Time and Causality

by

Russell Davidson

Department of Economics and CIREQ
McGill University
Montréal, Québec, Canada
H3A 2T7

AMSE-GREQAM
Centre de la Vieille Charité
2 Rue de la Charité
13236 Marseille cedex 02, France

russell.davidson@univ-amu.fr

Abstract

The understanding of causal chains and mechanisms is an essential part of any scientific activity that aims at better explanation of its subject matter, and better understanding of it. While any account of causality requires that a cause should precede its effect, accounts of causality in physics are complicated by the fact that the role of time in current theoretical physics has evolved very substantially throughout the twentieth century. In this article, I review the status of time and causality in physics, both the classical physics of the nineteenth century, and modern physics based on relativity and quantum mechanics. I then move on to econometrics, with some mention of statistics more generally, and emphasise the role of models in making sense of causal notions, and their place in scientific explanation.

Key words: Causality, time, quantum mechanics, relativity, econometrics

JEL codes: Z00

This research was supported by the Canada Research Chair program (Chair in Economics, McGill University), and by grants from the Social Sciences and Humanities Research Council of Canada and the Fonds Québécois de Recherche sur la Société et la Culture. Earlier versions of this paper were presented in April 2010 in the context of an ongoing series of seminars on Causality organised by Alain Trannoy of IDEP in Marseille, and at the Doctoriales of January 2012, also held in Marseille. My thanks go to Pierre-Henri Bono, who gave me valuable feedback on the earlier versions, and engaged me in a number of thought-provoking discussions.

May 2012

1. Introduction

The notion of Cause, or Causality, has given rise to much discussion in the philosophical literature, a little discussion in theoretical physics, and considerable application in econometrics. In this article, I try to present some of the ideas that have been put forth in these very different literatures, to try to find some of these ideas that are common across disciplines, and to explore their repercussions for econometrics.

In the next section, I lay out some general principles relating to causality, in particular making the distinction between causal necessity and causal sufficiency. I also mention some problems posed by the assumption of determinism, or of non-determinism, in causal reasoning. In section 3, I review the evolving notions about the mystery of time, as it was viewed in classical physics as it was developed up to the beginning of the twentieth century. Then, in section 4, I move on to the unexpected and still counter-intuitive insights of twentieth-century physics, coming from special and general relativity and quantum mechanics. In section 5, I look at the concept of Granger causality in econometrics, and mention some present-day controversies about the role that causality should play in economics and science generally.

2. Causality

There are various, often incompatible, definitions of causality to be found in dictionaries, philosophical works, and other places. Here I try to form definitions that are as simple and general as possible. Consider two events, A and B. An intuitive definition of the proposition that A causes B is:

(i) A and B are real, or true;

(ii) If A is not real or true, then neither is B; and

(iii) A precedes B in time.

This definition raises a number of issues. What do we mean by an "event"? There are several admissible answers: an action, a fact of nature, among others. A fact is true or not, an action is performed (real) or not. Our tentative definition is general enough to allow for various different possibilities.

In order to steer clear of some trivial cases, we want to suppose that the events A and B are logically *independent*. Thus we don't want to say that the conclusion of a mathematical theorem is *caused* by the premisses of the theorem.

It is important to distinguish between causal **necessity** and causal **sufficiency**. Necessity means that:

not A (written as $\neg A$) implies $\neg B$.

In words, without A, there can be no B. Logically, the condition is equivalent to the condition that B implies A; that is, A is a *necessary* condition for B. This is our condition (ii).

Sufficiency means that:

A implies B, or $\neg B$ implies $\neg A$.

In words, every time that A holds, B holds as well, inevitably, that is, A is a *sufficient* condition for B.

Sufficiency is logically quite distinct from necessity. Necessity leaves open the possibility that A holds without B. Sufficiency leaves open the possibility that B holds without A.

Example

This example, which involves me and the graduate student with whom I discussed the material of this paper, is not based *completely* on real life!

A: I, Russell, tease Pierre-Henri.

B: Pierre-Henri gets angry with me.

For causal necessity, Pierre-Henri does not get angry unless I tease him, but I can sometimes tease him without his getting angry with me. For causal sufficiency, Pierre-Henri must get angry with me every time I tease him, but he sometimes gets angry with me with no provocation on my part.

In this example, we could imagine an empirical study of the assertion that A causes B, with either causal necessity or causal sufficiency. The investigator should observe the regular interaction between Pierre-Henri and me, noting the occasions on which A or B occurs. If one day I tease P-H (A) and he does not get angry ($\neg B$), then causal sufficiency can be rejected. If another day P-H gets angry (B) without being teased ($\neg A$), then causal necessity is rejected. This study is made possible by the fact that our interactions are **repeated**. Note also that, in keeping with the Popperian idea (see Popper 1972) that, in order for a statement to have empirical content, it must in principle be *falsifiable*. That is, we must be able to conceive of an experiment, or an observation, that would lead us to reject the statement.

