. Then I will be perfectly capable of explaining why the two clocks run at the same speed.

Indeed. But this story is not available to those who defend a causal theory of time! For if a causal theory of time is true, then assigning 1/32,768th of a second to every alternation of the quartz crystal is merely a conventional choice for assigning certain labels to certain events. But no mere labelling convention can *explain* the fact that the two clocks keep showing the same state to the observer. (It would be like explaining the clockwise progress of the clock hands by pointing out that in that direction the positions on the clock face are labelled with progressively greater numbers.) Thus, a causal theorist must accept synchronicity as an unexplained brute fact.

Is it possible to explain synchronicity while holding on to the Humean claim that temporal concepts are causally innocent? At first sight, it may seem so. The claim that the quartz crystal alternates between its two states 32,768 times per second, does not seem to involve any *causal* concepts. But on the hypothesis of causal innocence, the problem of synchronicity rears its head in a different form. For we now observe that it is only causal processes that exhibit systematic synchronicity. Non-causal processes – for instance, the

appearance of the first, second, third and so on pedestrian in a given street on a given day – do not exhibit systematic synchronicity. Why is that? If the causal innocence thesis is correct, then if there is an explanation at all, it must be that we define causal concepts in terms of synchronicity. But no existing theory of causation does so, nor does this seem an especially promising direction for future attempts.

Explaining synchronicity thus seems to require a principle linking causation and time; e.g., the principle that identical causal processes always unfold during the same time interval. Neither the causal theories of time we have been looking at so far, nor theories based on the causal innocence thesis, provide us with the right kind of link. Of course, it is possible to just put synchronicity down as an unexplainable brute fact – but isn't that too high a price to pay for holding on to a metaphysical hypothesis?

## 4.3 A Kantian Approach

### 4.3.1 Introduction

The upshot of the argument from synchronicity is that neither a reduction of the causal to the temporal nor a reduction of the temporal to the causal has succeeded, since none of the extant theories is able to provide us with resources needed to explain the phenomenon of synchronicity. But even apart from the possibility that future theories might provide such resources, this argument can only indirectly imply that temporal and causal concepts are mutually dependent. To defend the mutual dependence thesis, as I wish to do, it is better to develop a direct argument to the effect that we cannot conceive of the world as temporal without also conceiving of it as causal and vice versa. We can find an argument with that conclusion in Kant's *Critique of Pure Reason*. Although the argument runs through several sections of the work (including the transcendental aesthetic, the transcendental deduction, and the schematism chapter), its core is to be found in the section on the second analogy of experience.

I first give Kant's argument (§4.3.2) and then show that it depends essentially on an acceptance of his transcendental idealism (§4.3.3). I am not convinced that transcendental idealism is wrong, but it certainly has few adherents. This means that Kant's argument won't be able to convince many people. For that reason, I develop an epistemological variation on the argument, which shows that one must accept either the mutual dependence thesis or radical temporal scepticism (§4.3.4). Since we don't want to be sceptics, we have to embrace mutual dependence.



### 4.3.2 Kant's Argument in the Second Analogy

Let us use the term 'apprehension' for a conscious mental image. Then our conscious sensation consists in the temporal succession of apprehensions. But this does not mean that we are apprehending an objective time order. On the contrary, the temporal succession of apprehensions is compatible with three completely different types of situation. First, there are situations like a dream, in which my apprehensions do not correspond to an object at all. I may, to use a perhaps rather worrisome example from my own experience, dream that Immanuel Kant is showing me his Star Trek museum while wearing a Vulcan mask, but then neither Kant nor the mask is grasped as a real object. Second, there are situations like looking over a house. One's gaze may first dwell on the door, then quickly sweep up to take in the roof, so that the apprehension of the door comes before the apprehension of the roof. Here there is indeed an object, but the order of the apprehensions does not correspond to an order in that object: the door does not come 'before' the roof in any objective sense. As Kant tells us:

In the previous example of a house my perceptions could have begun at its rooftop and ended at the ground, but could also have begun below and ended above; likewise I could have apprehended the manifold of empirical intuition from the right or from the left. In the series of these perceptions there was therefore no determinate order that made it necessary when I had to begin in the apprehension in order to combine the manifold empirically.  
(B237–238/A192–193)

Third, there is the situation where the order of the apprehensions *does* correspond to an order in the object that is apprehended. Kant's example (B237/A192) is that of a ship sailing downriver; perhaps even clearer is that of an apple falling from a tree to the ground. That the apple was first hanging in the tree, then halfway to the ground, and then on the ground, slightly the worse for wear, is not some merely subjective order of my apprehensions; it is the objective temporal order of the apple's states.

