“things are made to look the same only if we fail to examine them too closely”
Nancy Cartwright

Table of Contents

Abstract

Causal relations show us mechanisms that are responsible for events, give us control and provide proof for hypotheses. But causality is often badly defined and used in an intuitive way. Different disciplines of science use different definition, and often no definition at all. This causes confusion that can lead to the wrong conclusions. This confusion can also be used in order to hide or distract from the real causes of a event.

In this thesis I discuss the conceptual problems with causality that exist in philosophy. Many of them have, even after decades of discussion, not been properly solved. However, in most practical cases we can state conditions under which the problems can be ignored, and this allows us to ignore the conceptual problems for most practical purposes.
Combining the many definitions found in the literature, and taking the conceptual problems into consideration, I have composed the following definition of a cause:
“a cause is any member of any actual set of plausible conditions that are jointly sufficient to produce the effect. That the cause leads or contributes to the effect has to follow logically from the relevant laws.”

Within this definition we can distinguish several types of causes. These types of causes turn out to be the most important distinction to make if we want to draw conclusions from a cause-effect relationship. Examples from different disciplines of science show that determining the type of cause is often not easy, and authors can have strong incentives to misrepresent causal strengths. I have built a decision tree to help the reader decide what kind of cause he is dealing with, and to give a qualitative indication of a causal strength that can be determined from this relationship. The causal tree works for cases where all the causes of the effect are known to the user of the tree.

Causality needs a proper definition, which respects the conceptual problems that are relevant to the topic. An intuitive notion of causality is almost a guarantee for trouble. Even more important is the determination of the type of cause that we are dealing with, and the causal strength that comes with it. This cannot be properly determined without taking at least the most important other causes in the system into consideration.

Introduction

Causal relations play a large role in decision making, both in science and in everyday life, and are essential to control and manipulation of the world around us. As young children we learn about causality by observing regularities. We often have to guess what the nature is of those relationships between the events we observe, because we cannot always see all the causes and all the effects.

We don't grow up with the tools for critical analysis of relationships, but with rhyme and metaphor. At the same age we learn about causality, we get told stories and fairy tales that use analogies to give us the feeling that things are intrinsically related, while in reality they are only related in appearance. We do not learn to distinguish between unreal relationships, real relationships and causal relationships. This vagueness can be used or misused by marketing companies, politicians, or religious leaders, to influence our behaviour.

Looking around us we see people take decisions that are completely unlogical, with sometimes devastating consequences. Alternative medications that have been proven to have no beneficial health effects at all, are so popular that the demand for them drives species into extinction. Most of these claimed health effects are based on an analogy or a metaphor. Dried sea turtle shell gives you stronger nails, the powdered testicles of a tiger improve your libido etc. Some seem less destructive, like grass powder sold as super food because 'cats eat grass when they feel sick, and cats are closer to nature than we are, so grass must make you healthy' (NRC). Grass has no nutritional or medicinal value whatsoever, and cats eat it because it makes them vomit. Grass probably won't go extinct anytime soon, but the claims can still be bad if the user believes in them and therefore refrains from doing something that does improve his health. In the case of the grass, the 'healer' who advises the use of grass, claims it can cure cancer.

People who are badly trained in critical thinking seem to be an easy victim for fake causal relations. But are scientists doing any better?

In science, the use of causality is often a bit fuzzy. Different scientific disciplines have different definitions and uses of causality, and also different problems with causality. Often, no definition is given, and even within disciplines every researcher can have their own intuitive notion of causality. The vagueness of our understanding of causality can suggest a strong proof where there is none. In some papers, this fuzziness is misused to lead the reader to the wrong conclusions. In others, mistakes are made that may have lead the author to the wrong conclusion.

Many published scientific studies present claims about causal relationships that later turned out to be wrong. Ioannidis (2005)found as main reasons for such wrong causal claims that “a research finding is less likely to be true [...] where there is greater flexibility in designs, definitions, outcomes, and analytical modes. [...] Simulations show that for most study designs and settings, it is more likely for a research claim to be false than true. “

To be able to draw sound conclusions or to take rational decisions we need to have a thorough understanding of causality, both in science and in life.

The goal of this research is to gain insight into the diversity of thinking about causality in different scientific disciplines. In order to achieve this, I will ask the following questions:

- What are the conceptual problems we should take into account when thinking about causal relationships, and can we find conditions under which we can ignore them?
- What is the role of causality in different disciplines of science? Is causality well defined in those disciplines, and if yes, what differences are there and where do definitions overlap?
- Can we combine definitions to come to a generic definition that is useful for all disciplines of science?
- What kinds of causes can we distinguish and what do they tell us about the strength of causal relationships?

At some point in this paper I expect we will settle on an acceptable, or workable, definition of causality. By that time we can start looking at some cause-effect relations. We can look at the strengths of those relations, by figuring out if the cause is a direct or indirect cause, whether it is the total cause, or only a contributing cause. And if there are more, would one of the multiple causes alone be enough to cause the effect? What if we take away a few of the causes, would the effect still occur? These questions already give a hint to the number of different kinds of causes.

Having a better understanding of the meanings and applications of, and problems with causality, might enable us to develop a decision tree that features all types of causes. I will try to make a causal tree, that will make clear what kind of causality we are dealing with when we have a certain example in mind. When going through the decision tree, we will run into questions we have not asked ourselves yet, and hereby get a better understanding of the type of relationships we are dealing with. We can use it to test causal relationships and their strengths, and it can tell us how certain we can be of the conclusions we draw from those relationships. Hopefully, the decision tree will help us avoid mistakes in our own assumptions about causality, and discover them in the work of others.

Conceptual problems

If we talk about causality, we first have to define what we mean by it. There are many definitions, each with their own problems and limitations.

The trouble with the definitions

A causal relationship is always a relationship between two or more events, and that is as far as we get without running into trouble. Anything else we can think of can immediately be questioned.

One event can be said to cause another if the effect always occurs after the cause. Our observations tell us that must be correct, but we meet some major problems with that definition. It is a consistent observation that every morning when the rooster crows (event 1), the sun starts rising (event 2). But does that mean the rooster causes the sun to rise?

So maybe not all events that occur consistently before a second event are a cause to that second event, but if something *is* a cause of a certain effect, we think we can safely claim that the cause always has to precede the effect. But that relies heavily on our intuition of time as a something that flows forward. Is the direction of time something we can objectively prove from, for example, the laws of nature? Or by using logic? Or is it merely something we perceive, but that is not necessarily a part of the natural world?

And even if we are able to determine, or decide to accept, that the direction of time is forward, and a certain effect always occurs after the cause, we can never be sure if the next time we observe an event that has so far always been followed by a certain effect, will do so again. This is described by Hume (1739) in his work about the induction problem.

Of course, we might say, we have to find a mechanism by which the one event *has to* cause the other. We look for a necessity for the effect to follow the cause. If that necessity is there, we will know for sure that the event will always cause the effect in the future. But can we find any necessity in the laws of nature?

What are those laws of nature anyway? What can they tell us and how can we use them? Are they any use to us at all when describing causality?

We could define a cause by saying event A causes event B if without the occurrence of event A, event B would not have occurred. The counterfactual theories of causation work like that. But counterfactuals are highly controversial. How can we be certain that B would not have occurred if there had not been an event A?

We learn a lot about causal relations by trying to manipulate the world around us. But wait a minute, doesn´t the fact that we can manipulate things prove that the direction of time is forward? Since we can only know the past but cannot change it, and only change the future but cannot know it?

And what if the world is entirely deterministic and only depends on the state of the world on a certain moment, can we still manipulate the future by interventions in the present?
On the other hand, if the world is probabilistic, can an event still be a real cause if there is only a chance that it produces a certain effect?

Observation

Looking around us we see it all the time: something happens, and then, always, something else follows. This is how, when we are young, we learn about causality. It seems fine to assume that the first event is causing the second. But all we have observed is a correlation or a regularity. There are many cases in which one event always occurs after the other, also when the first event does not cause the second. Every time the cat hides under the bed, it will start to rain some minutes later. We know that the cat hiding under the bed does not cause the rain, but a thunderstorm that is already audible outside before the rain starts, makes the cat hide, and causes the rain as soon as it comes close enough. So two events always happening after each other does not guarantee that the first causes the second, but there may be a mutual cause. A mutual cause is a cause that causes both events, in this case the second event is a bit delayed. The two events may even be part of a series of events that are not causally related at all.

Most of our experience with causal relationships are through observation. But there are a lot of relationships we do not see, hidden causes, effects that are too far away, or too close, too fast or too slow for us to observe, or effects that occur so much later that we don't observe them. Observation does not give us all the facts about everything that happens.

A lot of events we observe seem related, but are not. Our brains are constructed to see patterns. Rhyme and metaphor can make us feel as if things are causally related, even if they are only related in appearance. In commercials we see a woman enjoying a coffee and having all the time in the world. We are told to buy the coffee, to have 'a moment for ourselves'. Buying the coffee buys you time, a wonderful view and an amazing sunset, is the message. Of course we know that the package only contains some instant coffee powder and no sunset, no extra time and no horizon behind a silhouette of palm trees.
In other cases, products are marketed with implicit promises of health benefits. Regulations now prevent manufacturers from making explicit health claims on packages if the health benefit has not been proven, but they have found clever ways of giving selective information that leads the consumer to the conclusion that the product must be good for his health, making use of the consumer's tendency to make causal mistakes.

Religion and superstition are filled with causal misunderstandings. F.B. Skinner (1947) has found superstitious behaviour in pigeons. The pigeons were fed on a regular interval. The moment of feeding was unrelated to the behaviour of the pigeons. Yet, the pigeons developed habits in order to get food. They assumed that whatever they had been doing right before the food came, was the cause of the arrival of the food. Each bird developed its own behaviour. Because the food kept coming at regular intervals and the birds kept doing the same rituals, the behaviour of the birds was reinforced by the arrival of new food. Every time they did a certain movement, the food came. Skinner assumed the birds were trying to influence the automatic food system, and had the feeling they succeeded. Skinner saw analogies with the behaviour of humans when they perform rituals. The movements the birds made would be a nice example of an event that precedes another event, being assumed (in this case by the birds) to be the cause of the second event, the arrival of food.
Skinners conclusions have been challenged by many other researchers who have tried the same experiment and found the same behaviour, but have reached different conclusions. Staddon and Simmelhag (1971) discovered that the pigeons did not do the rituals in the time interval after the food was received, but only right before the new food was expected. They conclude that since the time interval was short and of constant length, the pigeons knew when the food

would arrive, and them doing the rituals may not have been an attempt to influence the process.

Causal mistakes are also common in human societies. Some religions are based on one single causal mistake, like the cargo religions. Cargo religions are sometimes found in places where a society suddenly gets disrupted, for instance by colonisation or a sudden encounter with a different culture.
One of the most famous cargo religions came into existence on Melanesia. The Melanesians were put to work for the Australians to level strips of forest in order to build a runway for the Australian air planes. They witnessed the planes coming in and supplying the colonisers with cargo. In the Melanesian culture, it is very important to be competitive in gaining material wealth and showing it off (Schwarz 1962). The Melanesians saw that the Australians were doing rather well at that and felt inferior. They also wanted the wealth that was brought to the Australians by the cargo planes. They started building dummy airstrips, including a fake watch tower and a fake air plane made of branches, parked at the end of their runway. The runway was expected to bring in the cargo planes that would finally give them some cargo of their own.

Confusion about causality occurs most often when the real cause is unknown or invisible. Thousands of witches have been burned, mostly after epidemics, in times when the causes of the spread of disease were not known. A person being sighted in the wrong spot at the wrong time would be a more plausible cause of disease, if he were a witch, than invisible viruses whose existence were at that time completely unknown.

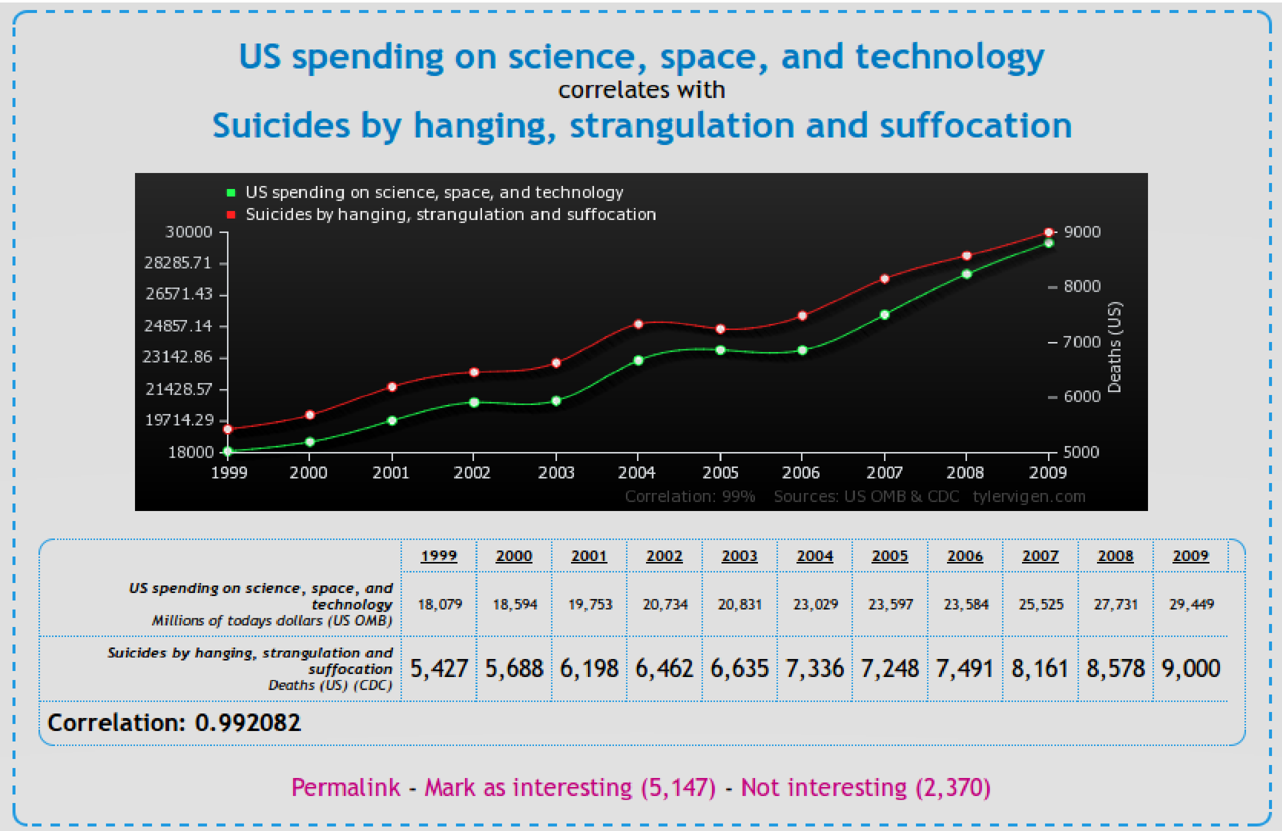

	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
US spending on science, space, and technology *Millions of todays dollars (US OMB)*	18,079	18,594	19,753	20,734	20,831	23,029	23,597	23,584	25,525	27,731	29,449
Suicides by hanging, strangulation and suffocation *Deaths (US) (CDC)*	5,427	5,688	6,198	6,462	6,635	7,336	7,248	7,491	8,161	8,578	9,000

Correlation: 0.992082

Permalink - Mark as interesting (5,147) - Not interesting (2,370)

Figure 1: Spurious correlation, from http://www.tylervigen.com

Figure 1 is an example of a spurious correlation that nobody would seriously think is a causal relationship. The only safe use for these kinds of correlations are in humour and education.

Time and intuition

In our every day lives, we make use of an intuitive notion of causality, that has everything to do with our notion of time. According to Reichenbach (1956), we respond emotionally to the flow of time, because it is not under our control, yet it does carry us steadily and irreversibly towards our deaths. Most people find that a disturbing thought. We know the past but cannot change it. We feel we can change the future, but cannot see it. We seem to live in an ever changing present, that travels with us through time. Ancient philosophers have done efforts to understand the nature of time. Parmenides argued there can be no change, all there is is the present. Zeno constructed the paradox of the flying arrow to prove that motion is impossible. In this paradox, the arrow is said to move if it goes from one point to the other. But if it is at rest in exactly one point on a certain moment, how can it jump to the next point? Without motion and change, is there a need for the concept of time? Does time still exist if there is no motion? Heraclitus argues the contrary: that there is only motion. A river is different every day, because the water in it is new. Heraclitus uses metaphors to make us feel that he is right. Intuitively, most people would agree with Heraclitus, but logically we cannot easily ignore the paradoxes of Zeno. None of them really answers any questions about the structure of time. Heraclitus tries to understand time by finding analogies with everyday experiences. Today, we would say that questions, even if they arise from everyday experiences can best be answered using the scientific method.

A direction of time

A causal chain needs time to develop from cause to effect. Before special relativity, simultaneity was often associated with the impossibility of causal connections. If two events happened

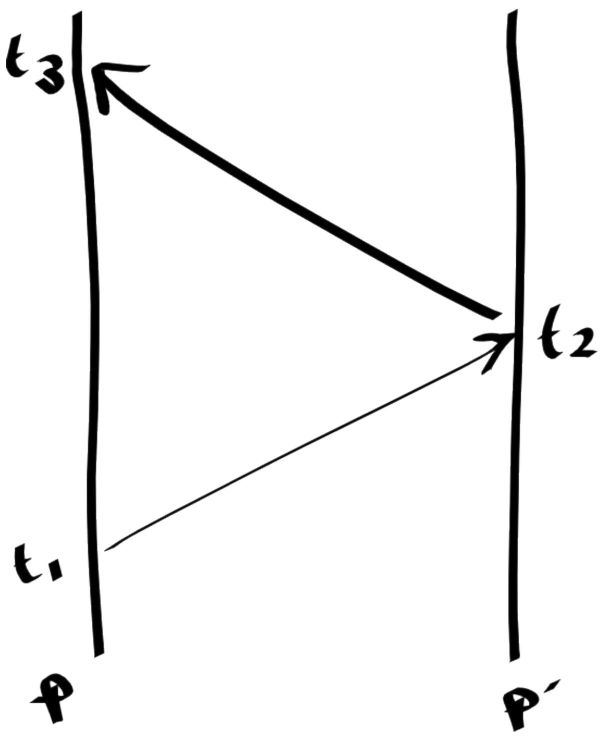

Figure 3: light signal travelling from point P to P' and back (from Reichenbach 1956)

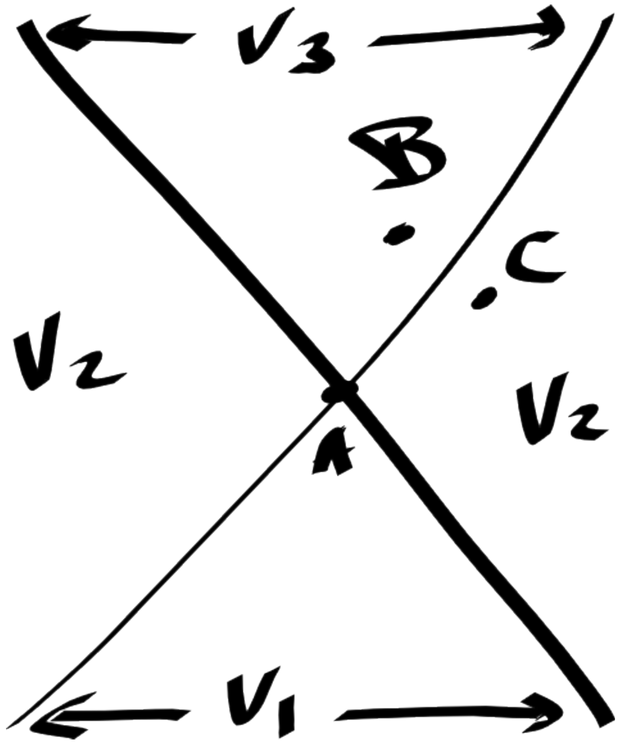

Figure 2: Einstein-Minkowski diagram of spacetime, where the volume V1 represents the past, V3 represents the future and V2 represent volumes of spacetime that are indeterminate as to the direction of spacetime (adapted from Reichenbach 1956)

simultaneously, the one could not cause the other. But since Einstein we know that, within a frame of reference, not only simultaneous events are excluded from causal relations. Light is the fastest possible way for a signal to travel. If within one frame of reference at t1 light travels from a point P in two dimensional space to another point P', gets reflected there at t2, back to point P where it arrives at t3, then all events between t1 and t3 are excluded from causal relationships (see figure 3).