But if A and B are not repeated, but rather unique, events, what sense can we make of the assertion that A caused B? I suppose here that condition (i) is satisfied, so that A and B both occurred. In order to make any sense of the statement about causality, we have to admit to our discussion *imaginary worlds* or even *universes*. We call such worlds or universes **counterfactual**. Without considering them, it is impossible to know what *might* have occurred if A did not, or if B did not occur.

But this remark leads to more questions than it answers, in the form of philosophical, or even physical, problems. What is the set of universes that these counterfactual universes inhabit? How can we delimit this set? Let's denote the set by X. Then we have a number of reasonable choices:

(a) X is the set of *logically* possible universes, that is, all universes that are not logically self-contradictory;

(b) X is the set of universes compatible with the laws of physics, as we know them;

(c) X is the set of logically and physically admissible universes that are sufficiently *similar* or *close* to the real world.

The last choice is no doubt the best, but, in order to implement it, what topology can we use to define a *neighbourhood* of the real world?

Determinism

A fictitious entity that is often referred to in discussions of determinism is **Laplace's demon**. Here is what Laplace had to say. It was later writers who ascribed the name of "demon" to his "intelligence".

> *Une intelligence qui, à un instant donné, connaîtrait toutes les forces dont la nature est animée, la position respective des êtres qui la composent, si d'ailleurs elle était assez vaste pour soumettre ces données à l'analyse, embrasserait dans la même formule les mouvements des plus grands corps de l'univers, et ceux du plus léger atome. Rien ne serait incertain pour elle, et l'avenir comme le passé seraient présents à ses yeux.*
>
> *An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes.*
>
> **Pierre Simon Laplace, A Philosophical Essay on Probabilities**

For Laplace's demon, and also for classical physics, the universe is perfectly deterministic. Everything is determined by the boundary, or initial, conditions, and by the laws of motion of classical physics, Newton's mechanics and Maxwell's electromagnetism. In fact, the deterministic approach is shared by Einstein's relativity, both special and general.

In the deterministic world view, there exists at each moment only one possible future. If this view reflects reality, then no causal relation between two events can make sense. Indeed, Bertrand Russell (1988), in an essay entitled *On the Notion of Cause*, says:

> *... the reason why physics has ceased to look for causes is that, in fact, there are no such things. The law of causality, I believe, like much that passed muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm.*

Some philosophers go so far as to claim that determinism excludes the possibility of free will, and consequently of moral responsibility. A famous defender of this view was Martin Luther, and it has been espoused more recently by, among others, the "hard incompatibilist" philosopher Derk Pereboom (2003), who however denies determinism rather than free will. But even in a deterministic universe, it is out of the question to have access to all of the information that Laplace's demon is supposed to possess. Thus by no means all future events are forecastable. Rather, we can, if not forecast, then at least *anticipate*, *imagine*, or even *fear*, futures different from the supposedly unique real future. We can perfectly well suppose that each of these imaginable futures obeys the laws of physics, and maybe even those of the social sciences, for instance, psychology.

Deterministic Chaos

It is well known that classical mechanics admits **chaotic** dynamics. What that means is that a change in the initial conditions, however small, can give rise to arbitrarily large divergences in the future evolution of the dynamical system. I endeavour now to formalise this notion. For a textbook treatment with much more detail, see Baker and Gollub (1996).

Let x_0 denote the initial state of a variable x. We suppose that the value of the variable evolves in time, in a purely deterministic fashion, according to some dynamical law. In continuous time, we would have a differential equation, $\dot{x} = f(x)$; in discrete time an iterative relation of the form $x_{t+1} = g(x_t)$. This dynamical system is called chaotic if, for all $\varepsilon > 0$ and for all $K > 0$, there exists a time T such that $|x(T; x_0) - x(T; x_0 + \varepsilon)| > K$ (in continuous time), or $|x_T(x_0) - x_T(x_0 + \varepsilon)| > K$ (in discrete time). This definition is simplified; proper mathematical definitions can be found in Aulbach and Kieninger (2001)

In other words, an infinitesimal change in the initial condition gives rise to arbitrarily large effects in the future. It is this aspect of things that gives rise to the well-known difficulties in weather forecasting. However, it is also what makes possible random number generators on the computer, which is nothing if not a deterministic device.

It should be clear that chaos authorises consideration of very different futures in our set X. In fact, prudence requires such! Specifically, before an event A occurs or does not occur, we are at liberty to consider futures with either A or $\neg A$, and B or $\neg B$. Then, by applying the laws of physics and any other relevant laws, we may be able to reject one or other type of causality by looking at what happens in our different futures.

