Chapter 2

Causation

Wesley C. Salmon

You awaken to the sound of your alarm clock; you get out of bed to turn it off. Causal relations enter your life at the beginning of your day, even before you are fully conscious. The noise of the clock caused you to wake up and get out of bed; turning it off caused the sound to stop. The same sort of thing is repeated innumerable times throughout the day. Under normal conditions – you hope today is normal! – turning the key in the ignition causes your car to start. After you get going, pressure of your foot on the brake pedal causes the car to slow down. Causation is involved in virtually everything anyone does, day in and day out, year in and year out.

Down the street, a carpenter hits a nail with a hammer, causing it to penetrate some wood. An electrician flips a switch, causing a bunch of lights to go on. At the city morgue, a forensic pathologist tries to figure out what caused a particular person to die. Millions of years ago, the last of the dinosaurs died; scientists have long wondered why. In recent years, they have theorized that the collision of a massive body with our planet caused conditions in which they could not survive. In this case, no humans were involved. Astronomers make extremely precise measurements of the movements of a star; they conclude that its behavior is caused by a planet in orbit around it. Perhaps conditions conducive to the emergence of intelligent life exist out there beyond the confines of our solar system.

So what's the big problem? We all understand causation, don't we? Why should philosophers write long articles on the subject? For an answer, we have to go back to the eighteenth century, and consider what David Hume had to say about it. Whether he was right or wrong, his writings on this subject are the most important.[1]

Hume's Problem

Hume discussed many examples, but his favorite is the behavior of billiard balls. Suppose that we have one ball lying at rest on the table, and another moving rapidly toward it. They collide. It's logical, isn't it, for the one at rest to start moving as a result of the collision? Well, no, he says, it's easy to imagine that the one at rest remains in that state, while the other ball returns in the direction from which it came. But what does Hume's imagination have to do with it? If it were a matter of pure logic, from the description of the cause we should be able to deduce the nature of the effect.[2] Hume shows that this entailment doesn't hold; alternative outcomes are conceivable without any contradiction. Suppose, for example that the ball at rest is securely bolted to the tabletop. In this case, if the first ball hits the second precisely head on, we surely expect it to come back to where it started. If there's any question, we can set up these conditions and perform the experiment.

Wait, you say. Hume's example didn't include anything about the ball being bolted to the table. True. But the only way we changed Hume's example was by adding some further conditions. Having done this, we see that the new description is logically consistent, because that's what would actually happen. Surely, what's actual is possible, that is, logically consistent. Now, if we have a logically consistent group of statements, and then remove one of them, the result *cannot* be self-contradictory. The *only* way to make a consistent set of statements self-contradictory is to *add* a statement that conflicts with something already contained in that set. Having added a statement to Hume's description of the situation, *without assuming it was consistent*, we can now take it away again. The result is a demonstration that what Hume imagined as logically possible is logically possible.[3]

The point of this example can be generalized. From a description of a cause, it is impossible to deduce what its effect will be. Hume gives many other examples to support this point. In addition, he makes the same point in reverse; from a description of an effect, it is impossible to deduce what its cause was. Place a large diamond (if you happen to have one handy) in the freezing compartment of your refrigerator. After it has been brought down to 32°F, take it out and put it beside an ice cube. How, Hume asks, could anyone who has not had prior experience with such objects deduce that one of them is produced – caused – by extremely high temperature and high pressure, whereas the other would be completely destroyed by such conditions? He concludes that distinct events, including causes and effects, are logically independent of one another. No valid deductive inferences can be made from the existence or nature of one object from nothing more than a list of the properties of the other.[4] Consider an example in which an entailment relation does obtain. Someone might say, "High unemployment is caused by a lot of people being out of work." The word "cause" is certainly out of place; what we have is a definition of "high unemployment," not a cause–effect relation. Definitions govern the use of terms; they do not describe physical relations among events.

Having established his point that causal relations are not logical relations, Hume asks whether we can find any physical relation, such as the power of one event to produce another, or a necessary connection of a nonlogical sort between causes and effects. Let's return to the billiard ball example and examine it carefully. Now we are not looking for a logical relation between cause and effect; we are trying to observe a factual relation between cause and effect. We observe three aspects. First, we notice that the cause comes before the effect; the collision with the moving billiard ball comes before the motion of the ball that was at rest. Second, there is contiguity; the collision and the initiation of motion are close together in space and time. Third, we notice that the same sequence of events will occur every time we set up the same conditions. "Beyond these three circumstances of contiguity, priority, and constant conjunction I can discover nothing in this cause" (1955 [1739–40], p. 187).