Figure 2 shows the Einstein-Minkowski diagram of spacetime, in which the volume V1 represents the past, V3 represents the future and volumes V2 are indeterminate as to the order of time. This means that if in the present we have an event A, it can be a cause of event B in V3, but not of event C in V2. In frames of reference that move with respect to each other, a cause can happen after its effect, but here we focus on events within a single frame of reference.

According to Reichenbach, special relativity assumes the structure of a causal net, which only gives time order, but no direction. From equations 1 and 2, in which *t* is time, *x* is position, *v* is velocity and *c* is the speed of light, we can see that if we replace *t* in the Lorenz transformation with *minus t,* nothing happens, except that *v* has the opposite direction.

$$x=\frac{x'-vt}{\sqrt{1-\frac{v^2}{t^2}}} \tag{1}$$

$$t=\frac{t'-\frac{v}{c^2}x'}{\sqrt{1-\frac{v^2}{t^2}}} \tag{2}$$

The laws look the same and can be used in the same way. They do not need positive time to be functional laws of motion. So if we reverse time, we only reverse the direction of motion, and we can use the same transformation laws. In the case of a reversible process of a ball thrown from A to B in reverse time, it would look the same as if it were thrown in the other direction in positive time (figure 4).

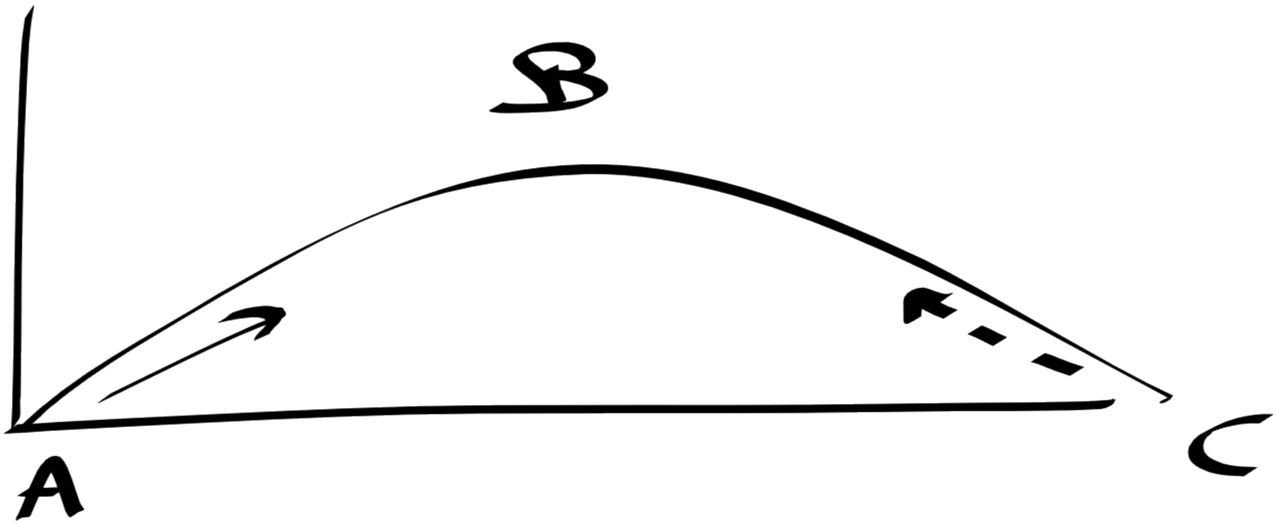

Figure 4: path of a ball in positive time (solid arrow) and negative time (broken arrow) (adapted from Reichenbach 1956)

We usually use time direction to define causal direction but Reichenbach (1921) does it the other way around. He developed a causal theory of time, in which temporal order can be reduced to causal order. He tries to define the direction of time by using the direction of causality. He therefore cannot use the direction of time to define the direction of causality, because that would make it circular. He tries to find a way to define the cause-effect relation without reference to time. He looks for an answer in the laws of physics.

The laws of physics are functional relations that say that if variables $x_1..x_n$ have certain values, then the value of x_{n+1} can be determined from them. They can therefore be written like in eq. 3.

$$x_{n+1} = f(x_1 \ldots x_n) \tag{3}$$

But eq. 3 can be solved for any of the variables x= in it, so it can be rewritten as eq. 4,

$$x_1 = g(x_2 .. x_{n+1}) \tag{4}$$

in which x_1 can be determined from the other values. So these functional relationships are symmetrical, and tell us nothing about order. Does this mean we cannot find a causal order in the laws of nature? Or that causality is a symmetrical relationship? According to Reichenbach, there must be more to causality than symmetrical relationships if we want to find an order in causal relations.

An order of time

To find this time order in these events, we must find something that is invariant if we invert time. Reichenbach's macrostatistical theory implements a causal theory of time. Reichenbach's theory does not depend on causes being temporally prior to effects. Instead he attempts to build a theory that will yield an asymmetry that can be used to define a relation of temporal priority. A key concept in this theory is causal betweenness. An event B is causally between the events A and C if the relations hold:

$1 > P(C|B) > P(C|A) > P(C) > 0$

$1 > P(A|B) > P(A|C) > P(A) > 0$

$P(C|A.B) = P(C|B)$

(Notation: $P(C|A.B) = P(C|B)$ means the Probability P of C given A combined with B is the probability P of C given B.)

If we throw the ball from A to C, and the point B is in between , it will still be between A and C in negative time. The betweenness is invariant to time reversal (figure 4). We can reverse all the arrows and still have B between A and C, but we cannot reverse one of the two arrows. This would no longer be compatible with our observation. Causal betweenness establishes a linear time order, but not a time direction.

Causal chains and causal nets

We have now found an order in events, and we can start building a causal net, with events being between other events. If we do that, we can either reverse all the arrows in the net, or reverse none. In a causal net we cannot end up at the beginning again. This is not a logical limitation, but an empirical one (which to me is not very satisfying, but it may well be as far as we get).

Causal nets are open. A closed causal chain would mean we move in a circle. From observation we know that we cannot end up in a certain point in time more than once. Reichenbach concludes from this that the time order of events is linear. So betweenness gave us an order, the fact that we never experience being at the same event twice gave us a linear order, but we have not yet found a direction of time.

Figure 5: Causal chain of events

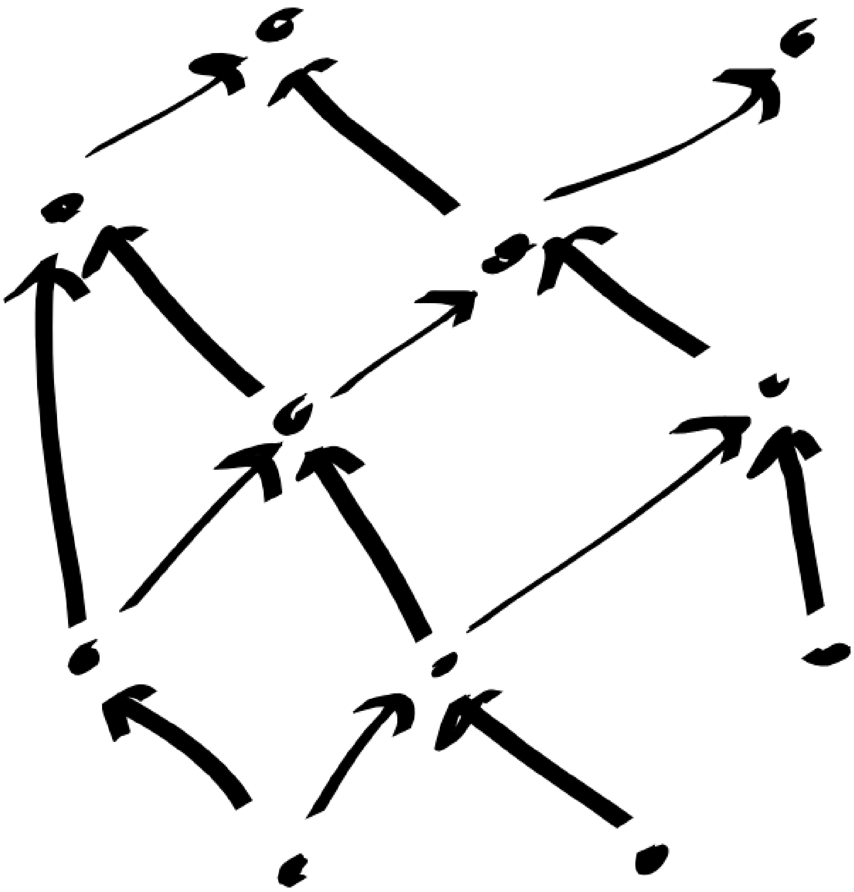

Figure 6: Causal net of events

The direction of increasing entropy

A thermodynamic equation alone does not tell us anything about a causal direction. It only gives us a relationship. If we look at the equation

$$p \cdot V = n \cdot R \cdot T \tag{5}$$

(In equation 5, *p* is the pressure (Pa), *V* is the volume (m3), *n* is the amount of substance of gas (moles), *R* is the ideal gas constant (8.314472 J·K−1 mol−1) and *T* is the absolute temperature (K).)

we see that whenever the temperature increases at constant volume, the pressure will increase. And the other way around. The equation does not tell us whether the temperature increase was caused by the increase in pressure, or the other way around. To know that, we will have to fall back on observations again, and check whether a piston has been lowered or a candle has been lit.

We have an intuition of which processes are reversible and which are not. If we watch a movie reversed, we see that mechanical processes look natural in reverse. A tumbleweed rolling through the streets of a deserted western town looks fine if played backwards. But a puff of smoke curling neatly into the bad guy's revolver looks impossible. Mechanical processes are reversible, most thermodynamical processes, like the mixing of gasses and the flow of heat, are irreversible. During these processes, the entropy increases. The second law of thermodynamics says entropy can never decrease in a closed system. Processes in which entropy increases, are irreversible. Completely reversible processes do not exist, because even in mechanical processes, there is always a little bit of heat exchange, with a small influence on the entropy.

The increase of entropy has to do with chance. The chance of finding a system in a state with an arrangement of molecules that is completely ordered, is much smaller than to find it in a state of disorder. Simply because there are many more disordered states than there are ordered states. According to Reichenbach, there are two kinds of laws, a strict, causal law, and a probability law. He argues that the second law of thermodynamics is a probability law. What kind of law is more fundamental? Should we see probabilistic laws as a limitation of our ability to know exact outcomes? Or are all strict laws limiting cases of probabilistic laws, the cases where the chances of a certain outcome are so high that we never see a different outcome?

If we assume the universe is a closed system, could we determine a direction in causality from the increase of entropy? If there is no way in which a closed system can have decreasing entropy, we could say that the direction of time of a system is the direction of increasing entropy.

Reichenbach says we can say that it is a matter of convention to select the direction of growing entropy as the direction of time, but “it is an empirical fact that in all branch systems the entropy increases in the same direction”. He developed the construction of a conjunctive fork to impose a direction or asymmetry upon the linear time order. He gives the example of a thunderstorm T, that is the common cause of a strong wind W and a fire F, which have a common effect E: a large surface of vegetation is destroyed (figure 7).

According to Reichenbach, the branch systems develop in the direction of growing entropy, due to the statistical isotropy of the universe. The chances are larger that the entropy increases than that the entropy decreases. The past produces the future, and not vice versa.

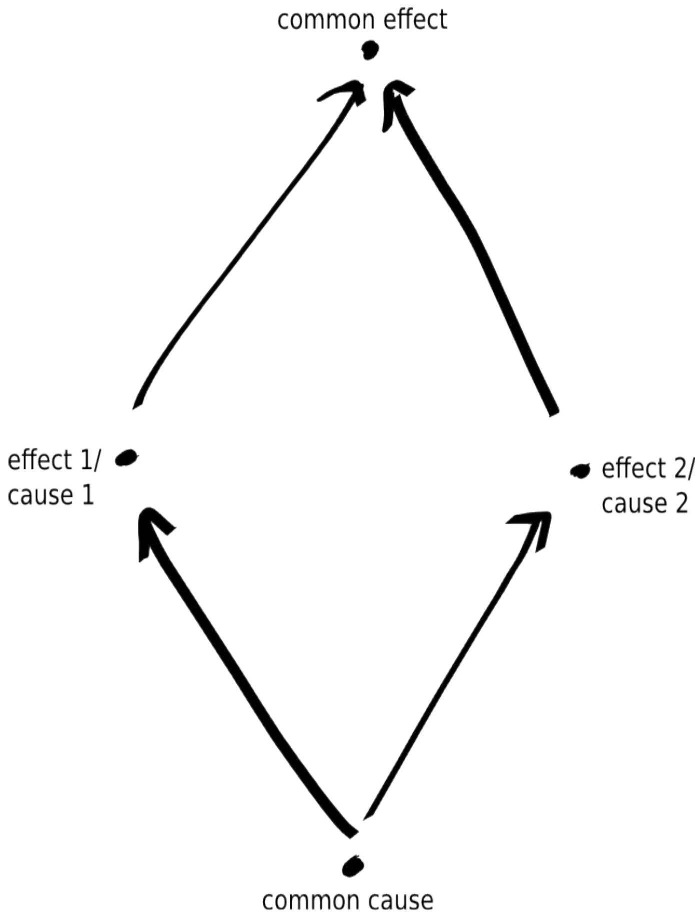

Figure 7: Causal fork featuring a common cause and a common effect. The direction of increasing entropy is upwards (adapted from Reichenbach 1956).

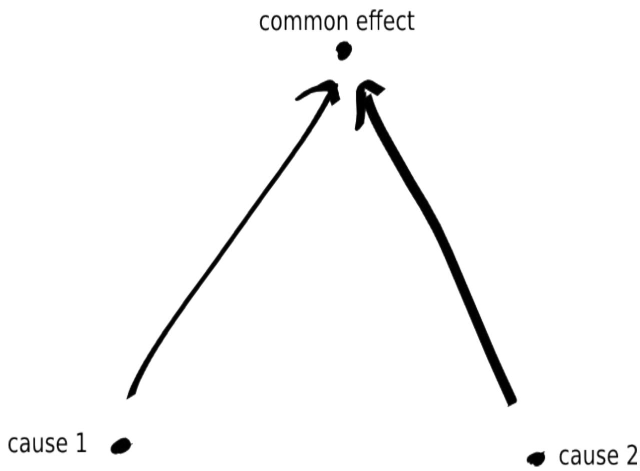

Figure 8: conjunctive fork showing two causes and their common effect, the direction of increasing entropy is upwards (adapted from Reichenbach 1956).

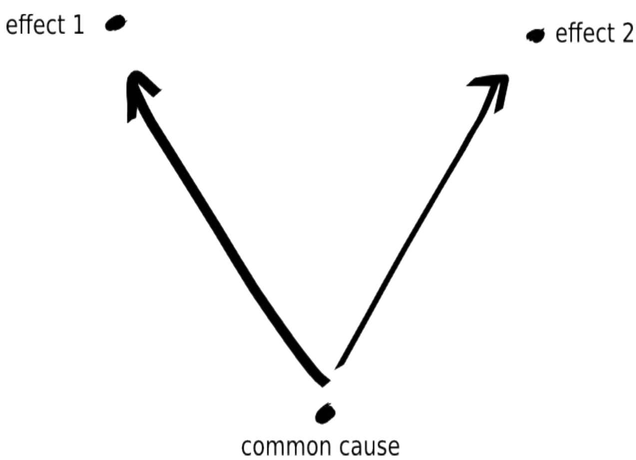

Figure 9: Causal fork featuring a common cause and its two effects. The direction of increasing entropy is upwards (adapted from Reichenbach 1956).

But there is a way to get systems in which entropy decreases. You can find them if you extend time to imaginary time, using Wick rotation. Wick rotation is a way to connect statistical physics to quantum mechanics using an analytical continuation of *t* to imaginary time τ *(which is $\hbar/k_B T$)*.

Imaginary time is a bit hard to imagine, but that does not mean it cannot exist. It also does not say anything about the direction of real time. It can be visualised in the same way as real numbers and imaginary numbers (figure 10). If you imagine real time to be a horizontal line running from the past to the future, imaginary time is a line running vertically, from the imaginary past to the imaginary future (figure 11).

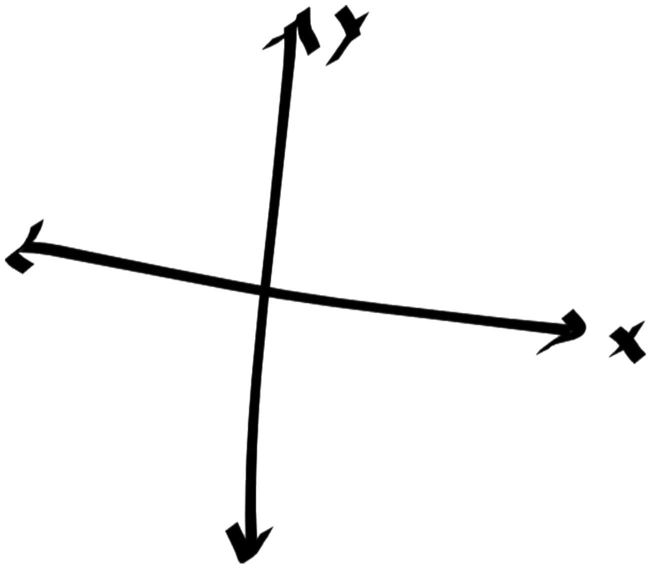

Figure 10: Complex numbers: C = x + iy

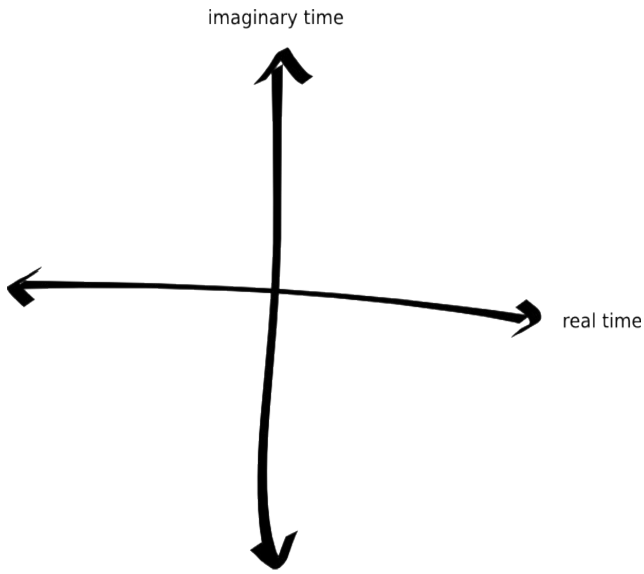

Figure 11: Complex time = t + iτ

Imaginary time makes it possible to have negative absolute temperatures, and therefore systems with decreasing entropy. We can no longer be sure that the entropy in closed systems always increases, so it is no longer completely safe to assume the direction of time and causality is the direction of increasing entropy. Reichenbach does accept the direction of time to be the direction of increasing entropy, because of the lower entropy state of the early universe. But in cosmology, imaginary time is used as more than a mathematical tool. Some theories in cosmology predict that after the universe has gotten into a state of high entropy, it will start to decrease its entropy again.

Stephen Hawking (2001) says about imaginary time: “One might think this means that imaginary numbers are just a mathematical game having nothing to do with the real world. From the viewpoint of positivist philosophy, however, one cannot determine what is real. All one can do is find which mathematical models describe the universe we live in. It turns out that a mathematical model involving imaginary time predicts not only effects we have already observed but also effects we have not been able to measure yet nevertheless believe in for other reasons. So what is real and what is imaginary? Is the distinction just in our minds?”