We can see, then, that neither the possible determinism nor the possible indeterminism of the universe is relevant for the analysis of causality by a conscious observer, except for Laplace's demon, if it were to exist. The essential capability is to conceive of different futures. This is analogous to what historians do, making sense of causal relations by considering counterfactual histories.

The fact is that determinism, even if it is only approximate, *helps* us to analyse different futures. If anything were possible, on the other hand, it would be a waste of time to try to anticipate anything at all.

Quantum Indeterminacy

We know that the introduction of quantum mechanics near the beginning of the twentieth century gave rise to a random element in microscopic physics. "Random" means non-deterministic. Does this new randomness facilitate the analysis of causality, or hinder it? According to the physicist David Deutsch – see Deutsch (1997) – quantum indeterminacy allows us to make sense of causality in a completely different way from anything allowed by classical physics with its deterministic chaos.

It must be noted at once that the equations of motion of quantum mechanics involve no random element. They are just as deterministic as the equations of classical physics. According to Deutsch, and a handful of other physicists, the indeterminacy is the result of the existence of an infinite number of **parallel universes**. This notion is termed the **many-worlds interpretation** of quantum mechanics, and is due originally to Hugh Everett (1957). His interpretation was scorned by other physicists to the extent that he abandoned theoretical physics after obtaining his Ph.D. But since then the interpretation has had many advocates.

What is random in this interpretation is not the "multiverse" dynamics, but rather the branch of the multiverse to which the observer belongs. If we consider the full set of the universes that constitute the multiverse, then there will be a certain proportion of these in which A is true, and another proportion in which $\neg A$ is true. If in even a single one of these universes, we find that $\neg A \wedge B$ is true, where the symbol $\wedge$ for set intersection can be interpreted here as meaning "and", then we can reject causal necessity, and similarly finding $A \wedge \neg B$ leads to rejection of causal sufficiency. In the same way, we can define what we mean by the *probability* of an event as the proportion of universes in which it occurs. This is of course what is done whenever one does a quantum-mechanical calculation. Unfortunately, this interpretation of quantum mechanics does not allow the possibility of an experiment that could reject causality of either sort.

Non-Causality

In all scientific disciplines, progress comes from the result of an experiment, or an observation, that leads us to reject a hypothesis. It is therefore important to be able to demonstrate non-causality, that is, the absence of any relation of cause and effect between two events, or types of events. I will say more about this in the context of econometrics.

Here, though, are a couple of problems that have vexed many people.

(1) Two naughty boys throw stones in the direction of a glass bottle. The first stone thrown by the first boy smashes the bottle. A second later, the first stone thrown by the second boy passes right through the shards of glass where, a second earlier, the bottle was standing. For this little tale, we can establish every causal link, both for necessity and sufficiency.

(2) Person X makes a trek into the desert. One of his mortal enemies, person Y, secretly substitutes a deadly poison for the water in X's water container. A second enemy, Z, steals the container from X once the latter is in the desert, believing that he is stealing X's water. X dies of thirst. Who is guilty of his murder?

3. Time

quid est ergo tempus? si nemo ex me quaerat, scio; si quaerenti explicare velim, nescio.

What then is time? If no one asks me, I know: if I wish to explain it to one that asketh, I know not.

Saint Augustin, Confessions, XI, 14, 17

Tempus absolutum verum & Mathematicum, in se & natura sua absq; relatione ad externum quodvis, quabiliter fluit.

Absolute, true, and mathematical time, of itself, and from its own nature, flows equably without relation to anything external.

Isaac Newton, *Principia*

Time has been regarded as a mystery for long ages. In fact, there is no such thing as the flow of time – it is just an illusion. Deutsch (1997) explains this very clearly in Chapter 11 of his book. Julian Barbour (2008) goes a good deal further, and claims that "time has no role to play as an independent element of reality". He shows how to formulate many of the laws of physics with no explicit mention of time.

Time is one-dimensional, but it has to be taken in conjunction with three-dimensional space for us to obtain a coherent concept, namely four-dimensional **space-time**. In classical physics, including non-quantum relativity, space-time is a static block. An "instant" is a three-dimensional cross section of this block. A physical object, microscopic or macroscopic, follows a **trajectory** in space-time which is such that, at each instant, the object occupies a region of the three-dimensional section that constitutes that instant. Movement is the change in this region from one instant to a later one. But, even if an agent is incapable of perfectly forecasting the trajectory, the trajectory is what it is, and the agent can't do anything to change it. Thus we can make no sense of causality in this framework.

Quantum mechanics could possibly improve this sorry state of affairs, but the fact is that physics has not yet reached the point where it can satisfactorily integrate quantum mechanics and general relativity (our currently best theory of gravity). The difficulty is that no one has yet found out how to quantise space-time. Without that, in the context of the multiverse, things are just as deterministic as in the classical world-view.