In the second analogy, Kant is concerned with understanding the difference between this third situation and the other situations. Thus, he is concerned with objective time order, that is, quite literally, with the temporal ordering of (states of) objects. But his question is not, or at least not primarily, how we can *know* that a certain succession that we apprehend is an objective succession. Rather, his question is how we can *think* an objective succession. What is it that allows us to conceive of the subjective temporal succession of apprehensions as corresponding, at least in some cases, to a temporal succession of the objects? In other words, how can we think of ourselves as being in a changing world, rather than just think of ourselves as changing?

Here is yet another way to approach Kant's admittedly abstract question. What is the difference between a subjective and an objective order? Answering this question requires us to make a distinction between our subjective mental states and the objects that are apprehended in those mental states. Kant tells us that:

appearance, in contradistinction to the representations of apprehension, can thereby only be represented as the object that is distinct from them if it stands under a rule that distinguishes it from every other apprehension, and makes one way of combining the manifold necessary. That in the appearance which contains the condition of this necessary rule of apprehension is the object. (B236/A191)

While unpacking all of this would take us too far, the core idea is clear. An object is that which makes it necessary for our apprehensions to be combined according to a certain rule – it is what ensures that our stream of mental images cannot change in any which way (as it *can* in a dream, where we are not apprehending objects). So to conceive of objects as undergoing change, that is, to conceive of an objective temporal order of events, *is* to conceive of them as necessitating a certain temporal order of apprehensions.

If we investigate what new characteristic is given to our representations by the *relation to an object*, and what is the dignity that they thereby receive, we find that it does nothing beyond making the combination of representations necessary in a certain way, and subjecting them to a rule; and conversely that objective significance is conferred on our representations only insofar as a certain order in their temporal relation is necessary. (B242–243/A197)

Now, the idea of a necessary temporal relation, that is, of one thing having to follow from something else, is for Kant simply the idea of *causality*; and any particular such rule is a causal law. Hence, the upshot of Kant's argument is that thinking objects as placed in time, that is, thinking an objective temporal order, *is* thinking objects under the concept of causality –, in other words, thinking the world as being subject to causal laws. A temporal world that is not also a causal world is simply inconceivable.

### 4.3.3 Assessment of Kant's Argument

What makes Kant's argument of particular interest for a defence of the mutual dependence thesis is that it delivers a conclusion at exactly the right level, namely, that of concepts and conceivability. If Kant is right, then we really can't think objective temporality without thinking causation. Even a causal theory of time doesn't deliver a result that is quite that strong, since it leaves open the possibility (emphasised by van Fraassen) that one could also give a temporal theory of causation, and hence doesn't absolutely prove that one

cannot have temporal concepts without having causal concepts. So *if* Kant delivers, he delivers exactly what we want.

Does Kant deliver? Of course, everything depends on whether we embrace his conception of objects. One could deny that there *are* objects in the Kantian sense. This would be equivalent to denying that there is any kind of necessity in our apprehensions: we just have one conscious image and then another, with the distinction between Kant's three scenarios – the dream, the house and the ship – having no real significance. Of course, this immediately shows how extreme such a denial would be. (Though arguably there are traces of it in Hume's *Treatise*.)

Of more interest is another way of going against Kant's conception of objectivity. It is crucial to Kant's argument that objects not only make certain combinations of representations necessary, but that this fact *exhausts* our grasp of what it is to be an object.<sup>11</sup> If we have an independent grasp of objects, the argument of the second analogy doesn't work, since there could then be a way of thinking objects as temporal that does not involve their connection by relations of necessity. In particular, it seems that we could just define objects as spatiotemporally located entities. If we did, Kant's questions wouldn't even make sense.

This becomes clear when we look at Peter Strawson's (Strawson, 1966, section 2.III.4) reconstruction of Kant's argument. Strawson believes that Kant's main line of argumentation can be separated from his metaphysics of transcendental idealism, a core part of which is precisely the conception of objects we are discussing. So Strawson, keen on giving an interpretation of Kant's argument in terms of a straightforwardly realist conception of objects as things located in space and time, tells us that Kant argues as follows. (1) There are objective temporal relations between events. (2) The causal processes of sensation generally don't switch these around; that is, if *a* happens before *b*, then we generally sense *a* before *b*. (3) The objective temporal relation between events necessitates a certain order of sensations. (4) So, we sense a necessary temporal relation between events. Unfortunately, the step from (3) to (4) is clearly invalid. Indeed, Strawson calls this 'a *non sequitur* of numbing grossness' (Strawson, 1966, 137).