The induction problem

Hume (1748) says causal relationships are discovered not by reasoning, but by observing that certain effects occur in combination with certain circumstances. But according to Hume's problem of induction, if causal relations are only the assumptions we make after many past observations, we have no logical reason to assume we can extend our expectations of causal relations to the future. Even if we have seen that a stone falling into a pond creates ripples in the water every single time, there is no logical reason to assume that the stone will keep doing that every time in the future. Did Hume really expect the stone to hit the water with some future throw, and witness how it does not make ripples in the water? Probably not.
Rudolf Carnap (1966) does not believe Hume intended to reject the concept of causality, rather he wanted to purify it. What Hume did reject was the component of necessity in the concept of causality.

The induction problem is very much related to the problems with counterfactuals. We will not solve it here, but get back to it in the next chapters.

The laws of physics

To distinguish a causal relationship from a correlation, we can demand that a cause has to be in some way connected to the effect, via a law or a plausible mechanism. Plausibility itself is a rather problematic concept. It is not the same as probability. We can consider whether or not a cause is a plausible one, without being able to assign a probability to it being a cause of the effect. Plausibility alone does not tell us anything about an event, except that if we think a possible cause is a plausible one, it may be worth our time to conduct some experiments to find out if it can really be the cause we are looking for.

But when is a cause plausible? Resnik (2003) says, as a rule of thumb, the plausibility of a hypothesis is increased by coherence, explanatory power, analogy, precedence, precision and simplicity. He gives the following definitions of these characteristics:

- Coherence. The hypothesis should be consistent with and supported by our background knowledge and theories. If a hypothesis requires us to reject widely accepted scientific theories and facts, then it is not plausible.
- Explanatory power. The hypothesis should be able to explain important facts and phenomena. Hypotheses that have no explanatory power are not plausible.
- Analogy. The hypothesis should posit causal mechanisms or processes that are similar to other well-understood mechanisms and processes. A hypothesis that posits radically new and unfamiliar mechanisms and processes lacks plausibility.
- Precedence. Events posited by the hypothesis should be similar to previously observed events, which set a historical precedent for the hypothesis.
- Precision. The hypothesis should be reasonably precise. Although there are limits to precision in science, a hopelessly vague hypothesis should not be regarded as plausible.
- Simplicity. The hypothesis should be parsimonious. Recondite and complex hypotheses are not as plausible as parsimonious ones.

None of Resniks criteria are without problems. In the chapter about laws we will see that a theory that explains well, is never precise. So we cannot satisfy those two criteria at the same time. The other criteria are rather conservative, and would only be useful in periods of normal science, but would fail when paradigm shifts occur. The simplicity argument has always been under discussion, because it is not an argument that follows from logic, or that is reflected in the beautiful and exciting complexity of nature.

Then what criteria can we give for plausibility? We can start with the demand that a plausible cause has to at least be logically possible. And then we can add that it should also be practically possible. Whether something is practically possible depends on the circumstances and on the laws that govern the process. But now we have described possibility, instead of plausibility. We are still missing some ingredient that gives a likelihood to the event being a cause. We will get back to that later.

If we want to base our judgment of causal relationships on plausibility and laws, we might want to ask ourselves: how do these cooperate? How accurate are the laws of physics? What is their explanatory power? Can they be used to get certainty about causal relationships?

What are the laws of physics good for?

Nancy Cartwright (1983) distinguishes between phenomenological laws and theoretical laws. For physicists, phenomenological laws are laws that describe what happens in a physical process, whether directly observable or not. Phenomenological laws function really well in describing reality. Theoretical laws are fundamental and explanatory. Cartwright argues that the fundamental laws in natural science only describe idealised objects, and do not describe reality. They are useful for explanation, but not for describing nature.

If a theory explains a phenomenon well, that does not mean it is true, according to van Fraassen (1980). There is no guarantee that if A explains B, and B is true, than A has to be true. We can think of many examples in which case an explanation that sounds perfectly reasonable, is not true.

Salmon (1984) also values causal explanation. He says causal explanation involves a spacio-temporal, continuous causal process. And all genuine explanation must be causal in this way. According to Cartwright, only if the explanation is causal it is necessarily true. But also in causal explanations, one can make mistakes, which would again make the explanation false.

For two thousand years, the theory of spontaneous generation has been generally accepted. Aristotle (1) wrote “with animals, some spring from parent animals according to their kind, whilst others grow spontaneously and not from kindred stock; and of these instances of spontaneous generation some come from putrefying earth or vegetable matter, as is the case with a number of insects”. An example of spontaneous generation was the mice who got generated out of the mud of the river Nile, as every time the river flooded, large numbers of mice appeared. Aristotle came up with an explanation for spontaneous generation that, to his contemporaries, sounded plausible: “Animals and plants come into being in earth and in liquid because there is water in earth, and air in water, and in all air is vital heat so that in a sense all things are full of soul. Therefore living things form quickly whenever this air and vital heat are enclosed in anything.” (Aristotle 2). In this case, it was true that every time the Nile flooded, mice were all over the place. The theory explained the phenomena, but the causal explanation was not true.

Cartwright's problem is that theoretical laws are not true for objects of reality. Phenomenological laws are. But theoretical laws can be true for objects in models. So theoretical laws do not directly tell us anything about reality, but they tell us something about objects in models, and models can tell us something about reality. This suggests that in an indirect way, they are useful for describing reality anyway. I find this a bit weak, because it limits the possible use of theoretical laws to applications where you can use a model.

Types of laws

An important issue in the philosophy of causation is the distinction between causal laws and causal relations. David Hume (1739) discusses the question of which are more basic. His conclusion is that causal laws are more basic, and causal relations are logically supervenient upon causal laws. This means there cannot be a difference in causal relations without there being a difference in causal laws. An exact similarity in causal laws guarantees an exact similarity in causal relationships if all the circumstances stay the same.

So this supervenience leads us to modalities like necessity. If the laws change, the outcome has to change.

Carnap (1966) says there is a difference between logical modalities (logically necessary, logically possible), causal modalities (causally necessary, causally possible) and many other kinds. He believes it is possible to have a logic of causal modalities. He tries to find this logic in order to get the discussion about causality out of the realm of metaphysics.

To explain the logic of causal modalities, he first explains the counterfactual conditional: if a certain event had not taken place, a certain other event would have followed. We will discuss some problems with the concept of counterfactuals later, but for now let us take them for granted.

A distinction must be made between genuine laws and accidental universals. A genuine law can be a sufficient justification for a counterfactual. Imagine a raindrop that fell from the sky and landed on your head. The law of gravitation is enough justification for the following counterfactual: if your head had not been there, and everything else had stayed the same, the drop would have continued its fall. But if we would use an accidental universal as a justification for a counterfactual, we could get strange results. Carnap uses the following example by Goodman (1947) of an absurd counterfactual: If the universal would be: all the coins in my pocket yesterday were made of silver, the counterfactual claim could be: if this coin had been in my pocket yesterday, it would have been made of silver.

Carnap proposes to put statements into two categories: statements that have a nomic (lawlike) form and statements that do not. It is not a rule that those statements have to be true. The statement “gravity decreases with the third power of the distance” is of the first kind, lawlike, but untrue, and therefore not a law. A lawlike statement that is also true is a basic law. Basic laws are universal. There are practical laws and technical laws that are not universal, and therefore not basic laws. They only hold under certain conditions. Some laws that were assumed to be basic, later turn out to be limited to certain conditions, or to be a special case of a new law that is at that moment thought to be basic. Newton's laws are an example of the last, they turned out to be a special case of the laws that followed from relativity, with low velocity as a condition. Newton's laws have now been reduced to practical laws, or maybe even technical laws.

Of course we can never know for sure if a law is true, or even if it holds universally. But that does not have to prevent us from using the concept of truth in defining what is meant by a basic law. If we are fuzzy about using the word “true”, we can also say “highly confirmed”. What scientists mean by a basic law is something that holds regardless of whether any human being is aware of it. So whether a statement is a basic law or not should have nothing to do with the level of confirmation. There are also derivative laws, basic laws that are restricted in space or time.

Carnap gives a definition of causal truth: a concept is causally true if it is a logical consequence of the class of all basic laws or derivative laws. This way, Carnap has shown that there is no need to discard of causality as a metaphysical concept. We can look at it in a scientific way and use the rules of logic in combination with basic laws to find causal truth.

Knowing a causal relation means predictability. In a deterministic world, if we were able to know all the details, a complicated series of events could have been predicted. If for now we ignore that we usually don't know all the details, and assume that all the relevant facts can in principle be known. If we ignore the practical problems with obtaining all the facts, and also the limitations in principle by quantum theory on knowing all the facts at a subatomic level, we would think we could predict anything. But we need the facts, but also the laws to make our causal predictions. Knowing all the relevant laws is also a problem.

To say event A caused event B means that there are certain laws of nature from which event B can be logically deduced when combined with the full description of event A. It is irrelevant to the truth of a causal statement whether the laws are known to us or not. Whether we know about gravity or not, a coffeecup that I knock off the table will still fall to the ground. It would be strange to think truth is dependent on what we know of the laws today.

Necessity in the laws of physics

Our intuition, as well as the induction problem that we have not yet found a solution for, tell us that for a causal relationship there has to be more than just two or more events always happening in a certain order. We want to be able to find a mechanism that connects these events, and prove there is this necessity.

The right place to look for the necessity that could solve the problem of induction for causal relations, or at least give us an excuse to ignore it, would be the laws of physics. What we mean if we say an event B is caused by an event A, is that there are certain laws in nature from which we can logically deduce, when we have the full description of event A, that event A will cause event B. But this definition does not mention necessity at all, and we just stated that we would like to see some necessity. Could necessity be included in the definition of a law? Do laws imply necessity? According to most empiricists, a law is merely a universal conditional statement. A causal law is no different. Whenever A occurs, B will follow. But if you say: Iron expands when heated, is that nothing more than one thing that follows the other? Carnap uses the example: When iron is heated, the earth rotates. Carnap says we don't call it a law because the earth's rotation has nothing to do with the heating of iron. But how can you know for sure? I think two reasons why we would not call this a causal relation is that the earth already did rotate before we heated the iron, and we know we cannot influence the rotation of the earth by influencing the temperature of our piece of iron.

Carnap says to call something a law of nature, there has to be more than just one event following the other. There has to be a necessity, a certain event *must* follow another. There has to be some sort of necessary connection between the one event and the other. But it cannot be a logical necessity, because logical necessity means logical validity, and something is logically valid by virtue of the meanings of the term that occur in it. For any law of nature, we can think of a series of events that would violate it, without any logical self-contradiction. Laws of logic hold under all conceivable conditions.

So what necessity can there be in laws of nature, if it is not logical necessity? And how are the laws of nature any better if they have necessity in them? Do they gain predictive power? According to Carnap, they do not gain any predictive power. Predictions are made based on the knowledge of a current state, combined with the knowledge of the laws that apply to the process. Necessity does not play a role in getting accurate predictions. It only gives an emotional feeling of certainty. Hume says that you cannot observe necessity, and because you don't observe it, you shouldn't assert it. Hume's view is the conditionalist view: a statement about a causal relation is a conditional statement that only describes an observed regularity of nature.

Counterfactuals and causation

If we could use counterfactuals, we could simply say: "A is a cause of B if without A, B would not have happened." But counterfactuals have a bad name in philosophy. They are often done away with as metaphysical and not suitable for serious discussion. But we use them a lot and they come in handy in theories of causation and in explanation.
Counterfactuals are necessary in many areas, like law and justice, for instance in law of torts. How can you hold somebody accountable for an action that caused somebody else harm, if you cannot say that without this action the harm would not have been done? Surely the fact that we cannot think of a way in which the harm would have been inflicted otherwise, can not be enough to say that it wouldn't?

The reason why counterfactuals are so controversial is that you can never be sure that a certain event would have happened if another event had not happened. It sounds like speculation to try to make assumptions like that. But are counterfactuals so much different from other predictions? You would have to think about how a counterfactual statement could be falsified. If we say "if the chair I am sitting on had not been here, I would have fallen to the ground", the statement would be falsified if my chair had suddenly disappeared and I were hoovering above the ground. We can also point out some mechanisms, gravity for instance, combined with the fact that I am sitting about 60 cm above a rather large planet, that cause a person to fall to the ground if a chair disappears. Without this additional information about the circumstances, the statement "if the chair I am sitting on had not been here, I would have fallen to the ground" would not only have been true, it would also have been false, if I had for instance been nowhere near any large body of gravity, or had been holding an enormous helium balloon.

Nelson Goodman (1947) distinguishes several kinds of counterfactuals, each with their own peculiarities. I feel his examples are rather technical and too linguistic for use in this thesis, but the problems he addresses are relevant.

Goodman identifies the main problem as that the consequent of a counterfactual never follows from the antecedent by logic alone. There are always assumptions and conditions that are not stated. And even if all conditions are stated, there is no law of logic that tells us the statement must be true.

Under what conditions can we accept a counterfactual as a useful statement? We could demand that the laws that can be applied are not empty laws. “All birds with four wings can fly” is just as true as “all birds with four wings are unable to fly”, simply because there are no such birds. Laws like that are useless. Also laws that state something that is too general to be false should be avoided. With these laws, we could make any true statement we want.
Relevant statements are the set of true statements that are both logically and non-logically compatible with the antecedent. But then we can still end up with statements that are all compatible with the antecedent, but incompatible with each other.
So we need more conditions: A counterfactual is true if there is some set of self-compatible, true statements that leads to the consequent, and there is no such set that leads to a consequent that is incompatible with the first. This criterion, however, can seldom be satisfied.

Then there is the problem with laws. If we look back at the silver coin, we see a law, “all coins in my pocket Tuesdays are made of silver”, which could be true. If the counterfactual is “yesterday was a Tuesday, if this copper coin had been in my pocket yesterday it would have been made of silver”, that would be strange, and we cannot help thinking that the right statement would be “if this copper coin had been in my pocket yesterday, not all coins would have been made of silver”. Something happened that made our law untrue, and it is therefore no longer a law. We need a real law, one that will allow us to draw the conclusions in the statement. A real law would have to be a sentence that can be used for prediction, a sentence that is both lawlike and true.

Hume (1739) defined a cause in the following way: “We may define a cause to be an object followed by another, where all the objects similar to the first are followed by objects similar to the second. Or, in other words, where the first had not been, the second would not have followed.” That last sentence sounds like a counterfactual theory of causation. Lewis (1973b) says we have moved on since then by distinguishing between accidental regularities and causal laws. We have now more clarity on counterfactuals. Lewis warns us that counterfactuals are still notoriously vague, but that does not mean we cannot give a clear account of their truth conditions.

Lewis (1979) says the present depends counterfactually on the past. If the past had been different, the present would be different. We accept that. If a giant meteorite had struck the earth yesterday, I would not be biking to the university today. Fair enough. But we seldom accept that the past depends counterfactually on the present. Which is a bit strange, to think that even if the present were different, the past would still be the same. How can the past be the same? Something in the past must have led to the present being different. We feel free to use whatever knowledge we have about the past in suppositions, but we do not feel free to use whatever we know about the future. We assume that facts we know from a previous time are counterfactually independent of this suppositions, and we can safely use them.
If we assume that present conditions have past causes, then whatever happened to those causes if the present were different? Did they suddenly fail to cause the present conditions?

We see a cause as something that makes a difference. A difference from when the cause had not been there. According to Lewis, counterfactuals are important for causality. He suggests we take counterfactuals at face value: “as statements about possible alternatives, somewhat vaguely described, in which the actual laws may or may not remain intact”.

He looks for a relationship of comparative similarity between possible worlds (Lewis 1973b).

Each world realises a counterfactual possibility of this world. Imagine that in our actual world I just knocked a coffeecup off the table, it falls and breaks. In Lewis' theory, there are worlds in which I did not knock the coffeecup off the table, and there are also worlds in which I did knock it off the table, but it did not break. There are even worlds in which I did knock the coffeecup off the table, and it did not even fall, but hoovered in mid air, refilled itself with fresh coffee and moved back onto the table. He orders possible worlds, according to how closely they are related to our actual world. Our actual world is closest to reality, and any two worlds can be ordered in accordance with their closeness in resemblance. The statement "If *A* were the case, *C* would be the case" is *true* in the actual world if and only if

1. there are no possible *A*-worlds; or
2. some A-world where *C* is the case resembles the actual world more than any *A*-world where *C* is not the case.

In this way, Lewis defines causal dependence between single events in the following way: "Where *c* and *e* are two distinct possible events, *e causally depends* on *c* if and only if, if *c* were to occur *e* would occur; and if *c* were not to occur *e* would not occur. (1973a)"

Lewis' theory is a useful way to judge the relative truth of counterfactuals, even without being a realist about those actual worlds. We can now use counterfactuals in our judgement of causal relationships, but they do not give us a 100% proof or certainty.

Manipulation

In the 17th century, scientist started to experiment, to manipulate processes in order to test hypotheses. Before that, it was thought to be good enough to just think about how things work, without testing in practice. Van Helmondt experimented with the spontaneous generation of mice and came up with a foolproof recipe: Leave a bag of grain in a dark place for a few days, and mice materialise out of the grain. In this process, the grain is replaced by the mice.

The manipulationist view on causality is that whenever we do an intervention in order to get a certain effect, we are manipulating a cause-effect relationship. The fact that we can manipulate a process is taken to be enough proof that the process we manipulate causes the effect we get.

Cartwright says that experiment is often our only chance to find out what the true cause is of an effect "We make our best causal inferences in very special situations—situations where our general view of the world makes us insist that a known phenomenon has a cause; where the cause we cite is the kind of thing that could bring about the effect and there is an appropriate process connecting the cause and the effect; and where the likelihood of other causes is ruled out. This is why controlled experiments are so important in finding out about entities and processes which we cannot observe. Seldom outside of the controlled conditions of an experiment are we in a situation where a cause can legitimately be inferred. "

Woodward's motivation for manipulating causes seems to be somewhat different than the search for truth. His motivation is that we spend an awful lot of our lives learning about causality, and what is the point of us having a notion of causality, as opposed to a notion of correlation otherwise than to have control? Philosophers (Cartwright for example, and

Reichenbach) say it is because causation plays a role in explanation and understanding, and according to Woodward, they downplay the practical gains.

Knowledge of causation is widely spread among humans and animals and occurs early in their development. Woodward gives the example of rats, who know that nausea can be caused by food (Garcia et al. 1966). Garcia et al. exposed the rats to radiation that caused nausea, after having fed them. The rats learned to avoid the food. When they exposed the rats to radiation that caused nausea after giving the rats light flashes, the rats did not learn how to avoid the light flashes. The conclusion by Garcia et al. is that the rats saw a connection between nausea and food, but not between nausea and light.

Woodward uses this example to prove that it is not just intellectual curiosity that makes us want to understand causality. It is the desire to manipulate things. We would never have developed notions of causation if we had not been capable of manipulation. In the case of the rats I do not agree that it is a good example, the rats were not dreaming of to manipulating things, they were probably acting on an instinct to avoid physical danger. If I encounter a giant, angry tiger, I will probably run away. I would not call that an attempt to manipulate the tiger. I would call it a desperate, instinctive and probably useless reaction. In the case of the rats, this experiment seems to point to some instinctive feeling that connects food to stomach, and stomach to nausea. Maybe the rats connect the food to the nausea, not because they have been sick because of food before, if that were the case, they could also learn a correlation between light and being sick, but because eating always gives a certain physical sensation in their stomach, where light does not. It could also mean the rats have an instinct about nausea that is not learned, but inherited. In that case, the rats' understanding of causal relations does not play a role.

The manipulationist pigeons in Skinner's superstitious-pigeon-experiment, however, could be said to do whatever they could think of to manipulate their food. Here, the pigeons had no instinct about what the mechanism was that caused them to get fed, and they developed unrelated behaviour. Superstition is a way of trying to manipulate the events we have no influence on and no instinct about.