Models

Science itself can be and has been regarded as an "object" in the real world, a sociological object probably. But the aim of science, according to its practitioners, is to *understand* the real world, whether it be a single universe or a multiverse. As humans, we don't have a good understanding of the way in which we can understand things, because we don't have a detailed understanding of the functioning of the human brain and of what we call consciousness. Despite this, we may observe that most sciences make use of **models** in order to arrive at understanding. Models can be good, bad, or indifferent, but what they do is to encapsulate what we understand of real-world phenomena.

Models can be thought of as **virtual reality**. This is evidently true when a model can be simulated on a computer, but it is fruitful to think of scientific models as virtual reality more generally. Now a given model may not conform to (real) reality. The laws of physics may not be respected, or the properties of human beings may be falsified, as in many economic models. But the very flexibility of models is what allows us to give causality its full force in the context of virtual realities. We can *calculate* effects, and thus say whether some virtual event causes another or not.

Deutsch (1997) has a fascinating discussion of virtual reality, and propounds what he calls the Turing principle:

> *It is possible to build a virtual-reality generator whose repertoire include every physically possible environment.*

As he says, this principle (assuming that we can accept it) is what makes reality comprehensible. In particular, the flexibility of virtual reality allows us to consider virtual worlds which are counterfactual to the real world, and so to determine experimentally whether one or other sort of causality can be rejected. Of course, most such experiments would require a virtual-reality technology far beyond anything currently available!

Some models are purely descriptive. A statistical model, for instance, might specify the probabilistic properties of a set of variables, and nothing more. But that may be enough for us to do forecasting, on the basis of the probabilistic structures of the model. Half a century ago, most physicists thought of quantum mechanics that way, as a mathematical recipe that could be used to predict experimental results. The "interpretations" of quantum mechanics that were then current were very counter-intuitive – recall that Everett's interpretation was treated with scorn, and still is, by most physicists.

However, this wilfully blind positivist approach is finally giving way to a thirst for **explanations** in physics. Yes indeed, physics gives better agreement with experimental data than any other discipline, but does it constitute a true **theory**? A theory must **explain**, by proposing a **mechanism**, or in other words a *causal* chain. Classical physics had no interest in any such thing. Its underlying model is purely descriptive, and explains nothing, even if it does allow us to compute the orbits of planets and to do "rocket science".

> *"It is wrong to think that the task of physics is to find out how nature is. Physics concerns what we can say about nature."*
>
> **Niels Bohr**

Most physicists would agree with Bohr's famous remark, although it is slowly becoming a more acceptable view to think that what we can say about nature should allow us to understand it.

4. The Relativistic Universe

Special relativity presents a considerable problem for our definitions of causality, because we want to require a cause to precede its effect. One of the less intuitive aspects of relativity is that it scraps the concept of **simultaneity**, as soon as we deal with two observers in relative motion. Suppose we have two events A and B – points in space-time. Then it is quite possible that the first observer finds that A precedes B, while the second observer finds that B precedes A.

Causality is as important a concept for physics as for the other sciences, and so it is not surprising to learn that a solution was quickly found. At each point in space-time, we can split up the rest of space-time into three non-intersecting subsets. The first is the set of those points such that the distance from our reference point to it is **spatial**. This means that it is impossible to go from one point to the other at a speed less than or equal to the speed of light. We can safely exclude the possibility of a causal link between two spatially separated points.

The second subset is made up of those points for which there does exist a trajectory from that point to the reference point which does not exceed the speed of light. These points are the "past" of the reference point, and may as such cause the reference point. The last subset is the "future"; those points that the reference point may cause. The three subsets are separated by the trajectories of photons (that is, light) which pass through the reference point.

Another problem has proved much more difficult to resolve. It has to do with the interaction between quantum mechanics and special relativity. After the invention of modern quantum mechanics, in the 1920s, theory indicated the existence of quantum states called **entangled**. One speaks of **quantum entanglement**.

To grasp this phenomenon, suppose that an atomic reaction creates two photons, at one and the same point in space-time, which go off in diametrically opposite directions, and with polarisations that are mutually orthogonal. The quantum uncertainty principle does not let us know in advance which one of the two photons will pass through a filter that blocks horizontally polarised photons. If one of them is observed and does pass, then it is definitively in a state of vertical polarisation, which implies that the other photon is equally definitively horizontally polarised. The two photons move in opposite directions at the speed of light, and so are spatially separated. But it appears that submitting one of them to a measurement by our polarising filter influences any possible measurement that may be performed on the polarisation state of the other one. This effect has been seen unambiguously in various experiments. The seeming paradox is associated with the names of Einstein, Podolsky, and Rosen (EPR), (1935).