Strawson's failure to find an even mildly plausible argument in Kant's second analogy lends strong support to the idea that Kant's argument about causation and time is dependent on his philosophy of transcendental idealism and on his conception of objects in particular. (This conclusion agrees with

<sup>11</sup> '[T]he relation to an object ... *does nothing beyond* making the combination of representations necessary in a certain way' (B242–243/A197), emphasis changed.

the analysis by Allison (2004).) Now, a lot can perhaps be said in favour of transcendental idealism, but it cannot be denied that it is very much a minority position in philosophy and even less popular among scientists. So for our current purposes, Kant's argument is of limited use. However, in the next section I will develop a variation on his argument that does not rely on transcendental idealism.

#### 4.3.4 The Argument from Scepticism

We start at a less exalted level than Kant: rather than asking how we can conceive of objects and events as temporal, we will assume that we already do so. Instead, we will ask the epistemological question of how we can *know* about the temporal order of events. Later on, we will have to return to the level of concepts.

How then can we know the temporal order of events? I remember sitting in the hospital, reading in Peter Strawson's book on Kant in preparation for writing a few sentences about it; I also remember sitting behind my desk and actually writing about the book. I also know that the reading happened before the writing. But how do I know this?

As Kant tells us, 'time cannot be perceived in itself' (B233). When we experience the world, we are not experiencing it *as* at a certain time. The experience of, say, seeing a blue square in 2022 just is the experience of seeing a blue square. It is indistinguishable from the experience of seeing a blue square at any other time. So, I cannot know based on nothing but the phenomenal contents of two remembered experiences which of them happened earlier and which later.

Hume claimed in the *Treatise* that memory is distinguished from perception by being 'fainter'. Now, if older memories were 'fainter' than newer memories, or if there were some other phenomenal difference imposed upon them, then this could perhaps be used as a clue to temporal order. But this doesn't seem to be the case. Very old memories can be painfully sharp and detailed, whereas I may have only the barest recollection of yesterday's breakfast. Of course, this is not to deny that the passage of time has an effect on our memories – I would have trouble remembering the names of all my classmates back in school, whereas I used to know them very well – but it *is* to deny that we can simply compare the contents of our memories and see which of them is earlier.

What about the fact that our memories do not need to be of instantaneous states, but can be of longer intervals? I remember hearing the famous ta-ta-ta-tam 'knock of fate' motif at the beginning of Beethoven's fifth symphony. This remembrance does not consist of four separate memories about which I can

legitimately wonder how their contents were related in time. I just remember hearing the motif, as such, in its temporal extension. Memories, then, can be extended in time. Does this provide us a purchase on time order? It surely does. But only of a limited nature: at the very least such memories require continuous consciousness, which means that we could never use this mechanism to find out about the time order of things that are separated by episodes of sleep. (In practice, the time that can be covered by a single episodic memory is of course much shorter. I do not have a continuous remembrance of my entire day.)

Here is another possibility: my memories may have been timestamped by the use of external clocks. Perhaps I looked at my watch in the hospital and saw it was the 8th of May, and then again behind my desk while writing about Strawson, where I saw it was the 9th. Surely this would establish temporal order? Well – as a matter of fact, I *didn't* look at my watch, but I nevertheless know the temporal order of the events. But even if I *had* looked at my watch, this would only serve to relocate the problem. For now I would have to explain how I can know that the event of my watch showing the 8th of May was temporally before the event of my watch showing the 9th. This may seem trivial, but it is not. I will return to the issue shortly.

So how *do* I know that I read Strawson's book before writing about it? In fact, there isn't much of a mystery here: reading the book was necessary for getting the knowledge needed to write about it. The reading was a causal condition of possibility of the writing. And causes always precede their effects. So the reading must have come before the writing. My certainty about the temporal order of events is based on my prior acceptance of two things: a particular causal link between these two events and the principle that causes happen before their effects.<sup>12</sup>

<sup>12</sup> Psychological research into assigning times to memories supports the idea that we arrive at such assignments through complicated processes of reasoning about how events hang together. Already in a 1993 review of theories about memory and time, William Friedman (1993, 58) concluded:

[T]ime information is stored not alongside memories for events but in our more general body of knowledge about time patterns. We know a tremendous amount about the temporal structure of our lives and the physical and social environment, from the daily cycle to the major events in a lifetime, and can rapidly extract the general temporal properties of some new experience [...] This general time knowledge allows us to interpret contextual information that happens to be associated with a memory. [...] Our ability to judge the time of events is not based on special time codes assigned to memories, an inherent temporal organization of memory, or age-of-memory codes created by the passage of time. Instead, the elementary information is the ordinary contents of memory.