If we were only concerned with prediction, Woodward says, it would be enough to know that whenever A occurs, B will follow. They don't have to be causally related for that. But I wonder, doesn't that also count if we want to manipulate things? Knowing whether a relationship is causal or not is only important for understanding, not for results of manipulation. For explanation it would make a difference to know whether something is merely a correlation or a causal relation, but not for manipulation.

By observation alone we can learn about correlations between two events. Manipulation is only a way to speed up the process of observation, by creating opportunities for observation that we would otherwise have to wait for, until they occur accidentally.

Skinner's pigeons only saw a correlation between the arrival of the food and the movement they were making at that point, they did not worry about the question whether or not the relationship was causal. They didn't seem to worried about any mechanism or law or explanation. The correlation was a good enough reason for them to develop certain manipulating behaviour, and as far as they could tell it worked.

Woodward (2003) says a causal explanation shows how what is explained depends on other, distinct factors, where the dependence in question has to do with some relationship that holds

as a matter of empirical fact, rather than for logical or conceptual reasons. This is a more practical approach, that, in my opinion, ignores some important conceptual problems we have with causality.
Causal relations are suitable for control and manipulation. They serve the purpose of explanation better than mathematical explanations. “In the theory of explanation, generality is not always a virtue.” Woodward seems to agree with Cartwright on this subject. The guiding assumption in Woodward's theory is that an adequate theory of causation and explanation should make clear how causal and explanatory information differs from mere description.
An important question is how a causal relation differs from a non-causal relation. We can determine a non-causal relation by looking only at the pair of events that are related. Is this font larger than this one? We measure both and determine which one is the largest. We do not need to study any other font sizes. We only compare a certain parameter that occurs in two instances of a font. To determine a causal relation however, we need a long history of encounters with similar events, to figure out a causal relation. If we see a stone hit a window for the first time in our lives, and have never before seen anything brittle break when hit by something hard, we may not see the causality between the window breaking and the stone hitting it.

Manipulation and theoretical science

The practical approach of Woodward's theory suggests that it is more suitable for experimental science that for theoretical science. But Woodward says it is applicable to both, because the ability to intervene with nature has become more important in the scientific understanding of nature, and theoretical and experimental science are deeply intertwined.

The fact that there are certain things that you cannot manipulate, not even in principle, is no reason to discard of the manipulationist theory in theoretical science (Hausman 1998). If you would treat causality in experimental science and theoretical science differently, that would mean there would be two distinct notions of causality, one for experimental science and one for theoretical science. This distinction would be hard to maintain if, due to some technological progress or some new way of looking at things, maybe a paradigm shift, a certain manipulation would become possible. Would the definitions in that case suddenly become the same? There has to be a definition for both experimental and theoretical science at the same time, and Woodward thinks the manipulationist method could work for both.

The manipulationist view on causality can be used in many scientific contexts. It extends into areas where there are no practical possibilities of manipulation. It is not practically possible to influence a planet's orbit, but we can safely say that if we would pick up the moon and place it twice as far from the earth, that would influence the tides of our oceans. But how about events that cannot be manipulated, not because of practical reasons, but for instance because they are in the past?

A plausible manipulationist theory will not deny that reliable causal inference is possible without experiment. There can be moral and practical reasons not to be able to do experiments. We can solve this by thinking of what a good experiment could be, without actually carrying it out, and then using non-experimental data to test our hypothesis.

Due to our limited capability to manipulate things in practice, manipulationsists rely heavily on counterfactuals. Woodward ignores the problems of counterfactuals by claiming that counterfactual dependence is no problem if the counterfactuals are properly understood.

However he does not give any criteria for determining when something is properly understood. He seems to find it good enough if things are understood in a common sense kind of way, but the purpose of this thesis is to look critically at this common sense understanding. He says an explanation should make it possible to answer a what-if-things-were-different-question. I can see no way to do that except by using counterfactuals.

Uses of manipulationist theory

Woodward says manipulation is the only way to get to explanation. He seems to think that you can only manipulate things if you know what you are doing, and that manipulating something will automatically give you a correct causal explanation. But you can also manipulate things if you don't know what you are doing. You can accidentally manipulate things without even knowing it, or you could be playing with correlations and drawing incorrect conclusions from your experiment. Galvani accidentally discovered that when he touched a frog's leg with two different metals, the leg would twitch. He experimented with it and came to the conclusion that he had found animal electricity, which he took to be the life force of the frog that was stored in its muscles. His manipulations did not give him any causal explanations of why and how touching the leg with the two metals could revive the frog's life force (Bresadolaa 1998).

Woodward, Carnap and Cartwright agree that non-causal relationships can be of great help in finding mechanisms. Non-causal correlations are still useful for finding the mechanism behind a phenomenon. Cartwright says correlations are important, not because they tell us what is going on, but they can help us find out what is not going on.

What does causal knowledge bring us that correlational knowledge does not? Woodward says it gives us control.
Neoclassical economist Kevin Hoover (1988) gives the following definition: A causes B if control over A renders B controllable. But there are many examples where this is not the case. We will later get to the field of biology, where, if you take away a cause, feedback mechanisms will create a new cause. The effect will still be there, but the cause is not. So control over the cause will not do anything to the effect. Yet, it is a true and only cause. Hoover's definition might work for a very limited set of simple causes. As soon as there are multiple causes at play, or more complex interactions, it no longer works.

Woodward distinguishes between practical and impractical theories of causation. An impractical theory is for instance David Lewis's counterfactual theory of causation, and the conserved quantity theory of Salmon (1984) and Dowe (2005): causal interactions involve the intersection of two or more causal processes, and the exchange of some conserved quantity, such as energy or momentum, as when two billiard balls collide. But this fails to describe the practical purpose of being able to distinguish between interactions that exchange momentum and those that do not. Woodward insists that there must be something that gives us more benefit. He wants a practical theory, like the manipulationist account.

The manipulationist view on causality does not add much to the discussion about causality. It is very much focussed on practical purposes, and Woodward almost seems to assume that causality is a tool that is there for us to be able to manipulate the world. He has not convinced me that manipulation gives us more insight into causal relationships than observation. The only advantage I can see is that it saves us time if we don't have to wait for an event to occur many times so we can make enough observations, but we can actively make it happen.

Manipulation and the direction of causality

Maybe we can use the notion of manipulation as an alternative solution for determining the direction of time and causality? If we can only change the future and not the past, that would mean that out interventions are limited to this future. If we have an event that we can interfere with, we will know that the effect will be in the future, and the direction of causality is forward.

Reichenbach shows that we are wrong. If we make an intervention, for instance we throw a ball from A in the direction of D, and we deflect the ball at point B in the direction of C, we can make the counterfactual assumption that if we had not intervened, the ball would have ended up at D. We can check this by throwing the ball another time, or a billion times, without intervention, and we see that indeed, it does end up at D every single time. We spend some time on throwing the ball a billion times with intervention, seeing it end up at C, and we think we have proven the counterfactual (see figure 12).

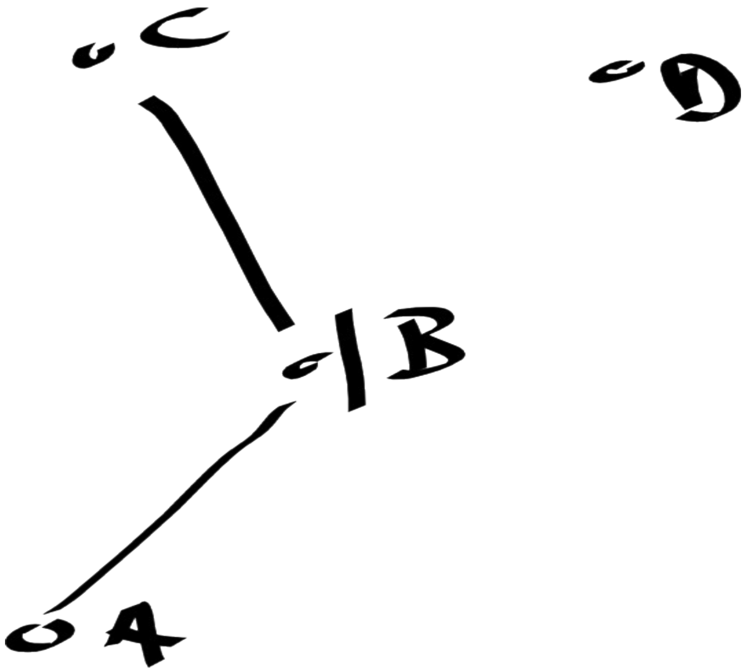

Figure 12: Tennis ball gets thrown at position A in the direction of D, and gets reflected at B in the direction of C.

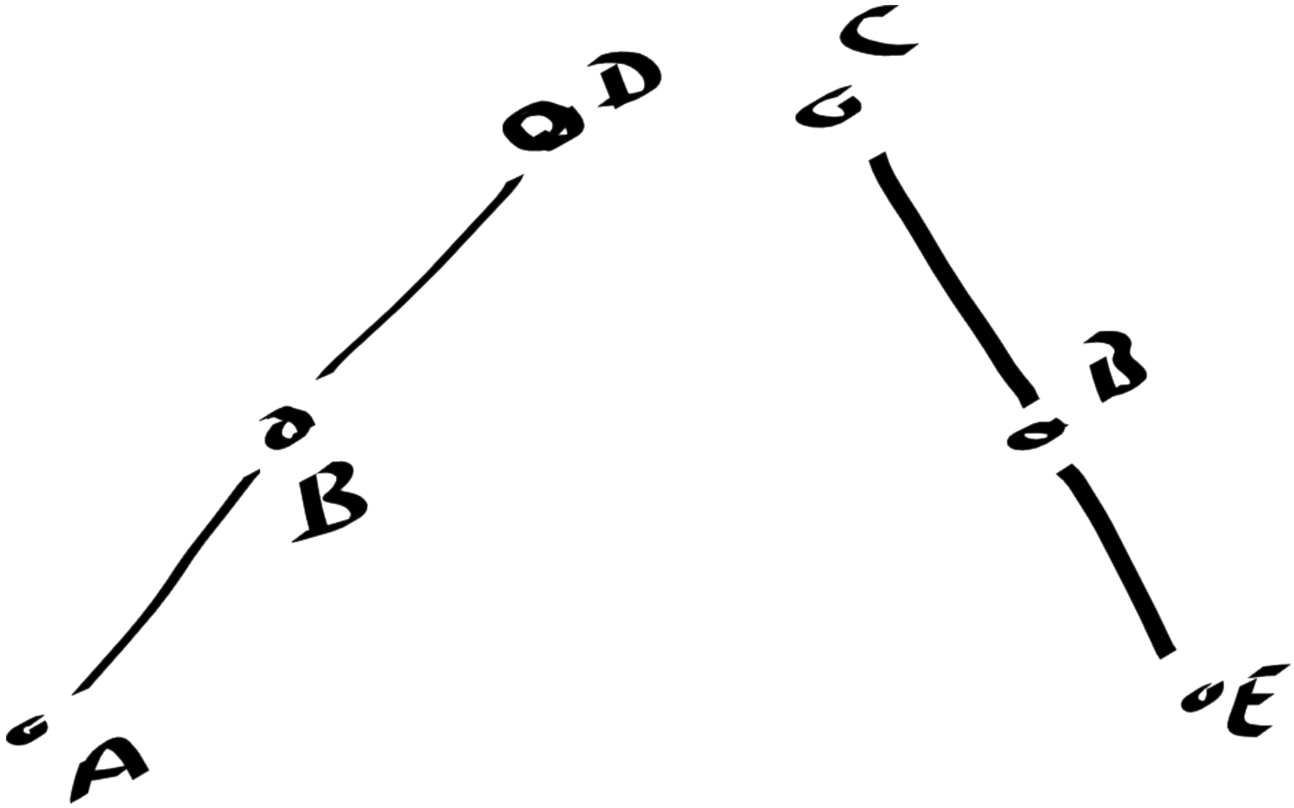

Figure 13: Some alternative paths of the tennis ball

But according to Reichenbach, we have made a mistake. We have assumed we can only leave the path AB unchanged. Reichenbach wonders if we can also design an experiment leaving the second part of the path unchanged. If we throw a ball from point E, and not intervene, it will also end up at C (see figure 13). We have introduced an alternative cause.

Looking from point B, the first part of the path represents the past, and the second part of the path represents the future. In designing the experiment so that the intervention only changes the second part of the path, we have assumed that we can only change the future with the experiment.

We cannot find a direction of time by using an intervention. In a closed causal chain it is not impossible to change the past, it is only impossible to change both the past and the future.

Manipulation does not give us a watertight excuse to assume the direction of time and causality is forward.

Determinism

In a way, classical mechanics has the simplest use of causality. It usually deals with simple and single collisions, that are idealised and completely deterministic. But this determinism, assuming it exists, raises some problems.

Although we cannot be certain of the future, we can use our understanding of the laws of physics, as well as our understanding of causality, to make predictions. Sometimes they are accurate, sometimes less so, but we feel that if we would know more about the state of the world and the laws that govern it, we would do much better in our predictions. Causal determinism says that the future is completely determined by causal laws. If we were able to know the exact state of every bit of the universe at a single moment, we could predict with absolute certainty the future up to infinity (Laplace 1814). Newtonian physics showed that natural processes can be predicted with high accuracy, using mathematics.

Kant saw that determinism was a danger to human freedom. If the future is completely determined by the present state of the universe, we have no way of influencing the processes around us. Kant tried to solve that problem by suggesting that the flow of time is something we perceive, but that is not a real feature of nature. He stated that the same goes for causality. In a reaction on Hume's problem of induction, Kant said we use the notion of causality to make sense out of the processes we observe, but it is not a property of nature (Kant 1783). I do not think he solves the problem of freedom by denying the existence of causal relations, because if we cannot use causal relations to influence a process, we are again limited to helplessly observing the world around us as we have no tools to influence it.

Probability

It is now assumed that the world is mostly indeterministic. This means that where for a deterministic system a change in cause A will always result in the same change in effect B, for indeterministic systems that is not necessarily so. A change in A might result in a change in B, with a certain probability. We call this probabilistic causality: the cause raises the probability of an effect. This in turn means that for some events there is no sufficient cause, because a sufficient cause is sufficient on its own to produce a certain effect.

But are there then events that are uncaused? As soon as the effect has occurred we can say that apparently there was a sufficient cause, otherwise there would not have been the effect. In the many worlds theory, however, there are worlds in which the effect has not occurred, and in which the same cause was not sufficient. Here, the same cause both be sufficient and not be sufficient. We could say that in an indeterministic world, there are only contributing causes, and that the term sufficient cause is not appropriate.

A way to formulate probabilistic causation is to say an event is a chance raising cause if it raises the chance of a certain effect to occur. A chance raising cause can be represented as follows:

$P(B|A) > P(B|!A)$

meaning the chance of an effect B occurring when cause A occurs is greater than the chance of an effect B occurring when A does not occur.

Most papers about probabilistic causation cover mainly chance raising causes. A paper by Glynn (2011) also covers some examples both of non-probability-raising causes and of probability-raising non-causation.

Glynn gives the example of a sudden drop in the reading on a barometer that can be said to be a probability raising non-cause. The chance of a storm (B) occurring when a low reading on the barometer (A) occurs, is greater than the chance of a storm occurring when a ow reading on the barometer does not occurs. Yet the barometer does not cause the storm. Both A and B are caused by a sudden drop in air pressure which is the common cause for both. There is a higher chance of a storm when the barometer gives a low value, but the reading is not a cause.

Another type of probability-raising non-causation follows from the chance raising cause, because instead of

P(B|A) > P(B|!A)

we can also say

P(A|B) > P(!A|B)

where the effect raises the probability of the cause. According to Glynn, that would imply the effect is the cause of its cause, making any case of non-causation a case of causation. I think that is only true if we ignore the temporal order of the events. A strategy to solve this is by keeping some background conditions fixed. This results in the following definition of a cause: A is a cause of effect B iff A raises the probability of B when certain background conditions are fixed.

Reichenbach (1956) suggests we can say that, if t_A is the time at which event A occurred, then any common cause of events A and B will have already occurred at t_A, and both A and B would be part of the fixed background. Also, if the presumed effect B had already occurred at t_A, before the cause A, it would be part of the fixed background conditions and would not be an effect of A.

Rosen (1978) gives an example of a non-probability raising cause: a golfer who hits a ball rather badly, but because the ball bounces off a branch of a tree, it ends up in the hole. The ball hitting the branch is a non-probability-raising cause of the ball ending up in the hole. The probability of the ball ending up in the hole after hitting the tree is smaller than the probability of the ball going into the hole without hitting the tree. I don't agree on that, because if you look at the events before the ball hit the tree, and see those as a fixed background, the fixed background includes the very bad aim of the golfer, and the trajectory of the ball that makes it very unlikely that the ball enters the hole without hitting the branch, then the ball hitting the branch increases the probability of it entering the hole.

The question is if deterministic systems are not just limiting cases of indeterministic systems, in which the chance of a cause to cause a certain effect is so large that we never see the case where the effect does not occur. That would mean that there are no sufficient causes, ever. In a many worlds theory, we can say that if for every possible outcome of a series of events there is a corresponding world, any total cause is always a sufficient cause. But for for this use of the many worlds theory, one has to be realistic about the many worlds, and for most people that may not be easy.

We should wonder if we can treat this world as a causal world at all. Maybe Woodward was right, and we have to be more practical about causality, and ignore some of the conceptual problems, to get to a workable definition.

Definition

We have seen a lot of problems with causality, many of which have not been completely solved. To be able to continue working with the concept, we need to make choices about what problems we can safely ignore under every day circumstances.

The direction of time and causality has not been completely settled as being the direction of increasing entropy. It could be that the future of the universe will give us decreasing entropy. However, in the universe as it is today we can safely assume the entropy increases. By the time the universe has reached its maximum size and starts to shrink again, we will have to reconsider, but let's make the temporal assumption that the direction of causality is the direction of increasing entropy.

The induction problem has not been solved either, but we can live with it in the same way as we can live with the problems with counterfactuals. Lewis has taught us the trick of considering other possible worlds, in each world one of the options is realised, and order them according to how closely they are related to our world. The world that is closest to ours, has the most likely option. We can never be completely, logically, sure, but we can be practically sure. This gives us an excuse to ignore the problems with counterfactuals in most cases, as well as the induction problem.

Lewis' (1973b) definition of a cause is: “a cause is any member of any actual set of conditions that are jointly sufficient, given the laws, for the existence of the effect. We allow a cause to be only one indispensible part, not the whole, of the total situation that is followed by the effect in accordance with a law.”

This is the best definition I have found so far, and we can use it as a basis for our own definition. Now let us modify it.

In the chapter about manipulation we have seen that there are situations in which when a cause gets eliminated, another cause gets 'created'. So in this case the first cause was not indispensible, yet before the intervention, it was a true and possibly even only cause. In more complex situations there are causes that are not necessary, like chance increasing causes or redundant causes. An example would be “smoking causes lung cancer”. You don't have to be a smoker to get lung cancer, but it does increase the chance to get it.

So the first adjustment of Lewis' definition would be taking out the demand that a cause be indispensible. The definition would become:
“a cause is any member of any actual set of conditions that are jointly sufficient, given the laws, for the existence of the effect. We allow a cause to be only one part, not the whole, of the total situation that is followed by the effect in accordance with a law.”

Now the definition suffices if we know what the cause is, but in case we don't know, we have to check for plausibility. We can use the many worlds trick, combined with the demand that the event being a cause is logically and practically possible, to determine whether or not a candidate cause is a plausible one. In this case, we would order the possible worlds, all having different causes, in accordance with how closely they are related to our world. The plausibility will only give us an increased likelihood of the event being a cause.

Using this version of plausibility, another iteration of the definition could be:

“a cause is any member of any actual set of plausible conditions that are jointly sufficient, given the laws, for the existence of the effect. We allow a cause to be only one part, not the whole, of the total situation that is followed by the effect in accordance with a law.”