The conclusion drawn by EPR in their paper is that quantum theory is incomplete. In order to avoid having to suppose that a message is transmitted instantaneously between the two photons, we must suppose (they claim) that the necessary information was there from the start, when the photon pair was created. This is a **hidden variable** theory. But the physicist John Bell (1964) produced a quite simple statistical argument that showed that it is impossible to reproduce the correlations predicted by quantum theory (and confirmed experimentally; see Aspect, Dalibard and Roger 1982 and Aspect, Grangier, and Roger 1982) by any number of hidden (or latent, in econometric parlance) variables. The usual conclusion drawn from the "Bell inequalities" is that quantum mechanics is intrinsically non-local. This conclusion is counter-intuitive, but then so are many other aspects of quantum mechanics, and it is accepted by many physicists, who believe that the paradox is real, and that there is a real incompatibility between special relativity and quantum mechanics.

This view was expounded recently in an article that appeared in the March 2009 issue of *Scientific American*. When I read this article, I was shocked, because I had been sure that, although there were problems with quantum mechanics and quantum gravity – and so with general relativity – there were no such problems if one is prepared to stick with special relativity. According to the analysis in Smerlak and Rovelli (2007), the EPR paradox is due to two assumptions, one the locality principle:

> *Relative to a given observer, two spatially separated events cannot have instantaneous mutual influence.*

the second "Einstein's realism":

> *There exists a physical reality independent of substantiation and perception.*

In order to get rid of the paradox, Smerlak and Rovelli use the framework of **relational quantum mechanics** – Rovelli (1996) – to do away with Einstein's realism, while maintaining a coherent theory of quantum mechanics. The main postulate of relational quantum mechanics is:

> *In quantum mechanics different observers may give different accounts of the same sequence of events.*

There is an interesting theological consequence of this postulate, namely that omniscience is impossible. Be that as it may, Rovelli maintains that

> *Quantum mechanics is a theory about the physical description of physical systems relative to other systems, and this is a* complete *description of the world.*

It is also a description that permits making sense of causal relations. But these relations are always only relative to some observer. However, the formalism of quantum mechanics rules out any possibility of contradiction in those cases in which two observers can compare notes after making their own observations on entangled systems.

5. Econometrics

Econometrics is supposed to be a scientific enterprise. Consequently, it rests on the use of models. Unlike those of classical physics, econometric models explicitly introduce random elements. What can we mean by "random" in this context?

The answer I normally give to this question – in teaching econometrics, I ask the question, and then give my answer – is that the random elements of an econometric model represent everything that we don't want to model explicitly. This might be because we *can't* model some things explicitly, on account of missing data, for instance. More commonly, it is because the phenomenon under study – in a moment I'll consider the example of household consumption – is subject to determinants that we can't quantify and that are probably of no interest to an econometrician whose goal is to understand economic mechanisms. We limit ourselves to a partial explanation, even at the cost of not being able to make reliable forecasts.

Econometrics, like all other statistical disciplines, relies on the notion of **repeated observations** of the same phenomenon. We saw earlier that repetition of a phenomenon in similar, but not identical, circumstances allows us to introduce the concept of causality in a coherent manner. In this way, we escape almost all the philosophical traps I have mentioned so far.

At the same time, inevitably, we introduce other complications. Most of the time, we deal with continuous variables, which means that the event B (the effect) must be replaced by a quantitative measure of one or more variables. Similarly for the cause, A. This makes it simpler to define what we mean by **non-causality**. We say that a variable X does not cause another variable Y if the earlier values of X have no influence on the later values of Y. This approach is called **Granger causality**, having been introduced by Granger (1969). A similar related approach is due to Sims (1972).

The random elements in the model allow us to introduce the **neighbourhood** of circumstances (worlds, universes) that surround the observed trajectories of X and Y. We no longer need to invent imaginary trajectories that *might* have existed in the real world, nor need we invoke the parallel universes of (one version of) quantum mechanics. It is enough to vary the realisations of the random elements in order to create, within a virtual reality, all the relevant circumstances needed to reject causal sufficiency.

Why not causal necessity?

We said that A is a necessary cause of B if B implies A. In propositional logic,

$$B \Rightarrow A \qquad \Leftrightarrow \qquad A \vee \neg B,$$

of which the negation is

$$\neg(A \vee \neg B) \qquad \Leftrightarrow \qquad \neg A \wedge B.$$

The set union symbol $\vee$ here means "or". This is *not* "exclusive or", so that, if $A \vee \neg B$ is true, it is not ruled out that *both* A and $\neg B$ are true. Causal necessity is rejected if

B occurs in the absence of A. Translating this into econometric terminology, the variable Y (associated with the event B) varies without any variation of X (associated with A). But normally Y has several determinants, which implies that Y can perfectly well vary with no change in the value of X. This is enough to reject causal necessity.

In economic theory, we often say *ceteris paribus*, other things being equal. In virtual reality, we can arrange things so that all the other variables, and also the realisations of the random elements, do not vary. This is a way to restrict the set of circumstances we consider for the purpose of establishing the existence or the non-existence of a causal link. However, if nothing but the two variables X and Y can move, then there is a deterministic functional relation between the two variables. In that case, we would never be able to reject causal necessity. Conclusion: it is causal sufficiency that is the useful concept in econometrics.