What Hume failed to discover is far more important than what he found. A pair of events can satisfy the conditions of temporal priority and spatiotemporal contiguity by sheer coincidence. For instance, a loud thunderclap might have sounded immediately before your alarm went off. The pair of sounds exhibits priority and contiguity, but they bear no causal relations to one another. So we have to appeal to constant conjunction to determine whether these events are causally connected. Repeated experience reveals that thunderclaps are heard when no alarm rings immediately thereafter, and alarms ring when there is no immediately preceding thunderclap. Now, Hume argues, if we observe just one collision of two billiard balls, there is no feature of the situation that reveals the power of the collision to produce the subsequent motion. Moreover, if we have just one case of the thunderclap immediately preceding the ringing of the clock, there is no observable factor whose absence allows us to perceive the lack of causal connection.[5] It is only repetition that enables us to tell the difference. There is no objectively observable aspect of the events that discriminates between causes and coincidences. If we observe additional identical collisions of billiard balls, we will not notice some new characteristic that reveals the causal relation between the events.

After careful extended argumentation, which we have barely sketched, Hume concludes that the constant conjunction, which reveals nothing about the causal relations in the physical situation, has an influence on our minds. If we observe the same pattern of billiard-ball collisions several times, we come to expect the pattern to be repeated. When we see the collision, "habit" – Hume's term – leads us to expect motion to occur in the ball initially at rest. Notice, however, that this conclusion puts the connection between cause and effect in the human mind, not in the physical world. Our idea of causal efficacy is what we now call a conditioned response. It is exemplified by a famous experiment, performed in 1905, by Ivan Petrovich Pavlov. In this experiment, Pavlov rang a bell just as he was feeding his dogs. After this process was repeated a number of times, Pavlov rang the bell without providing any food. The dogs salivated when the bell rang, showing clearly that they expected food. This, according to Hume, is exactly what causation amounts to.

This conclusion is shocking. If causal connections exist only in our minds, then there were no causal connections before humans or other forms of intelligent life (remember Pavlov's dogs) existed, or in places that are not available for observation by such beings. I happen to believe that the Grand Canyon in northern Arizona is an effect of erosion by the Colorado River hundreds of millions of years ago. Not even dinosaurs were present at the time. So, according to Hume, there was no causal relation between the flowing water and the erosion of the earth at that place and time. You can see this more clearly if you imagine that an event – like the one that destroyed the dinosaurs 65 million years ago – had occurred much earlier, permanently eliminating all higher forms of life on earth. In that case, no organisms capable of forming habits would have been around to impute a *causal relation* between the motion of water over stone and the erosion of the stone.

If you're not uncomfortable at this point, you should be. There just has to be some way to escape this paradoxical view of causation. One idea, proposed by John Locke before Hume's time, is that we can directly perceive causal power in certain circumstances. Hume's prime example involves two material objects existing outside of our minds. We should, however, consider cases where conscious intentions are involved. In your philosophy class, you decided to raise your hand to ask a question. Your hand went up. Didn't you feel the power of your will to cause your hand to move? No, said Hume, in answer to Locke. Lots of things had to happen between your wish to raise your hand and the motion of your hand. Somehow brain processes had to send a signal through appropriate nerves to the arm muscles that must contract in order for your hand to move. You were not consciously aware of these intermediate processes, so you couldn't have directly perceived the power to move your hand. In order to learn what's involved in this example, you'd have to study neurophysiology, a highly technical field of contemporary science. However, even learning about all of these physiological details wouldn't really help. Even if you learn that a certain chemical process stimulates an electrical impulse to travel in a nerve, Hume's question comes back at you. What have you found in this situation other than a constant conjunction between a chemical reaction and an electrical response? It's hard to see how you could have perceived a power of the kind that Hume and Locke were talking about.