A remaining problem is the status of laws. If we listen to Carnap, we should add: “To say event A caused event B means that there are certain laws of nature from which event B can be logically deduced when combined with the full description of event A.” I will incorporate that and rewrite the definition to make it less ugly:

“a cause is any member of any actual set of plausible conditions that are jointly sufficient to produce the effect. That the cause leads or contributes to the effect has to follow logically from the relevant laws.”

I feel we now have a definition that can work for most cases. We have incorporated a lot of the problems we have encountered. There is one problem we have not dealt with in this definition, and that is the problem of indeterminism. We will have to assume that the macroscopic situations we will use this definition for are roughly deterministic. I don't expect to run into any other conceptual problems in examples of practical cases.

Types of causes

Now that we have a definition, we can see what types of causes we can distinguish, and some examples of those causes.

Common cause

There can be causal confusion when two events always occur together in a certain order. We tend to assume that the one is the cause of the other. But they could also both be caused by a third event, that occurred earlier. This event is the common cause, sometimes called mutual cause, of the two events we observe.

A is a common cause of B and C if B and C are both caused by A.

Example: Every time a barometer shows a drop in its value, a storm follows. The storm is not caused by the barometer showing a decreasing value. Both are caused by the drop in air pressure.

Spurious cause

A spurious cause is an event that is taken to be a cause but later turns out not to be a real cause.

A is a spurious cause of B if, with hindsight, A turns out not to be the cause of B.

Example: The low value on the barometer in the previous example is a spurious cause of the storm.

Sufficient cause

A cause is a sufficient cause if the event alone is enough for the effect to occur. This does not mean it has to be the only cause. There can be many causes to a single effect that can all be sufficient causes. We will get back to the sufficient cause, the contributing cause, the necessary cause and the replacement cause in the section about epidemiology.

A is a sufficient cause iff a change in A or the value of A, and only A, results in a change in B or the value of B, or the probability distribution of B, if other variables remain fixed.

Example: A sufficient cause of the ripples on the surface of the water is the stone that just got thrown into it, if we keep other conditions fixed, like the conditions that the temperature of the water is above its freezing point but under its evaporation temperature, the water is located on a large body of gravity etc.

Contributing cause

A cause is a contributing cause if helps produce the effect. If there are many causes, not all of them are sufficient to cause an effect on their own. But some combined can be sufficient. These often only increase the chance to get an effect.

A is a contributing cause of B if a change in A produces a change in B, when other variables are fixed.

Example: Imagine you ride your bike on a cold day, using only one hand because the other one is operation your cell phone. A sudden gush of wind throws you off balance exactly at the moment you run into an icy patch of street. Your tires don't have much grip because you just put some extra air in. You slip on the patch of ice and get run over by the lorry behind you. The ice, having only one hand available for steering, the wind, the lorry and the high tire pressure are contributing causes of death.

Total cause

If a cause is a total cause of an event, it is the only (direct) cause to the event.

A is a total cause if a change in A will change B or the probability of B.

Example: This time you ride your bike on a cold day, using only one hand because the other one is operation your cell phone to check the stock markets. It is again windy and the streets are icy. A meteorite the size of the moon lands on your head. The meteorite could safely be called a total cause of death.

Necessary cause

A necessary cause is a cause without which the effect would not occur. This is not necessarily a total cause or even a sufficient cause.

A is a necessary cause of B if there is a possible change in A such that only if this change were made, B would change.

Example: a stone gets thrown into a pond and causes ripples in the water. There are no fish in the pond, no wind, nothing else to cause ripples. Without the stone being thrown, there would be no ripples. The stone being thrown is a necessary cause.

Direct/indirect cause

A cause can be direct or indirect. Every step in a causal chain is the direct cause of the event it is directly connected to. If A causes B, it is a direct cause of B. But if A causes B, and B causes C, then A is an indirect cause of C.

A is a direct cause of B when changing A will result in a change in B, even if other variables remain the same. All other causes are indirect causes.

Example: If a gun is fired in the air and the bullet hits a rope that has a concert piano hanging from it, the concert piano falls down, landing on a person and killing him, the gunshot is an indirect cause and the falling piano is a direct cause of death.

Replacement cause

An replacement cause is somewhat special, because it only occurs in processes with feedback loops. There are cases in which, for instance through an act of manipulation, cause ceases to produce a certain effect and another cause will get activated that does cause the effect. This does not mean there has to be a conscious being that decides the effect has to occur. Also, the effect can be a completely useless byproduct, there may not even be a need for the effect for the system to function. Figure 14 shows how a replacement cause can work. Under normal circumstances, a cause A would cause effect B. If A gets deactivated, effect B no longer occurs. The lack of effect B triggers a replacement cause A' to produce a new instance of the effect, called B'. We could call B' a replacement effect.

I have not been able to find a name for this type of cause in the literature, so for now I will call it a replacement cause.

C is a replacement cause of B if C only becomes a cause of B after another cause A has been removed.

Example: a certain gene causes an effect. If a biologist switches off the gene, another gene takes over the function of the first. This second gene did not cause the effect before the first got switched off.

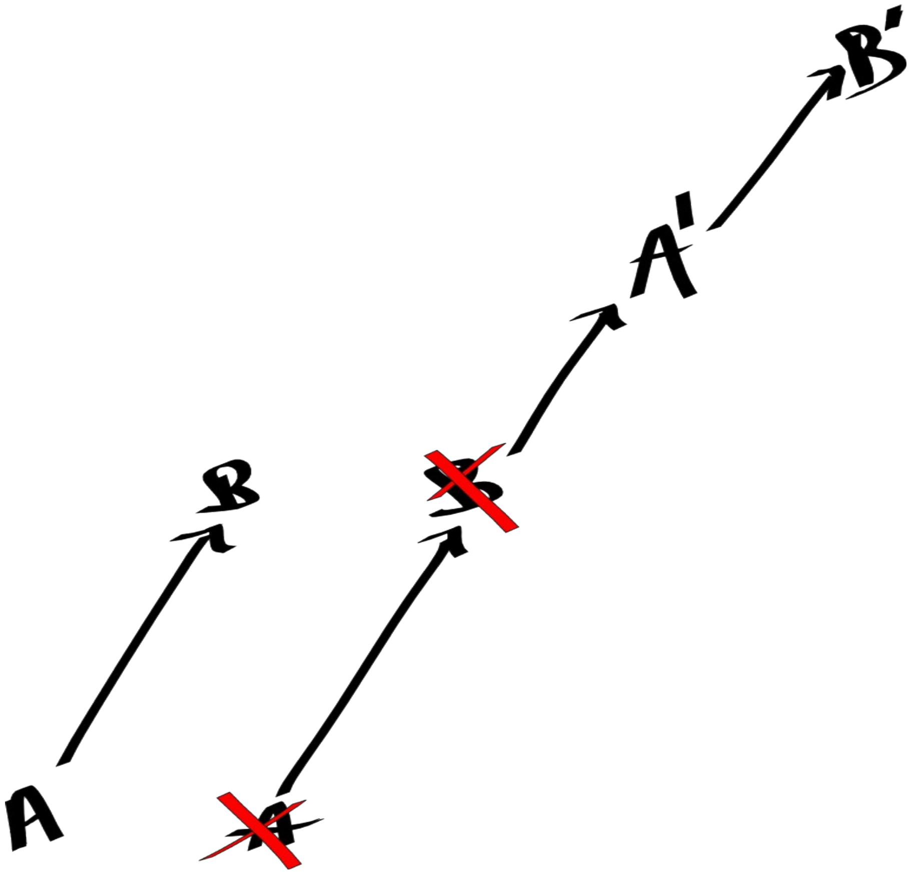

Alternative cause

An alternative cause is either the product of a thought experiment, or the result of lack of knowledge about the system.

C is an alternative cause of B if either A or C is a cause of B, but not both.

See the paragraph about the direction of time.

Example: in figure 13, E is an alternative cause for the tennis ball ending up at C, if there is no intervention at point B.

Connections between types of causes

There are many connections between different types of causes. A cause is never of one single type.

A total cause will always be a sufficient cause, except in an indeterministic world. However, if we assume there are many worlds, one for every possible outcome of a series of events, we can say that in the world in which the effect occurred, the total cause is a sufficient cause. That would solve the problem in every world.

A direct cause will always be a contributing cause, but not all contributing causes will be direct causes. An indirect cause can be a necessary cause but not a sufficient or total cause, remember the example of the gunshot and the piano.
Contributing causes can be sufficient causes or non-sufficient causes and they can either be necessary or not.

Causal events can change their cause-type in interaction with other events. A cause that used to be a total cause can become a mere contributing cause if another cause comes into play. Imagine for instance a single slit experiment, where photons go through a slit and leave a mark on the photographic plate behind it. The single slit is the total cause of the pattern that the photons make on a photographic plate. If we add a second slit, an interference pattern will emerge. The two slits together are the sufficient cause of the pattern. The first slit went from a total, necessary and sufficient cause to a contribution, necessary but not sufficient cause. And if we say the first slit is part of the historic background conditions that are fixed, we can say the second slit is a contributing, necessary and sufficient cause, given the fact that the first slit is already there. Now we add a third slit. The interference pattern will now be caused by three slits, each being contributing, non-sufficient and non-necessary causes (unless we fix the first two slits in the historic background conditions).

The next chapter is a collection of examples of causality in different scientific disciplines, and linked to the cause types above.

Use of causality per area of science

In some areas of science it is completely impractical to take certain conceptual problems into account. If we try to find the causes for a lethal disease, we will not put too much effort into determining the direction of time.

This chapter is about what definitions and causal types are used in different disciplines of science. It is not my intention to be complete, just to give some experience with thinking about causality on different subjects.

Complex systems and climate change

In many complex systems we see sudden transitions. A variable exceeds a certain critical threshold and the whole system completely changes its behaviour. In many of these systems, there is a hysteresis that prevents the system from going back into its previous equilibrium when the value of the variable drops back to the previous value. It has to decrease a lot before the system goes back into the first equilibrium.

An important example is the case of the possible change in the thermohaline circulation due to climate change. The thermohaline circulation is the circulation in the ocean that is caused by gradients in salinity and by the heating of the surface of the water. Salty water is denser, and therefore heavier, than fresh water, and sinks. This process is called deep water formation. Warm water is less dense than cold water. Together these factors, combined with wind, drive the thermohaline circulation (see figure 15).

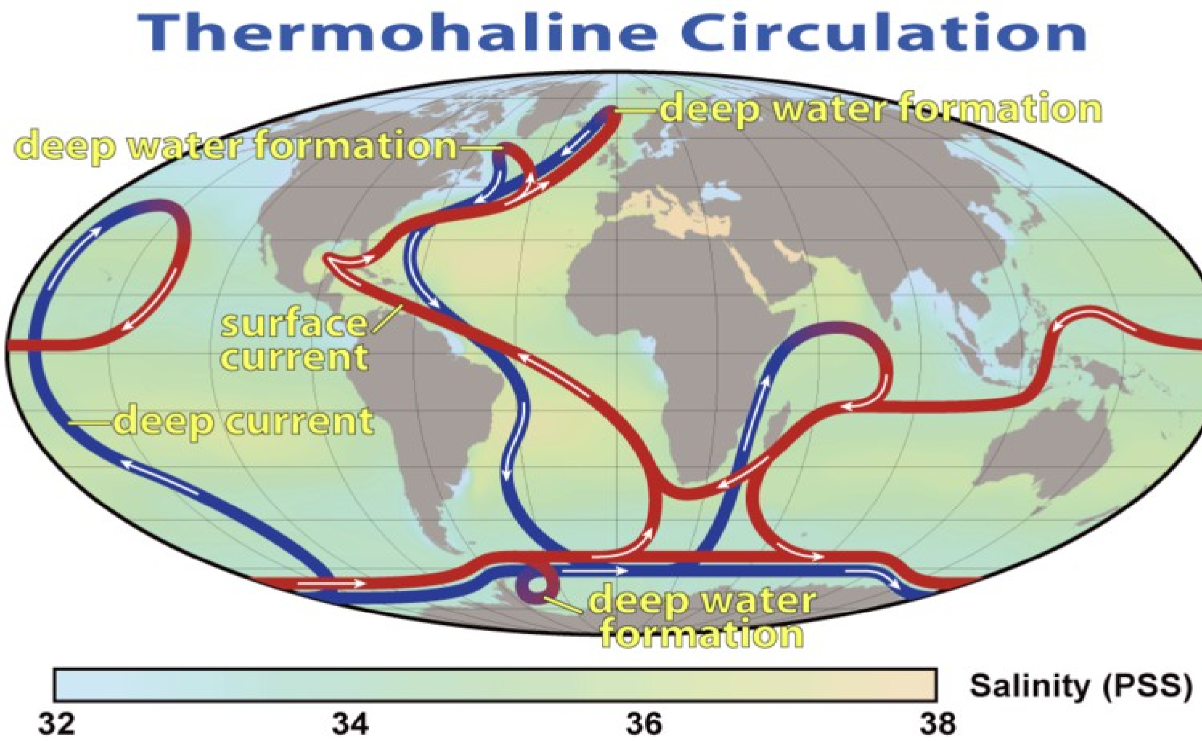

The thermohaline circulation is in danger of changing its intensity, or even reaching a complete shutdown, because of large amounts of fresh water entering the North Atlantic due to the melting of huge amounts of Greenland ice. When the salt level decreases to a certain critical value, a tipping point is reached and the northern deep water formation stops. There is a hysteresis (figure 16), so we cannot just add a pinch of salt to switch it on again.

The thermohaline circulation has a large influence on the temperatures on land on the mid and higher latitudes of the globe.

In the case of a complex system, there are many variables at work that act as contributing, non-necessary and non-sufficient causes. When the sum of the causes becomes sufficient, a transition takes place. Due to hysteresis, we cannot reverse the transition by getting the variables back to their initial values.

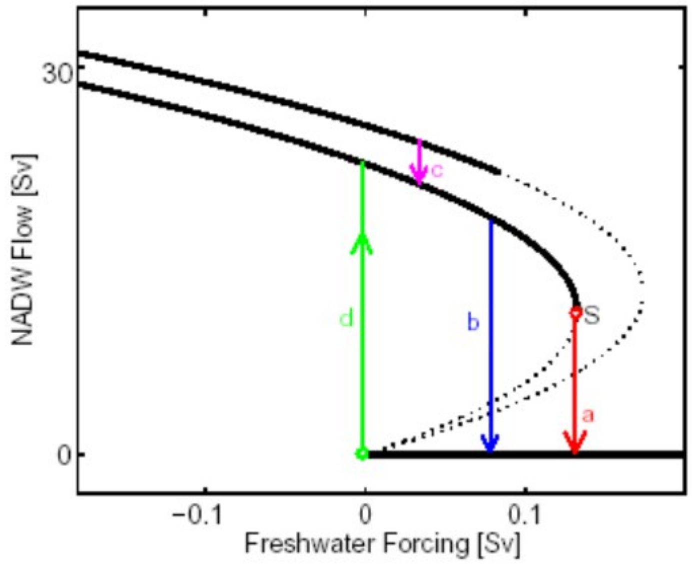

Schematic stability diagram of the Atlantic thermohaline circulation. The two upper heavy branches indicate the possibility of multiple states with different convection sites. Possible transitions indicated are: (a) advective spindown, (b) polar halocline catastrophe, (c) convective transition, (d) start-up of NADW formation. From Rahmstorf 1999; see Stocker and Wright 1991a for an earlier, similar diagram.

In this case, the cause can not be called an event. There is no single ice berg that drops into the ocean, but a slow and continuous process. The cause in this case is a process changing variables that eventually reach the tipping point. Reaching the tipping point is an event, but not a cause.

Statistical Physics

In a system that has many actors that together form the cause of a certain effect, for instance a gas in a container, where all the molecules that hit the wall of the container make the pressure in the container, we no longer look at independent agents. Rather we look for equilibrium explanations.

Equilibrium explanations can be called causal, even though we do not look at each individual event in the chain. We can manipulate the pressure in a container of gas by raising or lowering the temperature. Adding energy to the system will cause movement of the molecules, which raises the probability that the molecules hit the walls, which in turn causes an increase in pressure. We no longer say the individual molecules that hit the wall cause the increase in pressure, rather we say that the energy increase causes it, or even the act of lighting the candle under the container.

Elliott Sober (1983) identified some noncausal equilibrium explanations, in which an outcome is determined by the fact that there are many initial states that lead to the outcome state. He says this is in contrast to causal explanations that do track the chain of events.

In statistical physics there are also systems that have tipping points and hysteresis. An example is a phase transition from a metastable gas to a liquid, and from a metastable liquid to a gas (figure 17).

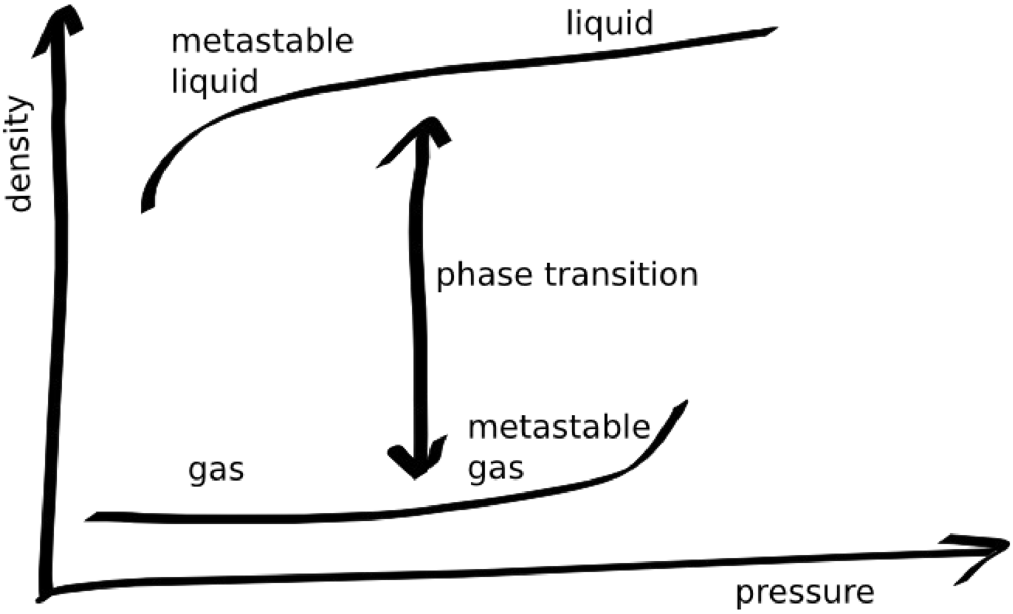

Figure 17: Phase transition between gas and liquid, with hysteresis

Epidemiology

Diseases are caused by complex interactions of several causes. Hardly ever is a disease caused by a single event. Epidemiology needs a model of causation that describes causes in terms of sufficient causes and their component causes. Multi-causality is one of the most important concepts.

Rothman and Greenland (2005) propose to view causal inference in epidemiology "as an exercise in measurement of an effect rather than as a criterion-guided process for deciding whether an effect is present or not."

Most diseases are caused by multiple events, where most causes are neither necessary causes nor sufficient causes. The causal chains can be extremely complicated. Rothman and Greenland give an example of a person who suffers a head trauma, which has no obvious consequences, except for many years later. On a winter day, the person walks on an icy path and falls. The causes of the fall are the ice on the path and the balance of the person that is slightly affected by the old head injury. The authors claim there is almost always some component cause that has to do with the person's genetic make up or environmental factors. The old accident with the head was an indirect cause of the fall, despite the many years that were between the first accident and the fall. In epidemiology, contributing causes can accumulate for a very long time before they maybe result in a disease. More often than not, there are no sufficient causes and no necessary causes. The statement “smoking causes lung cancer” is problematic. Smoking is neither a sufficient nor a necessary cause for lung cancer. Many other things can cause lung cancer and not every smoker gets it. Also, the term smoking is not binary. Is somebody who smokes only once a smoker? How about somebody who used to smoke for a year when he was 16? Do we have to distinguish between pipe smokers, cigarette smokers and cigar smokers? Are smokers who start smoking early in their lives more vulnerable than people who start at an older age?