Causal sufficiency

Causal sufficiency is the proposition that A implies B. Propositional logic tells us that the negation of this proposition is $\neg B \wedge A$. Translating this, we see that this means that X varies without producing the effect of a variation of Y. Once again, then, we can reject causal sufficiency if, *ceteris paribus*, Y takes on the same value whatever the value of X. This would mean that the deterministic relation between the two variables introduced by the *ceteris paribus* assumption admits one and only one value for Y.

This has finally led us to a testable proposition. The null hypothesis specifies a no doubt complicated relation among the full set of variables considered relevant for the model, along with a set of random elements, or disturbances, or shocks, as we usually call them (but not "error terms", please!) With this specification, the null hypothesis of non-causality requires that, for any configuration of the variables other than X and Y, and for any realisation of the random elements, the value of Y is uniquely determined, whatever the value of X. The alternative hypothesis allows the value of X to have an influence on that of Y.

So far, time has been excluded from the discussion. But, in order for there to be a causal link between X and Y, the variation of X must *precede* that of Y. In econometric parlance, it isn't the current value of X that has an effect on Y, but rather the past (or lagged) values of X. Thus a model with a causal element must be a **dynamic** model.

In economics, as also in ecology, our models often contain no **exogenous** variables. In a probabilistic model, an exogenous variable is treated as though it were deterministic. In other words, everything is done *conditionally* on the exogenous variable. With no exogenous variables, all the variables in the model are consequently **endogenous**, and so it can be the case that everything depends on everything else.

Vector Auto-Regression (VAR)

In econometrics, the most frequently encountered models of this sort are **VAR** models, where VAR = vector autoregression. In a model of this type, the current values of a set of endogenous variables are determined by the lagged values of the same set of variables and by the realisations of a set of random elements. In the current state of the art, one almost always postulates a linear relation among the variables. Now I will give an illustrative example that makes use of the household consumption function.

Here is the model:

$$c_t = \alpha_1 + \beta_{11} c_{t-1} + \beta_{12} y_{t-1} + u_{t1},$$
$$y_t = \alpha_2 + \beta_{21} c_{t-1} + \beta_{22} y_{t-1} + u_{t2}.$$

The two variables are c, household consumption, or, more likely, the logarithm of household consumption, and y, disposable income of households, or its log. This is a macroeconomic relation. The variables c and y are aggregate variables, and they represent **flows**. The time index t refers to a period of some given duration, typically a year, a quarter, or a month. The random elements u_{t1} and u_{t2} are realisations of a bivariate distribution with zero expectation. We may wish to suppose that the pair (u_{t1}, u_{t2}) is independent of all other pairs (u_{s1}, u_{s2}), with $s \neq t$. The quantities denoted by α_i, β_{ij}, $i, j = 1, 2$, are the model **parameters**, which are treated as deterministic constants.

If the parameter values are known, along with the bivariate distribution of the random elements, we can undertake a stochastic simulation – virtual reality – if we have the initial condition (c_1, y_1). This amounts to specifying a **data-generating process**, or **DGP**, and that is enough for us to be able to study all the statistical properties of the generated variables.

According to elementary macroeconomic theory, disposable income causes consumption. Normally, except for some sophisticated models, we don't imagine that consumption causes income. This hypothesis corresponds formally to the hypothesis that $\beta_{21} = 0$, and we have several ways in which we could test this hypothesis. We would refer to this hypothesis as that of Granger non-causality.

Just to amuse myself, I estimated the model on Canadian quarterly data. I could not reject the hypothesis that c_{t-1} does not cause y_t. What amused me was that neither could I reject the hypothesis that y_{t-1} does not cause c_t. Of course, this is a woefully inadequate model, and my amusing result just shows that you can get almost any false conclusion from an ill-specified model.

It can be seen that, with a bivariate VAR, two variables can cause each other. But with this toy model, it seems that there is no causality at all! With more lags of the variables, the dynamic structure would let us estimate the causal delays. We could get this from a model with **distributed lags**.

Econometricians versus Statisticians

Although the concept of Granger causality and its implementation by means of VAR models allows us to formulate hypotheses of non-causality, and possibly to reject them, it does not, or not always, satisfy our desire for **understanding** and **explaining** economic mechanisms by means of causal chains. For that, it is necessary to base the models used for estimation and inference on **economic theory**. Granger's approach has come in for little criticism on this ground, because Granger always maintained that the goal of his methodology is to help economic forecasting. To the extent that this goal is attained, the methodology must be justifiable on some level.