Classic Responses to Hume

According to Hume's successor, Immanuel Kant, these challenging notions sent him a real wake-up call.[6] By intricate and difficult reasoning, Kant concluded that our sense experience is necessarily organized in certain ways by our minds. His prime example is geometry; we must spatially arrange our perceptions in ways that conform to Euclidean geometry. He classified geometric knowledge as *synthetic a priori*. It is a priori because it can be established by pure reason alone, without the aid of sense experience. Euclid's axioms seemed self-evidently true, and the other propositions follow from the axioms with logical necessity.[7] Therefore,

geometry is a priori. But geometry is synthetic because it contains useful information about the physical world; even the ancient Egyptians used it for such practical ends as architecture and surveying. Thereupon, Kant extended his idea of synthetic a priori principles to include causality as well as geometry: "Everything that happens (begins to be) presupposes something from which it follows according to some rule."[8] Even if we were to accept this principle – which is extremely dubious in the light of modern science – it would not be very helpful. Suppose that lightning strikes a tree and we seek its cause. According to Kant it follows from *something*, but there's no clue about what that thing is or where to look for it. Moreover this unknown something is connected to the lightning strike by *some* rule, but again there's no hint as to what kind of rule it is. It's as if someone tells us that there is oil below the seas, without telling us how to find it or how to extract it.

In the middle of the nineteenth century, John Stuart Mill gave us five rather useful rules for discovering or proving cause–effect relations; in fact, they are still in use. First, the *method of agreement* applies to situations in which an effect occurs in many different circumstances, which have only one feature in common. For example, if all of the students in a particular dormitory, who have eaten the same dorm dinner, suffer severe gastrointestinal distress on the same night – and that is the only food that all of them consumed that day – there's a good chance that contaminated food in the dinner caused their digestive systems to behave so disagreeably.

Second, the *method of difference* is used in situations where all antecedent factors except one are the same – that is, some factor is present in some of these cases but absent in others. Moreover, an effect follows in cases in which the factor in question is present but fails to occur when it is absent. In these circumstances, there's a good chance that the factor that differs among the antecedent conditions is the cause of the one that differs among the subsequent circumstances. In one of my first forays into a shop to buy some music on CDs, for example, I made my selections and then headed for the men's room prior to checking out. *Stupid!* As I entered a short hallway, a loud alarm sounded; it nearly scared me out of my wits. A rapid analysis of the situation enabled me to conclude that taking CDs into rest rooms is not allowed and that I had set off the alarm. A quick survey of the situation showed that my crossing the boundary with the CDs was the only difference relevant to the sounding of the alarm. I'm happy to say that the staff readily concluded that it was only my ignorance that led me to carry them in; it was not a case of attempted shoplifting.

Third, the *joint method of agreement and difference* – as the name indicates – is a combination of the foregoing two, and it is more powerful than either by itself. To begin, you get many different cases where the result is present, and many others in which the result is absent. We can see how this works by an actual example similar to the dorm illness case. On an overseas airline flight on a jumbo-jet, two different dinners were offered – one was meat, the other seafood. Everyone who chose the seafood dinner became sick; no one who chose the meat was adversely affected. There can hardly be any doubt that some ingredient in the seafood dinner was tainted.

Fourth, the ***method of concomitant variation*** applies to situations in which the factors involved are neither completely absent nor completely present to a fixed degree. Instead, some causal factor varies in degree along with a variation in some effect. For example, you might find that the amount of natural gas you use in your household each month varies with the average daily temperature. Clearly there is a cause–effect relation between the outdoor temperature and the amount of gas you use.[9]

Although Mill didn't mention it, perhaps the most important kind of concomitant variation is a variation in frequencies of certain types of occurrences. Tests have shown, for example, that a smaller percentage of men who take an aspirin tablet every day suffer heart attacks than of those who don't. Taking aspirin reduces the risk of a heart attack. It isn't a case of all or nothing. Men who take aspirin do get heart attacks; men who don't take aspirin escape such attacks. It's a matter of percentages. Testing causal hypotheses through the use of such controlled experiments is a powerful tool of modern science that is widely used today.

Fifth is the ***method of residues***. If we have a list of possible causes of some phenomenon, and if we have in some fashion ruled out all but one, then the remaining possible cause is likely to be the genuine cause. For example, the planets in the solar system travel around the sun in nearly elliptical paths, but they deviate somewhat from perfect ellipses because of gravitational attraction of other planets. In the nineteenth century, it was noticed that the orbit of Uranus differed from the path calculated by taking account of the influence of other planets, especially Jupiter and Saturn. Since the other known planets could not account for all of the deviations, a previously unknown planet, Neptune, was postulated as the cause; it was observed telescopically not long thereafter.