Rothman and Greenland define a cause of a specific disease event as “an antecedent event, condition, or characteristic that was necessary for the occurrence of the disease at the moment it occurred, given that other conditions are fixed. In other words, a cause of a disease event is an event, condition, or characteristic that preceded the disease event and without which the disease event either would not have occurred at all or would not have occurred until some later time. ” That has similarities with the last and counterfactual part of Hume's definition “where the first had not been, the second would not have followed“.

They use the term sufficient cause in a different way from the definition above. A sufficient cause is not a single event, but a set of contributing causes, none of them have to be sufficient when occurring alone, that together form a sufficient cause. The concept of a sufficient cause seems to only be used with hindsight. One cannot predict at what exact moment a set of contributing causes is enough to be sufficient. They define a sufficient cause as “a set of minimal conditions and events that inevitably produce disease”. With 'minimal' they say they mean that all events in the set are *necessary* contributing causes. At a first glance this seems to be incompatible with the example of “smoking that causes lung cancer”, where we have just stated that there are many contributing causes for lung cancer, none of them being a necessary cause. Yet, if the cancer occurred, there must have been a set of contributing causes that together formed a sufficient cause. They solve this confusion in a diagram (figure 18)

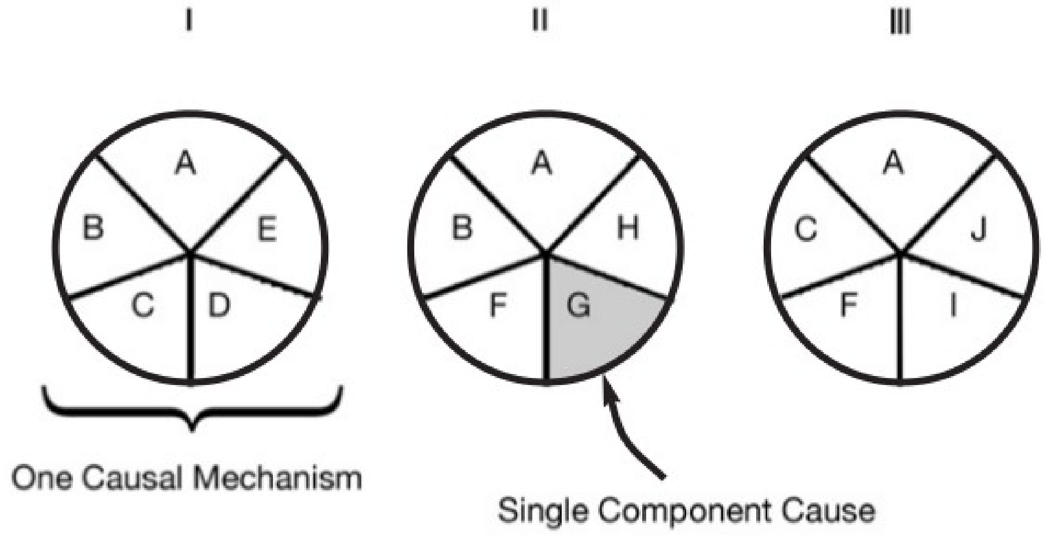

Figure 18: The sufficient cause as defined by Rothman and Greenland (Rothman and Greenland, 2005)

All three diagrams are sets that together form a sufficient cause. There is no redundant cause, every cause that would be redundant is moved to a new diagram. So in every diagram there are only necessary causes left. Some causes turn up in all of the diagrams. This can be a clue on how to solve the problem. If it is possible to eliminate the causes that are in every diagram, in this figure that would be cause A, eliminating those could be a good start.

Box 1: The Bradford Hill criteria

In epidemiology, attempts have been made to assess the strength of a causal relationship. Sir Bradford Hill (1965) suggested that the following aspects of an association be considered in attempting to distinguish causal from noncausal associations:

- Strength of association : A strong association is more likely to have a causal component than is a modest association
- Consistency: A relationship is observed repeatedly
- Specificity: A factor influences specifically a particular outcome or population
- Temporality: The factor must precede the outcome it is assumed to affect
- Biological gradient: The outcome increases monotonically with increasing dose of exposure or according to a function predicted by a substantive theory
- Plausibility: The observed association can be plausibly explained by substantive matter (e.g. biological) explanations
- Coherence: A causal conclusion should not fundamentally contradict present substantive knowledge
- Experiment: Causation is more likely if evidence is based on randomised experiments

The Bradford Hill criteria are outside the scope of this thesis, because they are very specifically designed to be used in epidemiology. They are however a very useful tool to assess causal strength. An interesting article to read about these Bradford Hill considerations is “The Bradford Hill considerations on causality: a counterfactual perspective ” by Michael Höfler (2005).

The fact that there are always multiple causes, including genetic make up and environmental factors, can be abused by anyone who has an interest in for instance selling products or defending themselves against law suits. We will get back to that later.

Biology

In biological systems, combinations of components form a complex organisation. A complete list of parts would be hard to get, but having it would also not give you a complete overview of what every part does (Lazebnik 2002). To understand biology, you have to understand the whole, integrated system. A system can be a cell, an ecosystem, an organism, a digestive system or whatever part you want to distinguish. Those parts usually also interact, so knowing one system is also not enough.

Feedback loops play an important role in and between sub-systems. Whenever a shortage occurs, a feedback loop ensures that new supplies are brought in. Enzymes are produced that function as catalysts to speed up production processes. Most often, the mechanism is that of negative feedback (Campbell and Reece 2005). Whenever there is a surplus, that surplus slows down the future production.

A way to discover how a biological system works is by intervention. If you want to know what a certain gene does, you can deactivate the gene and see what functions cease to be performed. In other words, if you deactivate this cause, you will find out what its effect was. The problem with systems back loops however, is that as soon as a cause gets deactivated, another cause will be created. In this case, another gene will take over the task of creating the effect. The last thing I want to suggest here is that there is some kind of necessity, or a will of the system to have a certain effect. These are all mechanisms that function automatically, and are the result of a long evolutionary process.

In this case the first gene could be a total and only cause of the effect. It could also be a sufficient cause of an effect. But it is not a necessary cause. If the cause gets eliminated, another cause will be born. The second gene starts out not being a cause at all, and ends up being what the first cause was, a total, sufficient but unnecessary cause. I call this second gene a replacement cause, for lack of a better word.

Chemistry

What kind of causality are we dealing with when a mix of two chemicals causes an explosion? Are both chemicals contributing, necessary but non-sufficient causes? Or is it the act of mixing the chemicals together that causes the explosion? We can compare the process with someone who sets a building on fire. We don't say it is the oxygen that causes the fire, nor the wood of the building. It is the act of igniting that is the cause. It is the adding of enough activation energy to get the reaction started. In the case of two chemicals being mixed and exploding, the movement of the molecules at room temperature provides enough activation energy for the reaction to get over the (apparently low) energy barrier.

Are catalysts contributing but non-sufficient (and non-necessary) causes? If you have a chemical process $A + B \leftrightarrow C$, with $A + B$ having a higher energy than C, and the reaction goes via $A^+ + B^-$, which has an even higher energy. You need to add a lot of activation energy, otherwise some molecules will react, but not a lot. To solve that problem, you can use a catalyst X to make the reaction work via a lower energy state than $A^+ + B^-$, say ABX, and increase the chance of a reaction (figure 19).

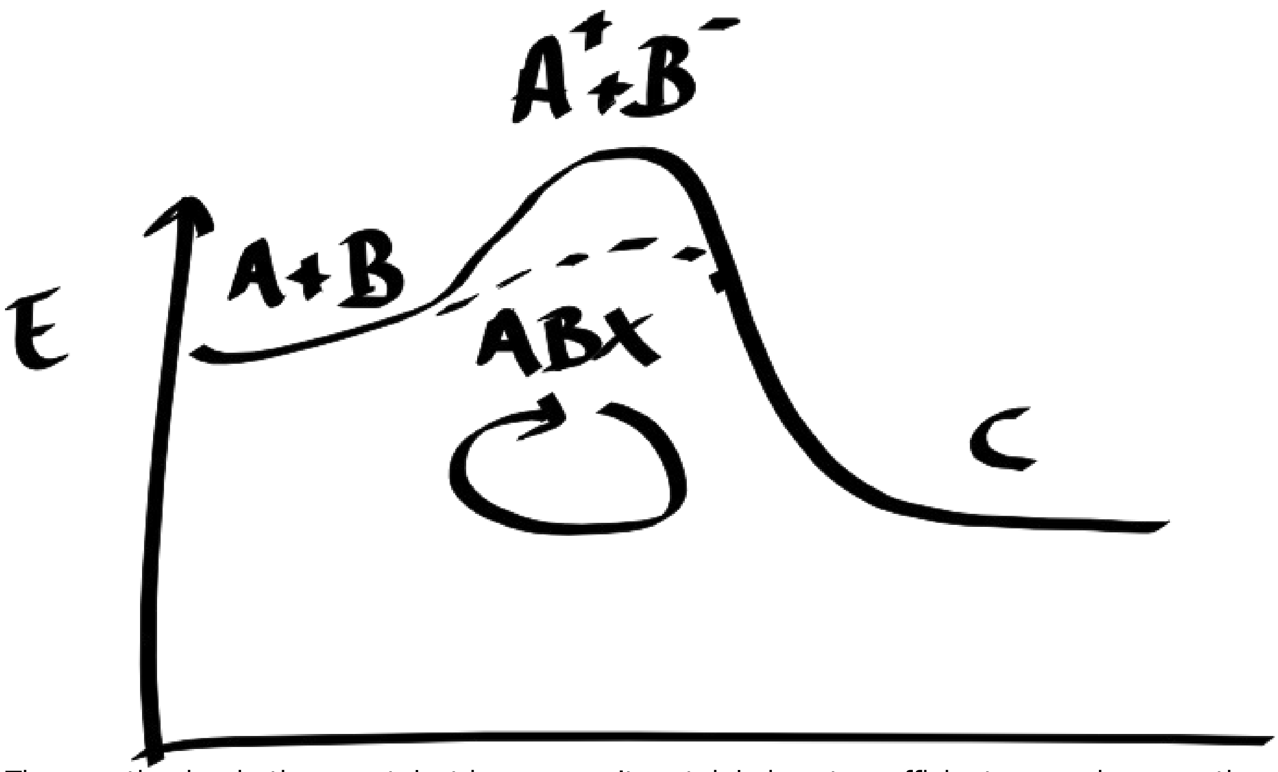

The question is whether a catalyst is a cause. It certainly is not a sufficient cause, because the catalyst alone does not do anything. We could call it a contributing cause, since it does contribute to the reaction. But is it a necessary cause? That may depend on the purpose of the experiment. If you want the reaction A+ B ↔ C to take place, then the catalyst is not a necessary cause. The reaction will take place anyway, and if you would have an infinite amount of time and patience, you could do without the catalyst. Necessary in this case means necessary for you to achieve your goal, rather than necessary for the effect to occur.

Law

In law, causation is mainly a problem in tort law (according to the dictionary, a tort, in common law jurisdictions, is a civil wrong which unfairly causes someone else to suffer loss or harm resulting in legal liability for the person who commits the tortious act).

Problems include counterfactuals (Mackie 1974) and multicausality.

If a cyclist falls on an icy path, due to a sudden gush of wind, while playing with his cell phone, can the owner of the path be held accountable for the fall, because he has not freed the path of ice? In other words, would the person not have fallen if the path had not been icy? Or would the other contributing causes have been sufficient for the fall? And how important a cause was the ice?

In law of war, there is the question of whether a person's inaction is the cause of a war crime taking place. This sort of thing happens in every war, but the most well known recent example is the Srebrenica genocide in July1995, where more than 8,000 Bosnians, mainly men and boys, were killed in and around Srebrenica, by units of the Serbian Republican Army.
The Dutchbat peacekeepers of the UN did not prevent that the Serbs captured the town and

killed the people.
In 2005 the Secretary-General of the United Nations (2005), Kofi Annan, said the following: “We can say -- and it is undeniable -- that blame lies, first and foremost, with those who planned and carried out the massacre, or who assisted them, or who harboured and are harbouring them still. But we cannot evade our own share of responsibility. As I wrote in my report in 1999, we made serious errors of judgement, rooted in a philosophy of impartiality and non-violence which, however admirable, was unsuited to the conflict in Bosnia. That is why, as I also wrote, “the tragedy of Srebrenica will haunt our history forever”.”

In this case, if we simplify the situation enormously, we could say there was one cause, the killers, but there were many contributing causes. A large organisation like the UN is too slow and inflexible for people in critical situations to take fast decisions that go against the philosophy that Koffi Annan calls unsuitable for the situation.

Ecology

In the last decades there has been a worrying trend of pollinator decline. In this case, a lot of causes play a role in populations of for instance honeybees.

Honeybees are threatened by a number of problems. They have insufficient food and too little variation in food, diseases, parasites, and they are exposed to pesticides (Van der Sluijs 2011).

Bees feed on pollen and nectar. Due to intensive farming, it has become difficult for bees to gather enough and diverse pollen. Not many of our food crops reach the stage of flowering before they get harvested (Van der Sluijs et al. 2013).

The two most common diseases are Nosema, which cause diarrhoea, and Deformed Wing Virus (DWV) which causes bees to be born without properly developed wings. Both diseases are always around in a bee population, and under normal circumstances the bees can manage the diseases. Only when the bees get weakened, do the diseases become a real problem (Van der Sluijs et al. 2013).

The parasite that bothers the bees is the Varroa Destructor, a mite that feeds on the larvae of the bees. Affected bees are born smaller and the mite infects them with Nosema and DWV. The parasite also negatively affects the immune system.

Since two decades, a new type of pesticide is used, which enter the sap stream of the plant, making the entire plant poisonous to insects. The pesticides are called neonicotinoids. Neonicotinoids are used in very small doses, but some are 7000 times as poisonous as DDT. They end up in the nectar and the pollen of a flowering plant, and get collected by many beneficial insects, such as honeybees. The bees either die immediately or take the pollen and nectar to the hive, where it is eaten by larvae or other bees (Van der Sluijs et al. 2013).

The low dose of neonicotinoids causes sub lethal effects, which can eventually lead to the collapse of the colony. Sub lethal effects include brain damage that makes it harder for bees to find the way to their food or back to the hive, which contributes to starvation of the larvae and decline in the population by disappearing bees. The ability to communicate gets affected, which disrupts processes in the colony. It influences the bees' grooming behaviour, and they do not clean up dead or sick larvae and bees, and most importantly do not resist the mites that give them diseases.
The neonicotinoids also damage the bee's immune system and decrease the bee's life span. In summer, healthy bees only live five weeks, but in winter they have to survive for five months,

and if they have a shorter life span, they will not make it through the winter.
Spiders, ticks and mites, like Varroa Destructor, seem to thrive in the presence of neonicotinoids. This issue is still waiting to be researched.
The sub lethal effects caused by the neonicotinoids enhance all the other factors that threaten bees (Van der Sluijs et al. 2013).

The pesticide industry, Bayer and Syngenta, say the decline in honeybees is caused by the varroa mite. Independent research says it is caused by the above mentioned factors, including the Varroa mite and the pesticides, and the other factors are worsened by the pesticides. The pesticide industry finances research that focusses on the damage done by the mite, and of course they do find genuine effects. The focus on the mite is a red herring, meant to distract from the effects of neonicotinoids (Van der Sluijs et al. 2014).

The mite is a contributing cause, but not a sufficient cause. So are the lack of proper food and the presence of diseases. The pesticide however is sometimes a sufficient cause, never a necessary cause, but it is also a contributing cause to the other causes. The causal net for this case is rather complicated (fig 20).

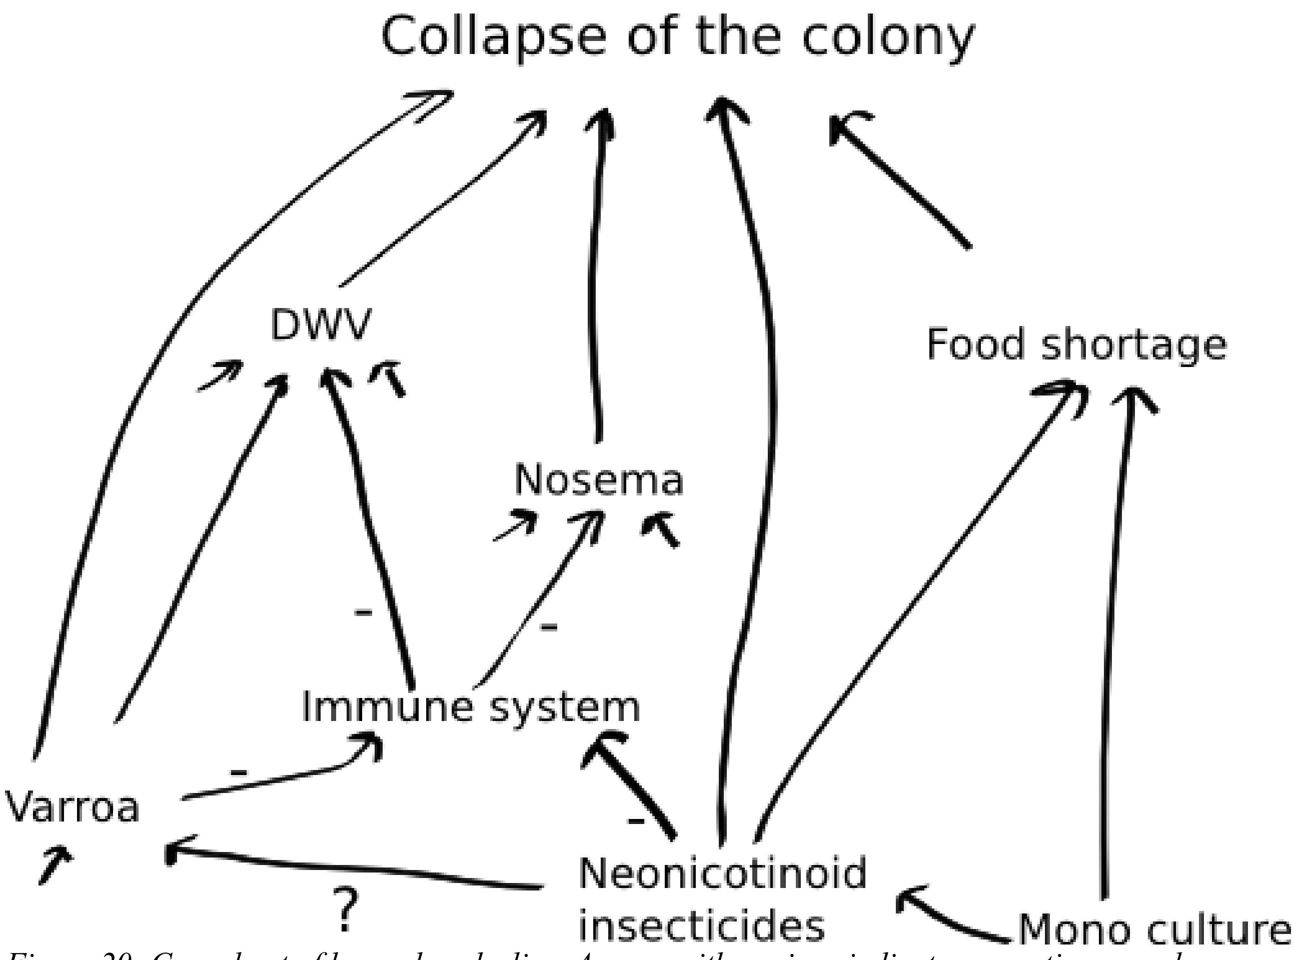

Figure 20: Causal net of honey bee decline. Arrows with a minus indicate a negative causal influence. Small arrows that come from nowhere indicate there are other contributing causes.

A causal net like this can make it hard for decision makers in politics to judge whether or not to allow or forbid the use of neonicotinoids, mainly because there are large economical consequences to these decisions, and a powerful lobby (Van der Sluijs et al. 2014).