But econometricians have always had a preference for **structural models**, in which the relation between the formal model and the underlying economic theory is clear. This preference led to the seminal work of the Cowles Commission. This research institute, founded by Alfred Cowles in 1932, developed an approach to econometrics called the **simultaneous equations approach**. Although the work of the economists who worked with the commission came in at the time for a good deal of criticism from other economists, their work is nowadays regarded as one of the cornerstones of modern econometrics. Since everything may depend on everything else, the Cowles econometricians developed models with several (simultaneous) equations that specified relations among the endogenous variables of the model. A good contemporary account of this work is found in Bennion (1952).

Historically, this led to models that were often thoroughly incompatible with the data, and the realisation of this led to a great many advances in econometric theory. Not enough for everyone, so that Sims (1980) felt moved to write an influential paper in which he points out the still great disparity between results obtained from the macroeconometric models current in 1980 and those obtained from more purely statistical models. Econometricians became more concerned with testing the statistical reliability of their models, and less concerned with the relation of these models to economic theory.

It is probably fair to say that structural models returned as the main focus of interest of many econometricians with the advent of the twenty-first century. Problems associated with the identification of such models and of their parameters assumed considerable importance, and stimulated much work intended to elucidate their nature and ways of solving them. Whether a model is structural or not makes little difference to how we can perform inference about causal sufficiency. The essential element is to be able to set up counterfactual situations by means of the model, as described above.

Statisticians have applied the term "causality" in more senses than that which we have considered here. Three of these are reviewed in Cox and Wermuth (2001). Their third, which they call "Causality as Explanation of a Process" is the one we have considered as fundamental. Their other two are "Causality as Stable Association" and "Causality as the Effect of Intervention". Depending on which of these interpretations of causality is adopted, different constraints arise as to what can be a cause, or, for that matter, what can be an effect.

For instance, in biostatistics and medicine, emphasis is often put on **randomised trials**, in which two groups of subjects are treated differently. One usually speaks of a control group, the members of which are not **treated**, and a treatment group, for which a particular treatment is prescribed. After some definite period, the members of both groups are examined for some particular property, which is thought of as the effect of being treated or not. Clearly, the idea is to be able to see whether the treatment causes the effect, and, perhaps, to reject the hypothesis that it does so. Here, if one can select the members of the two groups quite randomly, in a way totally unrelated to the treatment or the effect, then the distribution of effects within each group serves as the counterfactual distribution for the other. Here, causality is plainly thought of as the effect of an intervention, namely the treatment.

Even in medicine, a truly randomised trial can be difficult to achieve, for both practical and ethical reasons. In econometrics, it is even more difficult, although not completely impossible. However, “natural experiments” can arise for which an econometrician may be able to identify two groups that are “treated” differently, perhaps by being subject to some government programme, and to measure some effect, such as wages, that might be affected by the treatment. This can be fruitful, but, naturally enough, it requires the use of sophisticated statistical and econometric techniques.

In a polemical essay, Heckman (2001) maintains that econometrics has suffered as a result of too great an application of the methodology of mathematical statistics. He says that

> *Statistics is strong in producing sampling theorems and in devising ways to describe data. But the field is not rooted in science, or in formal causal models of phenomena, and models of behavior of the sort that are central to economics are not a part of that field and are alien to most statisticians (see, for example, Holland* 1986).

This is a strong statement of what I have called the preference of econometricians for structural models. Indeed, Holland, in the article cited by Heckman, distinguishes between what he calls the *scientific* and the *statistical* solutions to the **Fundamental Problem of Causal Inference**. The use of the term “fundamental” here is questionable, as the problem is stated firmly in the context of the measurement of treatment effects. The problem is that, for a given member of one of the two groups, control or treated, one observes only one effect, and so we cannot observe the difference between the effect with treatment and that without. Holland’s “scientific” solution to the problem is stated in language that would be confusing to quote here, but it comes down to the use of a structural model based on theory. The “statistical” solution, on the other hand, limits itself to estimating *average* effects, of one sort or another.

There has been considerable recent debate in the econometrics literature between supporters of views that are inspired more by one of the two solutions that Holland mentions than the other. I draw attention to two papers that have contributed to the debate, Angrist and Imbens (1999), and Heckman (1999) although there have been others since those two. A recent paper in which the use of causality in econometrics is extensively discussed is Heckman (2008). It is beyond the scope of this article to go into the issues dealt with there.

6. Concluding Remarks

Causality is a very general notion, so much so that it is necessary to be more specific about what sort of causality one means in any given context. However, one common element is the need to be able to construct counterfactual scenarios, whether it be for physical mechanics or econometrics, or indeed any of the other disciplines where causal mechanisms are used for the purposes of explanation and understanding.