While Friedman does not specifically mention causal considerations, it seems clear that the body of information he mentions will essentially involve causal knowledge.

This is not a peculiar feature of this example about memories, but true about temporal judgements in general. When we are asked to judge when a particular event took place, we always use causal information. When was this bronze knife made? Certainly after the invention of bronze, since inventing a technology is a causal condition for using it. When did this animal die? It ate plants; the plants got their carbon out of the atmosphere; and the radioactive carbon has since been spontaneously decaying and has not been replenished through any process. Given the truth of all these causal hypotheses, we can use carbon dating techniques to estimate the age of the bones. How long ago was it that I looked at my watch and saw it was the 8th of May? My watch tells me it is now the 10th of May, so assuming that its internal causal processes have been unfolding as intended by the maker, that must have been two days ago. When was the event that I am currently perceiving? Right now, since the causal processes of perception take very little time – unless the event is very far away, as in the case of celestial objects.

We could multiply examples indefinitely, but let us simply draw the conclusion that our temporal judgements are generally made with the use of causal presuppositions. The crucial presupposition is of course that causes happen before their effects. If we do not invoke that hypothesis, there is no basis left for making temporal judgements. Events could be located anywhere in time. We would have to be complete temporal sceptics.

Let us illustrate this by considering a wild hypothesis: that all the events that we label as the nineteenth century happen two thousand years in the future. The eighteenth and the twentieth century are contiguous in time – so 1718 is not 300, but only 200 years ago – while the nineteenth century simply lies two thousand years further down the line. We surely know that this hypothesis is wrong. But how do we know it? Again, as Kant said, we cannot perceive time itself. What we *can* do is tell an extremely plausible story about how events in the eighteenth century caused events in the nineteenth century and how events in the nineteenth century caused events in the twentieth century. We then use the principle that causes happen before their effects to dismiss the wild hypothesis. But without that principle, we couldn't dismiss it. A world that contains the exact same events as ours, only distributed differently through time, is empirically indistinguishable from ours. (We ignore the potential metaphysical impossibility of such a world.) The only way to avoid temporal scepticism is to accept a priori the principle that causes come before their effects.<sup>13</sup>

<sup>13</sup> At least in general. A mild amount of retrocausation could still be accepted without making temporal scepticism unavoidable.

By itself, the necessity of accepting this principle a priori doesn't disprove the causal innocence thesis. Hume, for instance, would immediately accept the principle, since it follows trivially from his definition of cause. But the fact that we need to accept this principle in order to stave off temporal scepticism *is* a problem for the causal innocence thesis. If causal concepts reduce to temporal concepts, then it cannot be the case that causal knowledge is a guide to temporal knowledge. It cannot be the case that before I define causation, I am in the dark about temporal relations; but then, merely by introducing a definition, I suddenly have access to knowledge about the world. So, given that temporal scepticism is false, so is the causal innocence thesis. And this is what we set out to prove.

A brief coda. What is the status of the principle that causes always come before their effects? As we have seen, temporal scepticism can only be staved off if we accept this principle a priori. Received wisdom has it that all statements that can be known a priori must be analytic. But we have rejected the causal theories of time, according to which the principle is analytic because of the definition of time; and we have also rejected the causal innocence thesis, according to which the principle is analytic because of the definition of causation. But this seems to exhaust the options for analyticity. So there is a case to be made that the principle is synthetic a priori – which would mean that Kant has, after all, the last laugh in this discussion.

## 4.4 Conclusion

Let us recap the argument. There are four possible positions with respect to the relation between temporal and causal concepts. There could be no relation; temporal concepts could be reduced to causal concepts; causal concepts could be reduced to temporal concepts; or these two sets of concepts could be mutually dependent. The aim of the chapter has been to defend the fourth possibility.

Two arguments were developed. First, the argument from synchronicity, which shows that the reduction of temporal to causal concepts is unsuccessful. It also narrowed the range of possible reductions in the other direction, but couldn't quite exclude them. Second, the argument from scepticism, which shows that temporal concepts are dependent on causal concepts. Together, these two arguments exclude all options except for the mutual dependence thesis. I therefore conclude that we cannot think time without thinking causation and that we cannot think causation without thinking time. Having temporal concepts and having causal concepts are two sides of the same coin.

As I explained in the introduction, contemporary theories of causation almost universally assume that we can understand time before we understand causation. If this is false, we need to rethink our approach to causation. Perhaps it is time to leave the Humean phase of analytic philosophy behind and enter the Kantian phase.