Possible conflicts between disciplines

In the practical disciplines of science, the thoughts about causality have a different focus from the more theoretical and descriptive disciplines. How we need to deal with causes depends on what the goal is. If the goal is to explain, we can concentrate more on the different types of causes than on the size of the impact of a cause on an effect. But if the goal is to manipulate, causal strengths are important. We want maximum results of a manipulation. Only a complete picture of the relevant causes can give us a correct understanding of what we need to do to achieve a certain goal. Sometimes taking the cause away does not even solve the problem, for instance if there are replacement causes. It can also be the case that by taking away something that isn't a cause, like in the case of the catalyst, you can reduce the effect considerably.

In epidemiology, there is never a single cause. A sufficient cause is defined as a set of causes that are together sufficient to create the effect. In law, and mainly in law of tort, one has to find a single component sufficient cause to be able to tell if, had the cause not been there, the effect had been there. If this single cause is not found, it is hard to hold somebody accountable for any damage done to another person.

In the case of law suits by smokers who got lung cancer against the tobacco industry, this problem plays a large role. The industry can defend itself by saying that smoking is neither a sufficient nor a necessary cause of lung cancer. And that the person's genetic make up and history play a role. All of that is true. Yet we know that smoking plays a large role in the development of lung cancer.
All we can say about smoking is that it raises the chance of a smoker getting lung cancer, but how much it raises that chance also depends on the person.

The tobacco industry has been able to sow causal confusion for decades, by stating that smoking is nor a sufficient, nor a necessary cause, which is true, and by funding scientific research that would prove that other factors cause lung cancer, like asbestos or genetic make up. There are companies that specialise in creating scientific doubt for products that have a "marketing problem" and on topics like global warming, CFC's, the ozone hole and acid rain. The same companies that wrote scientific reports for the tobacco industry, and managed to get them peer reviewed, are now hired by the pesticide companies to do the same (Oreskes and Conway 2010). Causal confusion is an important tool for them. They use the slogan "Doubt is our product".

For politicians, legislators and other decision makers, it is hard to find out which cause out of a set of contributing causes contributes the most. We need a tool to determine causal strength. It is hard to make something that generates a value for causal strength, but it could be enough to generate a comparative causal strength. If we have that, we can decide what cause should be eliminated in order to solve a problem, based on the causal strength and other factors, like the elimination of which cause is the most feasible.

Causal decision tree: how do these types fit into one script?

The decision tree on page 48 up to and including page 50 is an attempt to combine the types of causality that we have distinguished into a single schematic. By answering simple questions about a case, the user can go through the tree and find out what type of causality he is dealing with. The questions are marked light blue. From the questions, arrows lead the way to a next question. The answers to the questions are printed next to the arrows. The user ends up with a description of the cause type he is dealing with, which is marked red for “not causally related” and green for “strongly causally related”, and anything in between.

The tree can give accurate insights under the condition that the examples that the user feeds it are of systems where the conceptual problems can be ignored. So there should be no problems with the direction of time, and the system should not be sub atomic.

The tree gradually becomes more complex. It starts out dealing with a system of only two events, then a causal chain, and eventually a causal net.

Decision tree

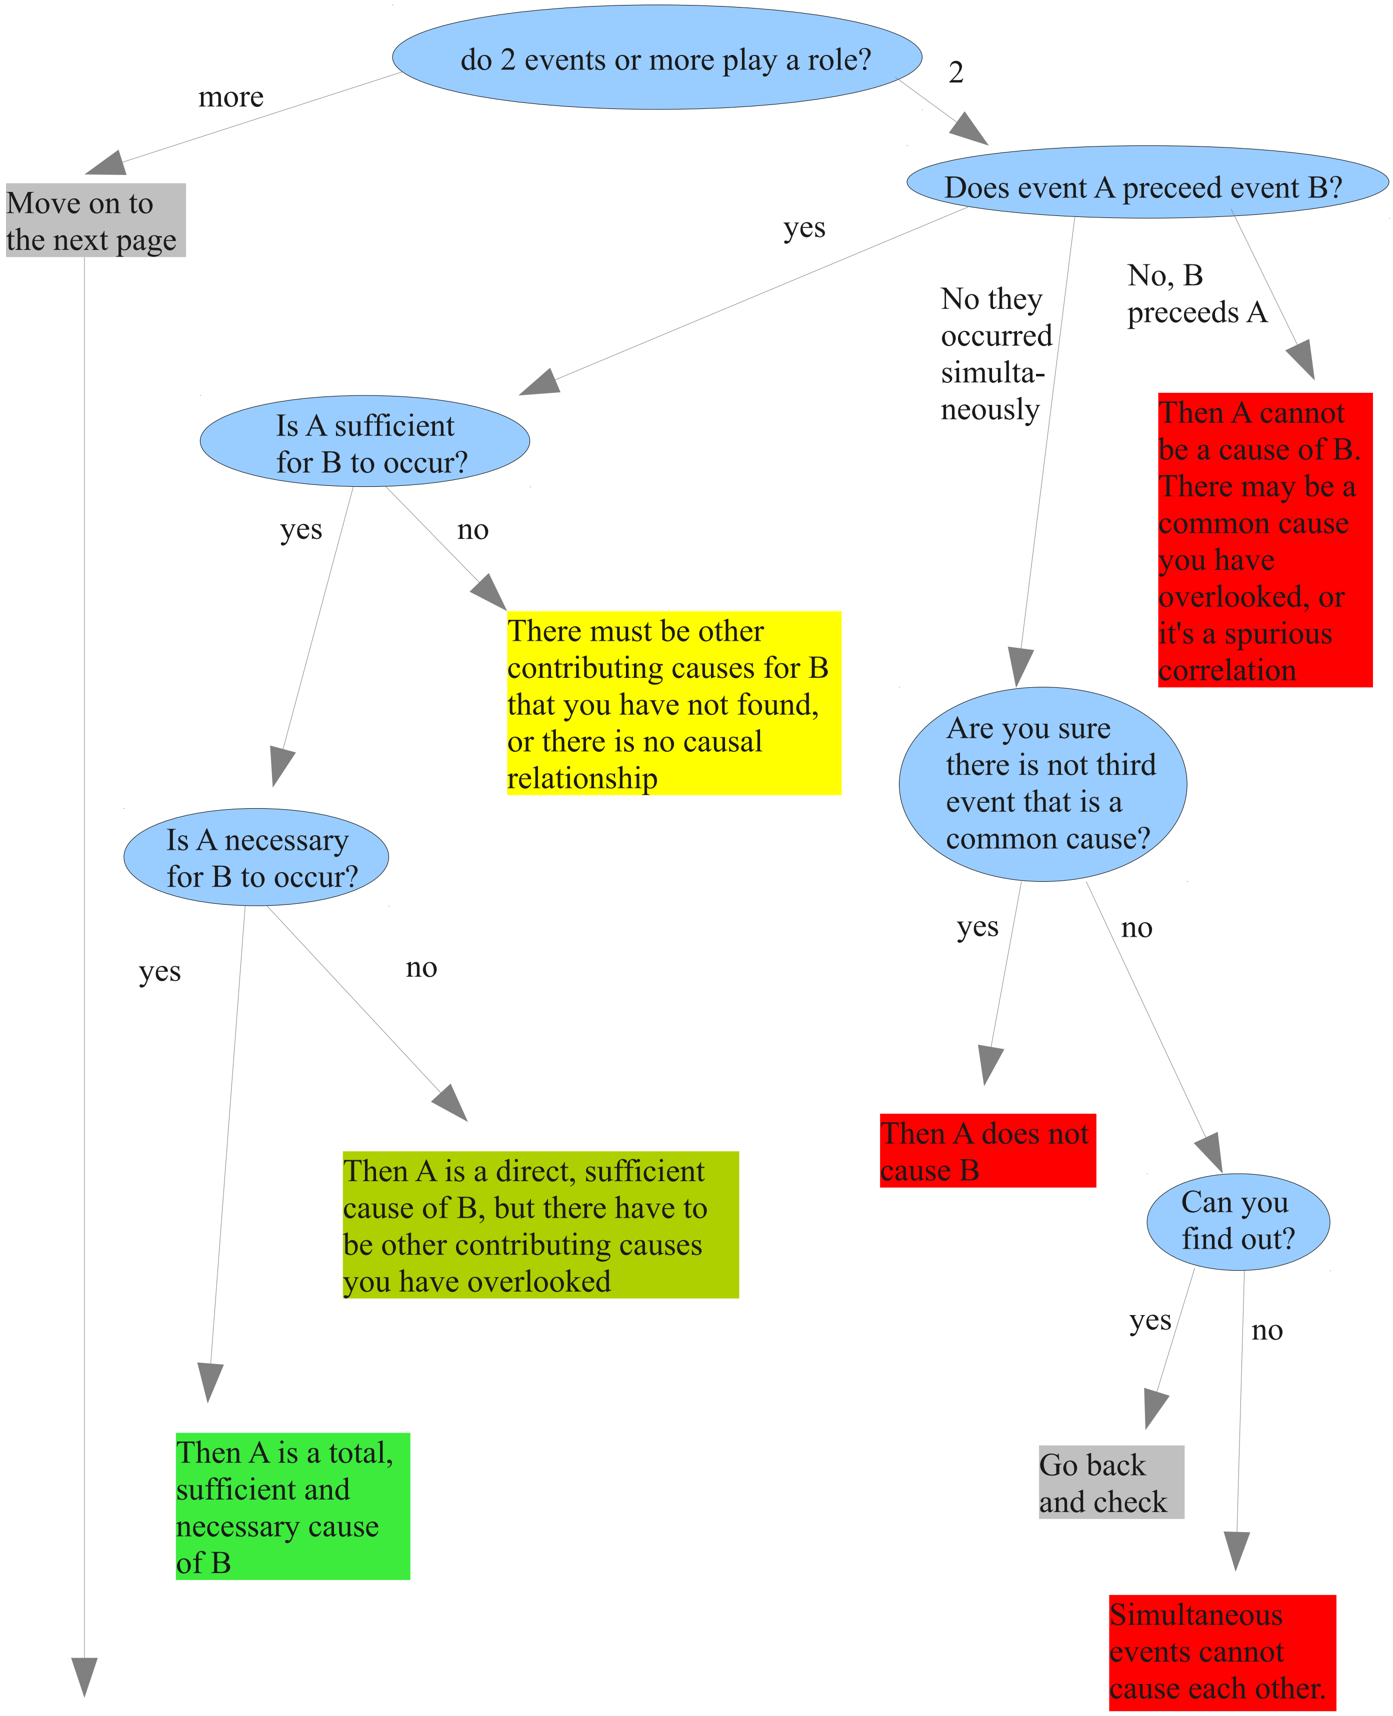

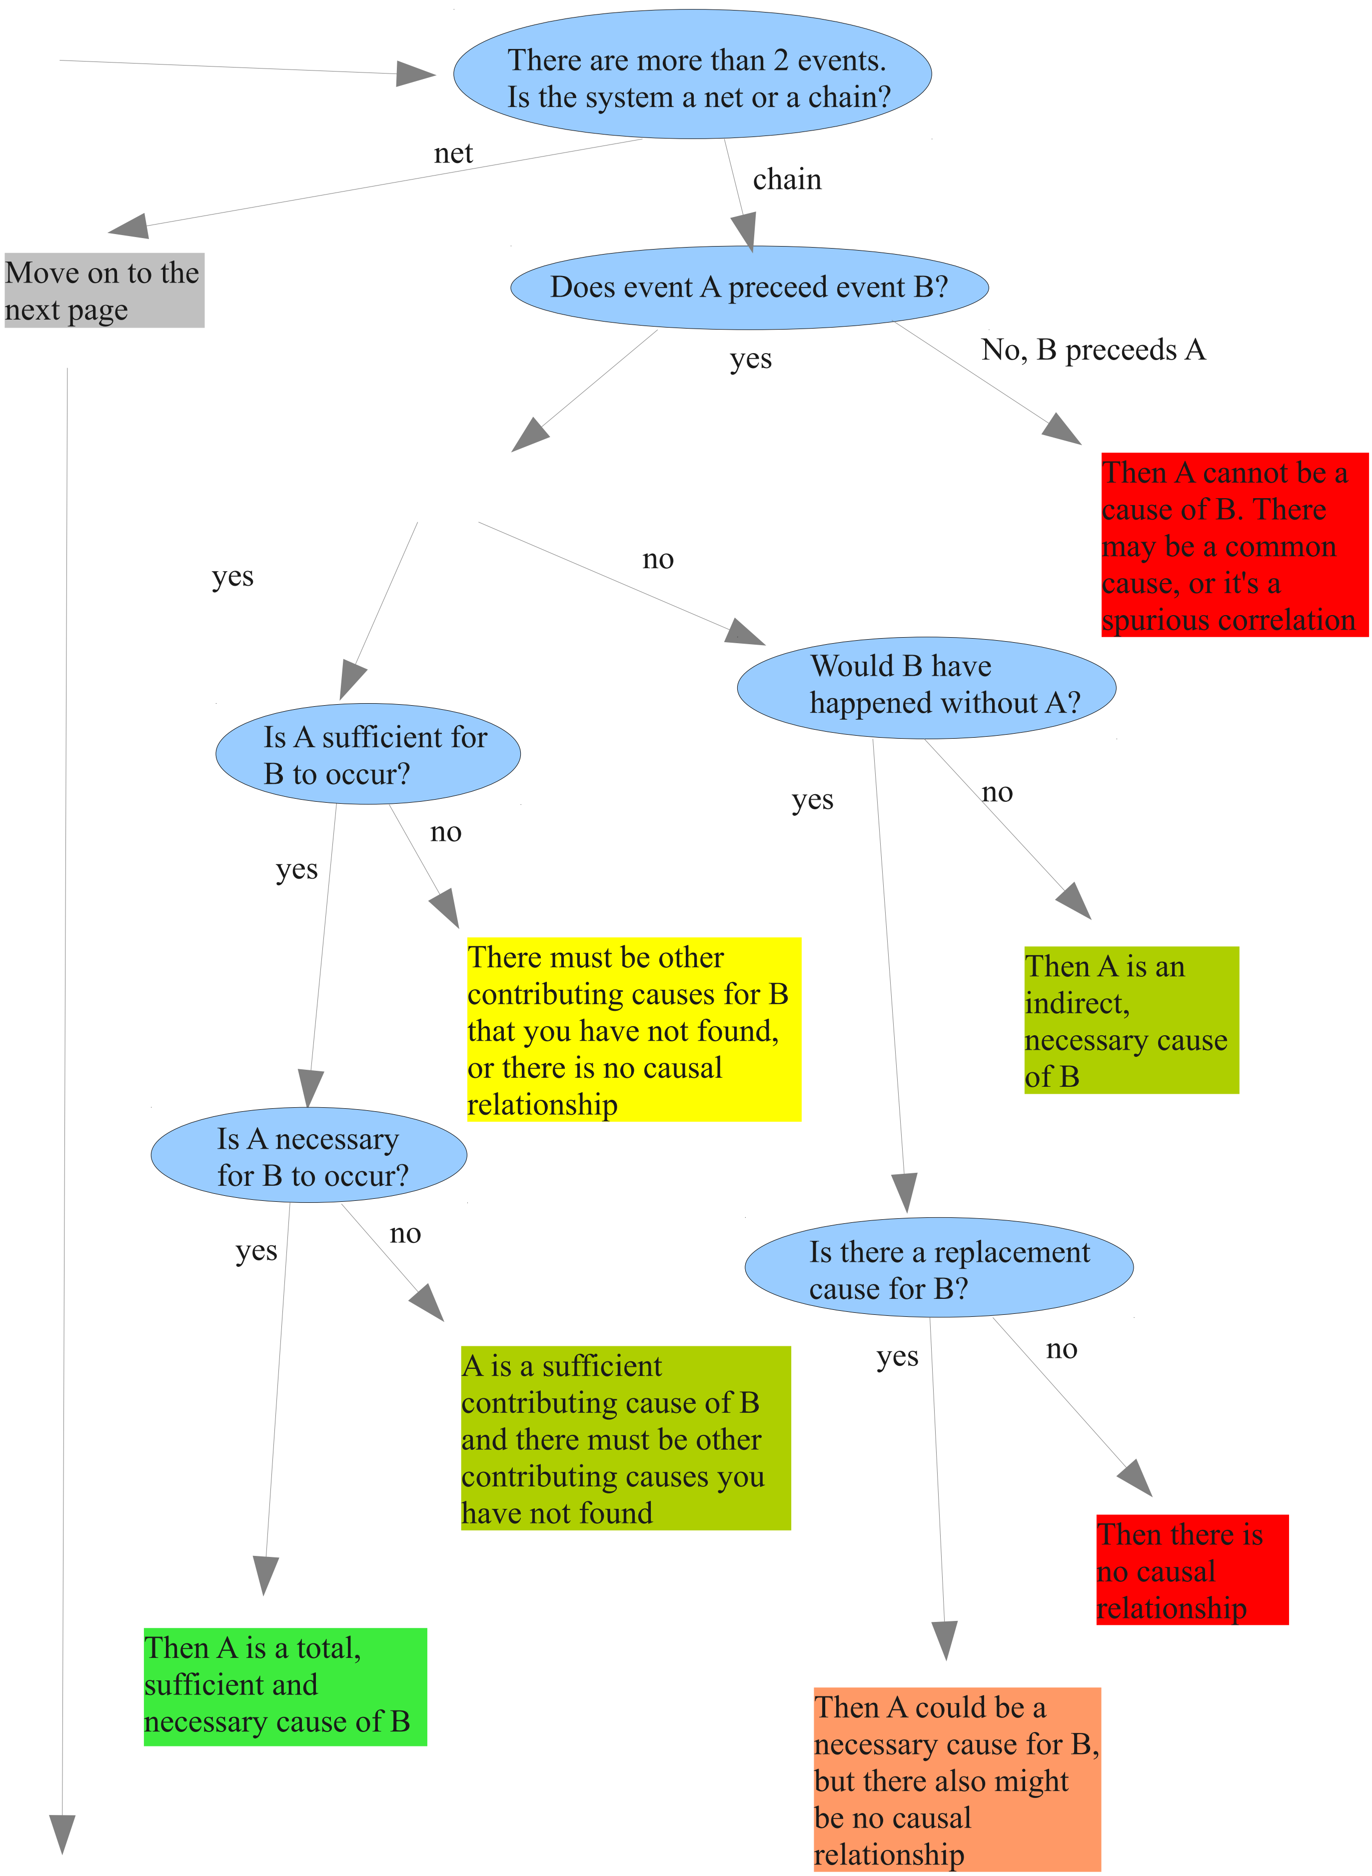
There are more than 2 events.
Is the system a net or a chain?
net
chain
Move on to the next page
Does event A preceed event B?
yes
No, B preceeds A
Then A cannot be a cause of B. There may be a common cause, or it's a spurious correlation
yes
no
Would B have happened without A?
Is A sufficient for B to occur?
yes
no
no
yes
There must be other contributing causes for B that you have not found, or there is no causal relationship
Then A is an indirect, necessary cause of B
Is A necessary for B to occur?
yes
no
Is there a replacement cause for B?
A is a sufficient contributing cause of B and there must be other contributing causes you have not found
yes
no
Then there is no causal relationship
Then A is a total, sufficient and necessary cause of B
Then A could be a necessary cause for B, but there also might be no causal relationship