Like most things, causality is relative. Holland (1986) makes this point explicitly: if we say "A causes B", we mean that A causes B *relative* to some other cause that includes the condition "not A". If we interpret causality as the effect of intervention, then the definition of the other cause must hold constant, in the sense of *ceteris paribus*, everything that could not be changed by intervention. Thus some researchers deny that a person's race or sex could be the cause of an effect, since no (non-surgical!) intervention could change either of those attributes. In an econometric model, on the other hand, we often want to see the effect of race or sex on a dependent variable, such as earnings. The reason for Russell's (1988) denial of any meaning for causality in classical physics is precisely that, in a deterministic universe, only one trajectory is possible, and so no counterfactual world or universe can be constructed.

In physics, a proper interpretation of causality remains a subject of debate, although few physicists would want to deny its importance. But in other experimental or empirical disciplines, the concepts of models and of virtual reality make it possible to make meaningful, and potentially falsifiable, causal assertions, and thus improve our understanding. Where scientific debate can legitimately arise is in circumscribing the relevant relativity, and choosing which things to vary in a counterfactual situation and which things to hold constant.

References

Angrist, J. D. and G. W. Imbens (1999). "Comment on James J. Heckman, 'Instrumental Variables: a Study of Implicit Behavioral Assumptions Used in Making Program Evaluations' ", *The Journal of Human Resources*, 34, 823-7.

Aspect A., J. Dalibard, and G. Roger (1982). "Experimental Test of Bell's Inequalities Using Time-Varying Analyzers", *Physical Review Letters* 49, 1804-7.

Aspect A., P. Grangier, and G. Roger (1982), "Experimental Realization of Einstein-Podolsky-Rosen *Gedanken-experiment*: A New Violation of Bell's Inequalities", *Physical Review Letters*, 49, 91-4

Aulbach, B. and B. Kieninger (2001). "On Three Definitions of Chaos", *Nonlinear Dynamics and Systems Theory*, 1(1) 2337.

Baker, G. L. and J. P. Gollub (1996). *Chaotic Dynamics: an Introduction*, Cambridge University Press.

Barbour J. (2008). "The Nature of Time", submitted to the `fqxi.org/essay` essay competition, and available at
`platonia.com/nature_of_time_essay.pdf`

Bell, J. S. (1964). "On the Einstein Podolsky Rosen Paradox", *Physics*, 1, 195-200.

Bennion, E. G. (1952). “The Cowles Commission’s ‘Simultaneous-Equation Approach’: a Simplified Explanation”, *The Review of Economics and Statistics*, 34, 49-56.

Cox, D. R. and N. Wermuth (2001). “Some Statistical Aspects of Causality”, *European Sociological Review*, 17, 65-74.

Deutsch, D. (1997). *The Fabric of Reality*, Penguin Books.

Einstein A., B. Podolsky, and N. Rosen (1935). “Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?”, *Physical Review*, 47, 777-780.

Everett, H., III, (1957). ‘ “Relative State” Formulation of Quantum Mechanics’, *Reviews of Modern Physics*, 29, 454-62.

Frost, P. and J. Wolstenholme (1863). English translation of Isaac Newton’s *Principia*, Macmillan. Digitised by Google and available at `http://books.google.fr/books/about/Principia.html?id=1x9LAAAAYAAJ`

Granger, C. W. J. (1969). “Investigating Causal Relations by Econometric Models and Cross-spectral Methods”, *Econometrica*, 37, 424-38.

Heckman, J. J. (1999). “Instrumental variables: response to Angrist and Imbens”, *Journal of Human Resources*, 34, 828-37.

Heckman, J. J. (2001). “Econometrics and Empirical Economics”, *Journal of Econometrics*, 100, 3-5.

Heckman, J. J. (2008), “Econometric Causality”, *International Statistical Review*, 76, 1-27.

Holland, P. W. (1986). “Statistics and Causal Inference”, *Journal of the American Statistical Association*, 81 (396), 945-70.

Laplace, Pierre Simon. *Essai Philosphique sur les Probabilités*, digitised and available at `http://archive.org/details/essaiphilosophi00laplgoog`

Pereboom, Derk (2003). *Living without Free Will*, Cambridge University Press.

Popper, Karl (1972). *Objective Knowledge: an Evolutionary Approach*, Clarendon Press.

Rovelli, C. (1996). “Relational quantum mechanics”, *International Journal of Theoretical Physics*, 35, 1637-78.

Russell, Bertrand (1988). *Mysticism and Logic and Other Essays*, republished by Rowman & Littlefield Publishers.

St Augustin (401). *Confessions*, translated by Edward Bouverie Pusey, available at `http://www.sacred-texts.com/chr/augconf.htm`

Sims, C. A. (1972). “Money, income and Causality”, *American Economic Review*, 62, 540-52.

Sims, C. A. (1980). “Macroeconomics and Reality”, *Econometrica*, 48, 1-48.

Smerlak M. and C. Rovelli (2007). “Relational EPR”, *Foundations of Physics*, 37, 427-45.