There is no question that, even though Mill's methods have their limitations, they are useful for distinguishing genuine causal relations from relations that only appear to be causal. Hume didn't help us much with this aspect of causation, but he was interested in a different kind of problem. Hume was asking what causality is. What precisely is the power of a cause to bring about an effect? What kind of relation binds effects to their causes? These are profound questions, and philosophers are still trying to find answers to them. Mill's methods are not particularly helpful in answering the kinds of questions Hume raised. They don't go much beyond the notion that causation amounts to constant conjunction, as Hume had said, but they are useful tools for determining which conjunctions are constant and which are not.

Mill's great work, *A System of Logic*, is *the* classic nineteenth-century work on causation. The twentieth-century classic is J. L. Mackie's *The Cement of the Universe*. Its title is a phrase used by Hume; it poses the question of what holds the universe together. It attempts to show the nature of this cosmic glue that we call "causation." Mackie offers penetrating analyses of the works of Hume, Kant, and

Mill, as well as many others. To understand his work, we must introduce some standard terminology.

Sentences having the form, "If *A* then *B*," are called ***conditional statements***. They say that if the condition *A* is present, then the consequence *B* will also hold. In this case, *A* is called a ***sufficient condition*** of *B*. For example, in the title of an old song – "If you've got the money, honey, I've got the time" – your possession of money is ***sufficient*** for the availability of my time. Another kind of conditional statement has the form, "If not-*A* then not-*B*." In this case, it says that if the condition *A* is *not* present, *B* will *not* hold. In this case, *A* is a ***necessary condition*** of *B* – *B* will not occur unless *A* does. The concluding line of the aforementioned song is, "If you run out of money, honey, I'll run out of time." This says that your possession of money is ***necessary*** for the availability of my time. If both of these statements are true, we say that *A* is a ***necessary and sufficient condition*** of *B*. According to the song, then, your possession of money is ***necessary and sufficient*** for the availability of my time.

Conditional statements of the foregoing sorts are often construed as generalizations, i.e., the symbols "*A*" and "*B*" are taken to refer to classes or types of events. Given this interpretation, to say that *A* is a sufficient condition of *B* means that in every case in which an event of *type A* occurs an event of *type B* will also occur. Similarly, to say that *A* is a necessary condition of *B* means that whenever an event of *type A* fails to occur, no event of *type B* will occur. For example, having one's head cut off is a sufficient condition of dying; it is not a necessary condition because there are many other causes of death besides decapitation. Having oxygen to breath is a necessary condition for staying alive. In the absence of oxygen, a person dies. But having oxygen to breathe is not a sufficient condition of life; without food and water as well a person cannot continue to live.

A philosophical account of causation according to which causal relations can be analyzed entirely in terms of sufficient conditions, necessary conditions, or any combination of such conditions is a ***regularity theory***. Those who hold regularity theories understand such conditions as generalizations, and they usually require that if an instance of *A* is a cause of an instance of *B*, it precedes that instance of *B*. Hume's theory, as explained above, makes his view a regularity theory. In addition to temporal priority and constant conjunction, he also requires spatiotemporal contiguity. However, it isn't clear from his various statements whether he regards causal regularities as sufficient conditions, necessary conditions, or both. That's not especially important because any simple regularity theory faces serious difficulties.[10] Many regularities that fit the description cannot be considered *causal* regularities; for example, day regularly precedes night, but day doesn't cause night. The regular succession of day and night is caused by the rotation of Earth.

Although Mackie does not, himself, advocate a regularity theory, a major part of this theory would qualify as an extremely sophisticated regularity account.[11] It can best be explained by a concrete example. Suppose that a barn burns down. Speaking loosely, we would say that there are many possible causes, e.g., a lighted cigarette dropped by a careless smoker, a stroke of lightning, deliberate arson, a spark from a workman's torch, and many others. Since there are many possible causes, none of the foregoing can qualify as a necessary condition. Suppose, for the moment, that the careless disposal of a cigarette is (a part of) the actual cause. It cannot, by itself, be a sufficient condition, because other factors must be present. It must land on some inflammable material, such as dry straw, rather than on a clean concrete floor. This straw must be located near other inflammable material, such as wood, in order for the fire to spread. No one who would have put out the fire before it spread could have been present. And so on. A fairly complex *set of conditions* must be fulfilled in order to make up a genuine sufficient condition. At the same time, the dropping of the cigarette is an indispensable part of this set of conditions.