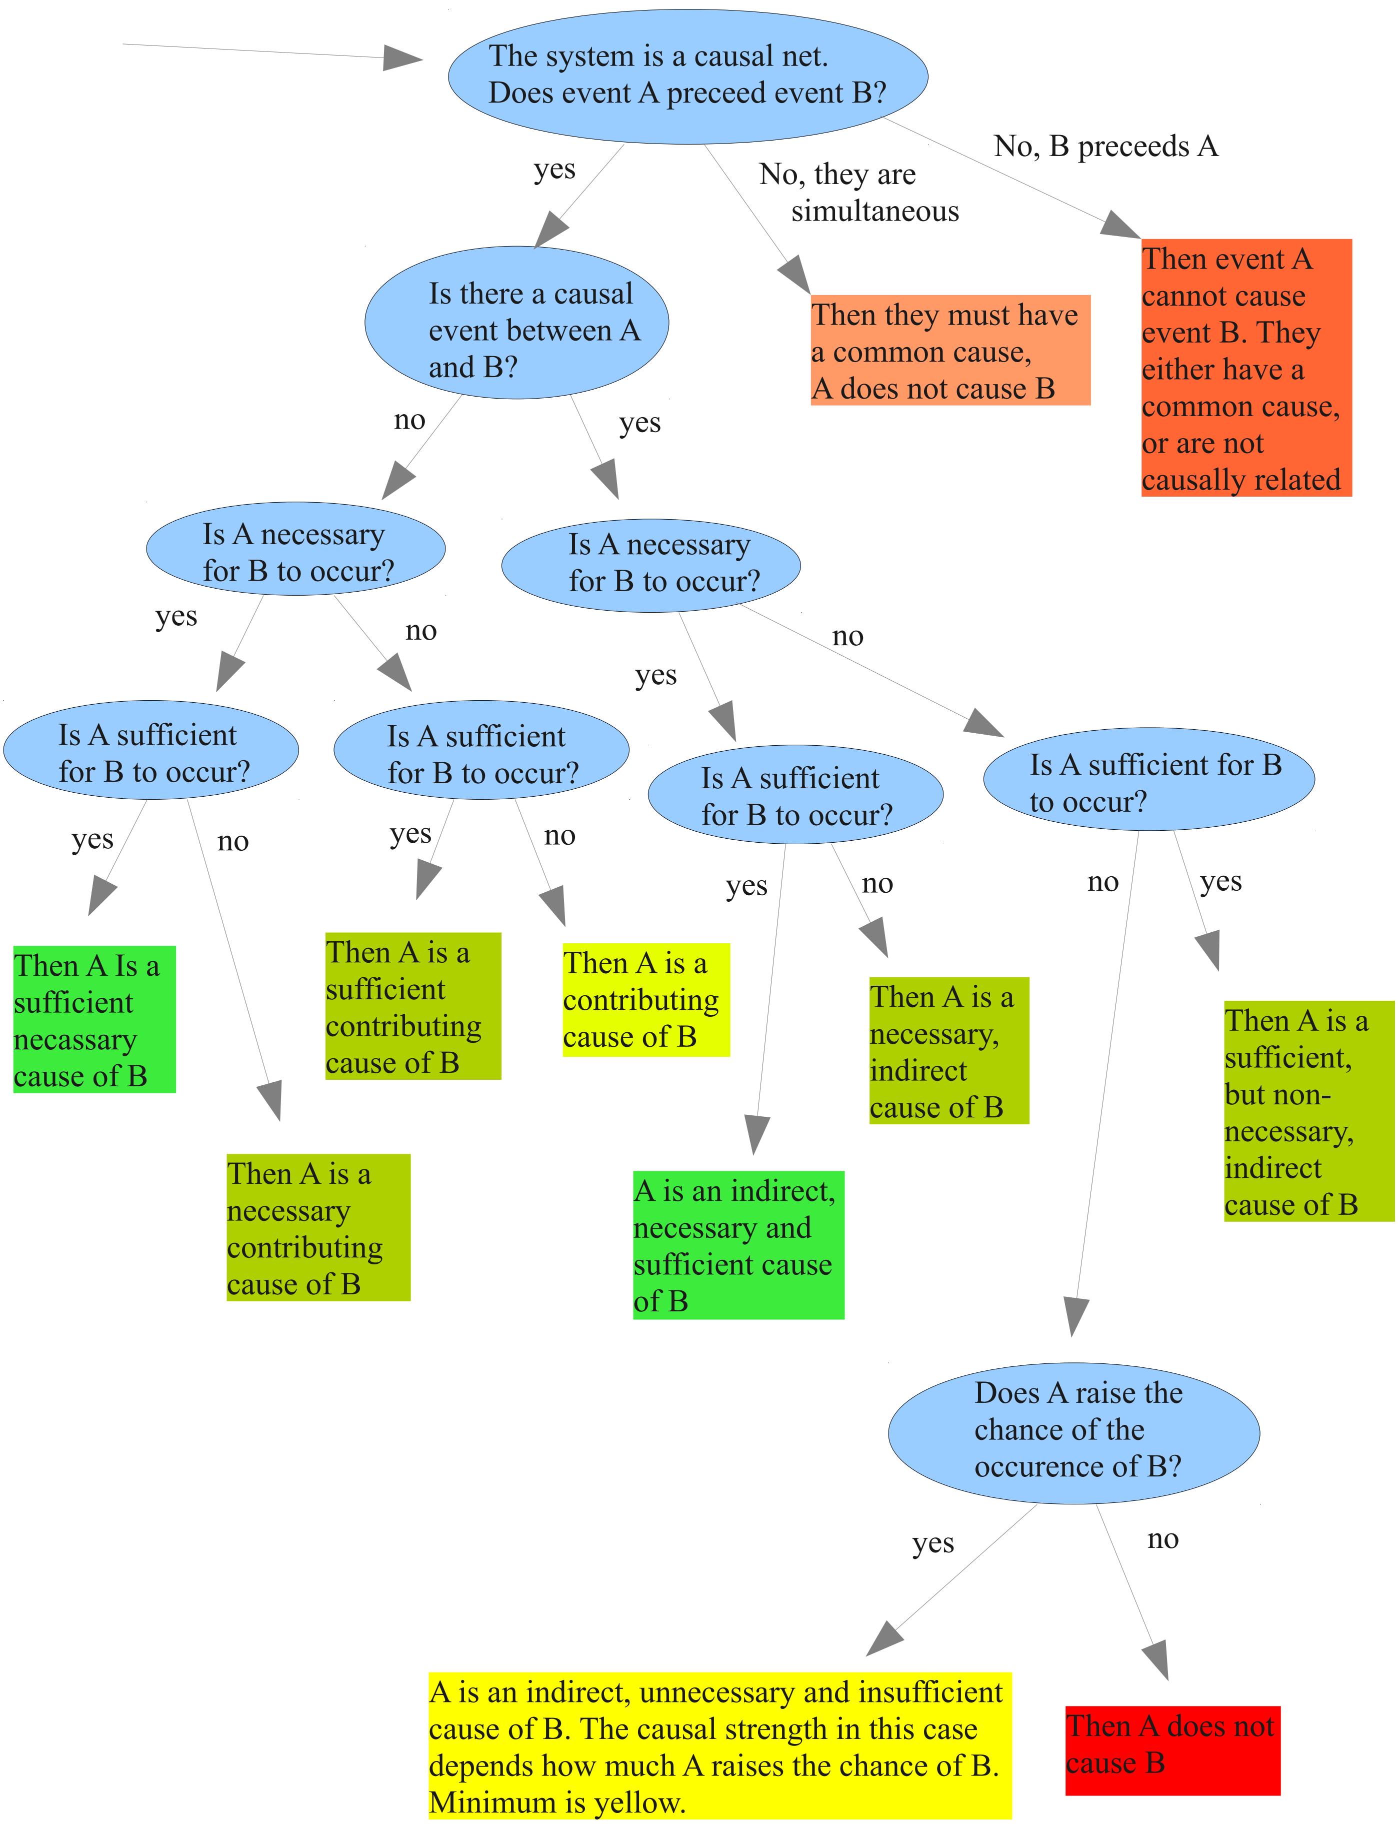
The system is a causal net.
Does event A preceed event B?
yes
No, they are simultaneous
No, B preceeds A
Then event A cannot cause event B. They either have a common cause, or are not causally related
Then they must have a common cause, A does not cause B
Is there a causal event between A and B?
no
yes
Is A necessary for B to occur?
Is A necessary for B to occur?
yes
no
yes
no
Is A sufficient for B to occur?
Is A sufficient for B to occur?
Is A sufficient for B to occur?
Is A sufficient for B to occur?
yes
no
yes
no
yes
no
no
yes
Then A Is a sufficient necassary cause of B
Then A is a sufficient contributing cause of B
Then A is a contributing cause of B
Then A is a necessary, indirect cause of B
Then A is a sufficient, but non-necessary, indirect cause of B
Then A is a necessary contributing cause of B
A is an indirect, necessary and sufficient cause of B
Does A raise the chance of the occurence of B?
yes
no
A is an indirect, unnecessary and insufficient cause of B. The causal strength in this case depends how much A raises the chance of B. Minimum is yellow.
Then A does not cause B

A system with only two events

In a system with only two events, most problems are easily dealt with or quickly run into a dead end. A cause is always a direct cause and a total cause, so it has to be a necessary and sufficient cause, unless there is a third event that the user has missed. In that case, it is not a two-event system but either a causal chain or a net and the user will have to move on to the next page and start over.

A causal chain

A causal chain is already a bit more complex, adding the possibility of indirect causes. The system still excludes multicausality, and therefore all causes are sufficient. They do not have to be necessary though, because of the possibility of replacement causes. In a system with only two events, there is no possibility for a replacement cause because there is no feedback loop. As we have seen in figure 14, the non-existence of effect B causes the replacement cause A' to produce a new instance of the effect, called B'. There are now at least three events in the system, four if you include A. One could argue that the replacement cause should also be an option in the two-events-system, because it only becomes a part of the system at a later time, when the replacement cause gets activated. Bin that case, the two-events-system becomes a causal chain, the moment the replacement cause gets activated. I have chosen to leave the replacement cause out of the two-events-system, but this choice is a bit arbitrary.

A causal net

In a causal net, the system becomes more complex. But some things are simpler: we don't have to ask whether there is a replacement cause anymore, because a non-necessary cause can exist without there being a replacement cause. After all, the net structure guarantees that there are multiple causes for each effect. In the causal chain that is different. The chain structure guarantees there is only one cause for each effect at any time, so if a cause is not a necessary one, there has to be a replacement cause.

What is this tree good for?

By asking the user the right questions in the right order, the tree helps determine what types of causal relations the user is dealing with. It leads to a causal type, usually a mix of types, that has a colour code that indicated the causal strength of the relationship. Red means no causal connection at all, green means there is a strong causal connection. The colours I have assigned to the types are still a bit arbitrary. I have given a necessary non-sufficient cause the same strength as a non-necessary sufficient cause, as a first attempt. But do those really have the same causal strength? The tree needs some critical feedback from possible future users.

The causal tree is very simple but so far it has survived many examples. I can imagine that it will at some point fail and will have to be adapted or expanded.

Let's run an example through the tree. For instance smoking causes lung cancer. We are dealing with more than 2 events, namely the smoking, lung cancer, the smokers genetic make up etc. These form a causal net, so we end up on page 50.
Does the smoking precede the lung cancer? Yes.
Is there another causal event between the smoking and the lung cancer? Yes, there is some biological process between them.

Is smoking necessary for lung cancer to occur? No.
Is smoking sufficient for lung cancer to occur? No.
Does smoking raise the probability of the occurrence of lung cancer? Yes.
Outcome: “Smoking is an indirect, unnecessary and insufficient cause of lung cancer. The causal strength in this case depends how much smoking raises the chance of getting lung cancer. “
In this case we know that smoking raises the probability of getting lung cancer rather a lot, so the colour of the causal strength will be close to fully green.

As it is now, the tree works well for cases in which the user knows the causes. A (rather complex) expansion could be to make it possible to say something about causality in cases where we don't know exactly what the causal relationships are, and even quantify the causal strength of those relationships. I did not go in that direction because the target user I have in mind for this causal tree is the advanced reader of scientific papers, who is not necessarily an expert on the subject. User can for example be science journalists, who have to communicate the content of a study to the public and to decision makers. To identify causes and quantify causal relationships, the user has to be an expert. Also, I feel that numbers and statistics should never replace rational judgment, and if you assign a value to a causal strength, that has the tendency to be taken as proof or at least to be something about which no discussion is needed.

Conclusions and discussion

Causality is an important concept for explanation, decision making and scientific proof. The vagueness of its definition, if there is one, within and between scientific disciplines causes confusion, both accidentally and intentionally.

There are some conceptual problems with the concept of causality, like the direction of time, the induction problem, the problem of counterfactuals ands indeterminism. Most of these problems originate in physics or philosophy and have little to do with practical cases from other disciplines. None of them have been solved properly, even after many decades of discussion. Most problems could be overcome or ignored by making assumptions that fit our observations. The uncertainty about the direction of time can be worked with by assuming that at this point in the development of the universe, within a single frame of reference, the direction of time is the direction of increasing entropy. The induction problem and the problems with counterfactuals we found a workaround for, using Lewis' many worlds comparison trick. The indeterminism of the world has its effects mainly on the sub atomic level, and as long as we apply our definition to larger objects, we do not expect to get into any problems with that issue.
This does not mean that we have solved the conceptual problems for the scientific disciplines they stem from. We have only found excuses for not taking them into consideration in other disciplines.

Using and combining the definitions found in the literature, it is possible to formulate a definition of a cause that should work for most scientific disciplines and that does not ignore conceptual problems that we might run into in practical applications. This definition is:
“a cause is any member of any actual set of plausible conditions that are jointly sufficient to produce the effect. That the cause leads or contributes to the effect has to follow logically from the relevant laws.”

In most of this thesis I have treated the cause as if it were an event, something that takes place at a certain moment. This is however not always the case. Maybe we should see it as a process of changing conditions, rather than a single event. Sometimes the process in slow and complex, sometimes it is simple and fast. The example of the thermohaline circulation shows us how a continuous process causes the circulation to shut down.
In the definition I have avoided the use of the word event.

Types of causes used in the literature are necessary causes, sufficient causes, total causes, contributing causes, spurious causes, replacement causes and direct and indirect causes. In some cases different disciplines of science used different definitions of cause types. In epidemiology, the sufficient cause was defined as a whole set of contributing causes, that together formed the sufficient cause. In other areas, but especially in law, the search is always for a single component sufficient cause.

Most of the confusion about causal relations, intentional or not, has been with the types of causes, and especially whether or not a causal relation is a strong one. This is not solved with a better definition for causality in general. I have built a decision tree to help the reader find out what kind of cause he is dealing with. The tree also gives a quantitive indication of the causal strength of the relationship.

The decision tree clarifies what the causal relations are in a certain example. It works for examples where we know what the causes are. This rather limits its usability, but there are many cases left, like the neonicotinoids case, or the smoking-causes-lung-cancer case, where it can be of help. It gives a rough estimation of causal strength, indicated with colours, from red for “no causal relationship” to green for “sufficient, necessary and total cause”. The tree could be expanded for use in cases where we don't know exactly what the causal relationships are, but should not become too complex for non-experts to use.

The case of neonicotinoid insecticides shows us that it is important to know the strength of a causal relationship, in order to take well informed decisions. Both the neonicotinoids case and the smoking-causes-lung-cancer case show that it can be lucrative to cast doubt on causal relationships, or to point at other contributing causes and prove that they are causes, in order to divert the attention from a cause that earns a company a lot of money. This can be avoided if the causal strength of a relationship is clear.

Vagueness in causal relationships has far reaching consequences. The concept of causality deserves to be taken seriously, not only in practical sciences, but also in philosophy.

References

Aristotle 1, History of Animals, Book V, Part 1

Aristotle 2, On the Generation of Animals, Book III, Part 11

Bresadolaa M (1998) Medicine and science in the life of Luigi Galvani (1737–1798). Brain Research Bulletin Volume 46 5:367–380

Campbell N, Reece JB (2005) Biology. Pearson Education Inc., San Francisco USA

Carnap R, Gardner M (1966) An Introduction to the Philosophy of Science. Basic Books Inc.

Cartwright N (1983) How the laws of physics lie. Clarendon Press, Oxford

Chisholm R (1946) The Contrary-to-Fact Conditional. Mind 55:289-307

Dowe P, Noordhof P (2005) Cause and Chance, Causation in an indeterministic world. Routledge USA

Van Fraassen B (1980) The Scientific Image. Clarendon Press, Oxford

Garcia J, Ervin F, Koeling R (1966) Learning with Prolonged Delay of Reinforcement, Psychonomic Science 5:121-122

Glynn L (2011) A Probabilistic Analysis of Causation. Brit. J. Phil. Sci. 62:343–392

Goodman N (1947) The problem of counterfactual conditionals. The Journal of Philosophy 44:113-128

Hausman D (1998) Causal asymmetries. Cambridge University Press, Cambridge UK

Hawking S (2001) The Universe in a Nutshell. Bantam Books, USA

Hesslow G (1976) Two Notes on the Probabilistic Approach to Causality. Philosophy of Science 43:290–92

Hill AB (1965) The environment and disease: Association or causation? Proceed Roy Soc Medicine 58:295-300

Höfler M (2005) The Bradford Hill considerations on causality: a counterfactual perspective. Emerging Themes in Epidemiology 2:11

Hoover KD (1988) The new classical macroeconomics. Blackwell, Oxford

Hume D (1739) A Treatise of Human Nature

Hume D (1748) An Inquiry concerning human understanding

Ioannidis JPA (2005) Why most published research findings are false. PLoS Med 2(8): e124

Kant E (1783) Prolegomena to Any Future Metaphysics http://eserver.org/philosophy/kant-prolegomena.txt. Accessed 1/7/2004 1:15:26 PM

Laplace PS (1814) Essai philosophique sur les probabilités, Paris

Lewis D (1973a) Causation. The Journal of Philosophy, Vol. 70, No. 17, Seventieth Annual Meeting of the American Philosophical Association Eastern Division (Oct. 11, 1973), pp. 556-567

Lewis D (1973b) Counterfactuals. Blackwell Publishing, Malden USA

Lewis D (1979) Counterfactual Dependence and Time's Arrow , Noûs, Vol. 13, No. 4, Special Issue on Counterfactuals and Laws (Nov., 1979), pp. 455- 476

Mackie JL (1974) The Cement of the Universe. A study of Causation. Clarendon Press, Oxford

NRC (2014) Superfood is Marketing, http://www.nrc.nl/handelsblad/van/2014/maart/15/supervoedsel-is-marketing-1356346 accessed 29-05-2014

Oreskes N and Conway EM (2010) Merchants of Doubt, How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming. Bloomsbury, New York

Reichenbach H (1921) Axiomization of the theory of relativity, 1969, Berkeley and Los Angeles, University of California press

Reichenbach H (1956) The Direction of Time. University of California Press, Berkeley, CA

Resnik DB (2003) Is the precautionary principle unscientific? Studies in History and Philosophy of Biological and Biomedical Sciences (34) 329–344

Rosen DA (1978) In Defense of a Probabilistic Theory of Causality. Philosophy of Science, 45:604–13

Rothman KJ, Greenland S (2005) Causation and Causal Inference in Epidemiology. Am J Public Health. 2005;95:S144–S150. doi:10.2105/AJPH.2004.059204

Salmon W (1984) Scientific Explanation and the Causal Structure of the World. Princeton University Press

Skinner BF (1947) Superstition in the Pigeon. Journal of Experimental Psychology #38

Van der Sluijs JP (2011) Bijensterfte: een nieuw risico. Magazine nationale veiligheid en crisisbeheersing mei/juni 2011

Van der Sluijs et al. (2013) Neonicotinoids, bee disorders and the sustainability of pollinator services. Current Opinion in Environmental Sustainability 5 doi:10.1016/j.cosust.2013.05.007

Van der Sluijs JP, Wildschut D, Maxim L (2014) Bijen en insecticiden: late lessen uit vroege waarschuwingen. Milieu - Dossier 20(3):37-40

Sober E (1983) Equilibrium explanation. Philosophical studies 43:201-10

Staddon JE Simmelhag VL (1971) The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. Psychological Review Volume 78 1:3-43

Ton O (2009) What happened to the Cargo Cults? Material Religions in Melanesia and the West. Social Analysis, Spring 2009, Vol. 53 1:82-102

UN, Press Release SG/SM/9993, https://www.un.org/News/Press/docs/2005/sgsm9993.doc.htm, accessed 16-06-2014

Woodward JF (2003) Making Things Happen, a Theory of Causal Explanation. Oxford University Press