Mackie (1974, p. 62) adopts the acronym INUS – standing for an Insufficient but Nonredundant part of a condition that is Unnecessary but Sufficient – to designate such causal factors as the careless tossing of the cigarette. The complex of conditions containing the improper disposal of a burning cigarette is a sufficient, but not necessary, condition of the burning of the barn; the dropping of the lighted cigarette does not, by itself, constitute a sufficient condition for the burning of the barn, but it is a necessary part of that complex of conditions.

A natural question arises at this point. Why do we pick out the dropping of the cigarette as the cause, rather than another indispensable part of the entire sufficient condition – e.g., the presence of dry straw? According to Mackie, this decision often depends on what is considered usual or unusual in the circumstances. In our particular case, it may be that dry straw is usually present; what is unusual is its contact with a burning cigarette. In some barns, however, the presence of dry straw might be an unusual circumstance (depending on what the barn is normally used for). If the barn had been of the latter kind, the presence of dry straw might be cited as the cause. It, too, is an INUS condition. As Mackie clearly notes, the selection of one INUS condition rather than another as *the cause* depends strongly on the context. Human interests, purposes, and knowledge play a large part in the selection of the cause.

As we saw, Hume's analysis locates causation "in the mind (imagination)." Mackie's aim is to find causation "in the objects" – i.e., in the external world as it exists independently of the human mind. Contexts are determined by the human point of view; conceptions of causation that depend on context don't succeed in finding causation "in the objects." Realizing this point, as Mill had done before, Mackie tries to analyze "the *full* cause" in terms of the set of all of the possible sufficient causes, each of which is spelled out in detail as a conjunction of factors that have to be present in order to insure the sufficiency of each term. In our discussion of the burning of the barn, we obviously left out many possible causes, e.g., spontaneous combustion in hay stored in the barn, burning debris from a nearby forest fire falling on the roof, being struck by a meteor, etc. Mackie (1974, p. 76) realizes that we don't usually know all of the possible causes, so the *full cause* will be represented by an "*elliptical* or *gappy* universal" statement. However, this doesn't solve the problem. Such universals aren't really statements at all; they

are *forms of statements* containing blanks that have to be filled in. The blanks reflect our ignorance of the exhaustive list of sufficient conditions – the list that represents the full cause.

Regularity theories in general face a further problem, namely, *causal preemption*. It can be illustrated by the barn example. Suppose that our careless smoker tosses his lighted cigarette toward the dry straw (and that all the other attendant conditions are fulfilled). Suppose, however, that lightning strikes the barn, which is unprotected by a lightning rod, just before the cigarette reaches the straw. The tossing of the cigarette is still an INUS condition by Mackie's formal definition, but the lightning is the cause of the fire. The literature on causation is full of examples of causal preemption; I don't think any version of the regularity theory has a satisfactory resolution of this difficulty.

One of Hume's definitions of causation suggests a different analysis, namely, one in terms of *counterfactual conditionals*. A counterfactual conditional statement is an "if . . . then . . ." statement whose antecedent clause is false. An old nursery rhyme (which I'm extending a bit) contains a sequence of counterfactual conditionals: "For want of a nail, the shoe was lost; for want of a shoe, the horse was lost; for want of a horse, the rider was lost; for want of a rider, the message was lost; for want of the message, the battle was lost; for want of the battle, the war was lost; for want of a victory, a kingdom was lost. All for want of a nail." If the nail had not been missing, the horseshoe would not have been lost – the missing nail *caused* the horseshoe to be lost. If the horseshoe had not been lost, the horse would not have been lost – the missing horseshoe *caused* the unavailability of the horse. And so on.[12]

The problem with counterfactual conditionals is that they are highly context-dependent. In the preceding case, the series of counterfactuals conjures up an overall picture of a military situation with many implicit assumptions about the surrounding conditions. It assumes that there was no means of communication except by the delivery of a message by a rider; it could not have been carried by a runner or sent by a carrier-pigeon. It assumes that no other horse and rider were available to carry the message. It vaguely assumes conditions of battle in which the receipt of a message means the difference between victory and defeat. Because of the dependence on contextual factors, it is extremely difficult to specify objective conditions that determine the truth or falsity of counterfactuals. Therefore, any attempt to analyze causality in terms of counterfactual statements will fall short of locating causality "in the objects." If there is any question about this, consider the difficulty of specifying precisely the situation in which the lack of a horseshoe nail could *actually cause* the loss of a kingdom. Philosophers have made many attempts to solve the problems raised by counterfactuals, but none has been clearly successful.

The approaches to causation that we have considered so far all share a fundamental feature, namely, they analyze the relation between causes and effects in terms of *sentential connectives* – that is, terms like "if . . . then . . ." or "unless" whose function is to join complete sentences together to form more complicated sentences. Regularity views, such as those advocated by Hume, Mill, and Mackie,[13] rely on statements of the form, "If anything is an *A*, then it is a *B*." Counterfactual theories require statements of the form, "If *X* had not occurred, then *Y* would not have occurred," or, equivalently, "*Y* would not have occurred unless *X* had occurred." Sentential connectives are *logical* terms. Perhaps we should be looking for a different kind of relation.[14] Hume had looked, without success, for "necessary connections" and "secret powers." We should look again. As we shall see, the connections need not be necessary and the powers need not be secret. I believe they can be discovered in the physical world and that *they can be objectively characterized entirely in terms of noncausal concepts.* This is the task to which I now turn. Mackie – in his effort to find causation "in the objects" – tried without success to accomplish this goal. He did not succeed. Your philosophy instructor probably believes that it can't be done.

Causation in the Objects – Causal Processes

You see a friend across the street, but she is looking in a different direction, so she doesn't see you. You shout "Hi, Mary!" The sound waves you created travel toward her; she turns and waves to you. Light reflected from her hand reaches your eyes. You know that she has heard your greeting and has recognized you. Stripping this encounter down to its bare essentials, we have three events, namely, your shout, Mary's turning and waving, and your seeing her wave. They are connected by two processes, namely, the sound waves that travel from your mouth to her ears and the light rays that travel from her hand to your eyes. The three events are *causally connected* by these processes. In the end, I shall claim that processes of various types are precisely the causal connections that Hume sought and failed to find, but several key concepts need to be clarified in order to explain this idea.

To begin, we must understand, at least in a fairly rough way, what is meant by the term "process." According to the dictionary, the core meaning of this term designates something that goes on continuously over a span of time. Frequently, it transpires over a spatial distance as well. Some examples will help. As we have already noted, sound waves and light rays are processes. Material objects, such as Hume's famous billiard balls, are processes. An airplane flying through the air is a process, and so is the shadow it casts on the ground on a sunny day. A paperweight that is motionless on your desk is a process because it endures through a span of time – well, it's motionless with respect to your room, but it moves with the Earth as the Earth rotates on its own axis and in its orbit around the Sun. A football being passed by one player to another is a process; so is the motion of the image of the same football if you are watching a movie of the game.

As this term is commonly used, a process is something that is related to some particular goal, whether or not it is designed for this purpose. For example, the process of annealing is designed by humans to make steel softer and less brittle by heating and cooling it. The process of sedimentation produces fertile plains by a river's depositing of silt in a valley. This process was not designed by humans. While we recognize that many processes are picked out because of notable or desired results, we shall not limit our concept to them.[15]

Our notion of a process is similar to Bertrand Russell's concept of a *causal line*. "A causal line may always be regarded as the persistence of something – a person, a table, a photon, or what not. Throughout a given causal line, there may be constancy of quality, constancy of structure or gradual change in either, but not sudden change of any considerable magnitude" (1948, p. 459; pt. VI, ch. 5 is entitled "Causal Lines"). Our concept of process will become much clearer as we proceed. For now, we need to distinguish two kinds of processes, namely, *causal processes* and *pseudo-processes.* Unfortunately, Russell did not make this distinction. We will aim to show how causal processes transmit *causal influence*, whereas pseudo-processes do no such thing. Suppose that you are driving along a road on a sunny day. Both your car and its shadow are processes, but the car is a causal process and the shadow is a pseudo-process. One way to make the distinction is by noting what happens to these two processes if they encounter obstacles. If the shadow meets a stone pillar standing at the roadside, it is temporarily distorted, but as it passes beyond the pillar it resumes its former shape as if nothing had happened. If, however, your car collides with a stone pillar, it will bear the marks of the encounter long after it has moved past the pillar (assuming that it is capable of going on). This means that the causal process (the car) can transmit the marks of the collision beyond the place where it occurred, whereas the pseudo-process cannot transmit any such mark.

Because the distinction between causal processes and pseudo-processes is so fundamental to this approach to causation, let's look at another example. In my office hangs a picture of the famous lighthouse that stands at the port of Genoa, Italy. Like the many other lighthouses all over the world, it is lit at night and its beacon rotates, sending beams of white light in all directions. As I stood at the window of my hotel room one evening during a visit to Genoa, I watched as it sent a beam periodically in my direction. As the light turned, it cast a moving spot of light across the wall of my hotel. As it pointed in other directions the spot of light moved across clouds in the distance. The light sent out in any given direction by the beacon was a causal process; if a piece of red glass had been placed in the beam anywhere between the lighthouse and my hotel window, I would have seen a flash of red light. The color of the white beam would have been changed in its encounter with the red filter, and that change (a mark) would have been transmitted by the pulse of light traveling from the beacon to my window. In contrast, if I had covered my window with transparent red plastic, I would have marked the spot of light traversing my window, but the mark would not have persisted as the spot of light moved on. The mark could have been imposed at one place, but it would not have been transmitted by the moving spot. The moving spot was a pseudo-process; it could not transmit that mark or any other mark.

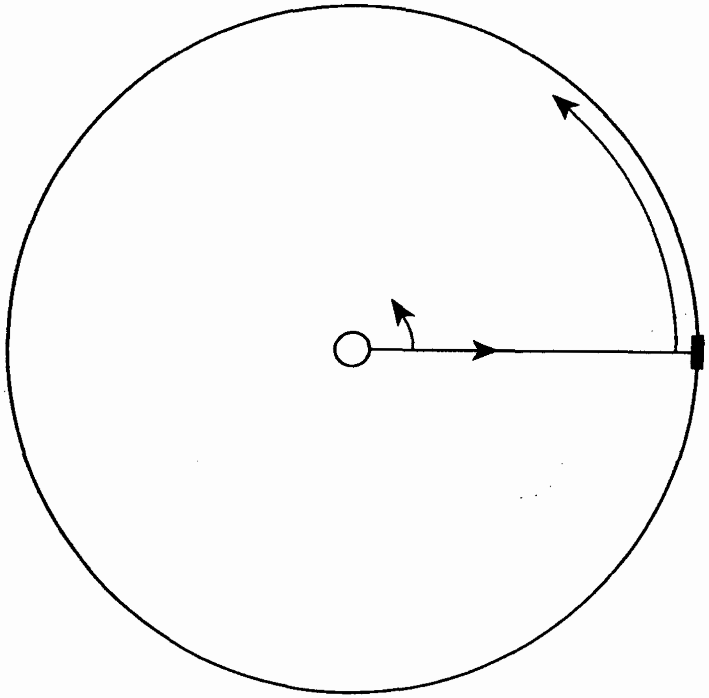

Figure 2.1 Rotating beacon

Causal vs. pseudo-processes

The processes we observe on a movie screen or a TV tube are pseudo-processes. If a passionate (quite possibly drunk) fan of one particular football team – desperately hoping to prevent completion of a pass – were to draw a pistol and fire it into the image of the football on the cinema screen, a hole would be made in the ball at one particular point, but the hole *in the ball* would not persist as the image moved across the screen. Again, the mark could be made, but it could not be transmitted. Shooting the image of the football on the screen would have no effect whatever as far as the completion of the pass was concerned. A similar shot at the image of the football on a TV screen would simply terminate that pseudo-process and all of the others that would have constituted the rest of the TV game as well.

In Hume's lifetime, no fundamental reason existed for the distinction between causal processes and pseudo-processes. The situation changed radically early in the twentieth century – 1905 to be precise – when Albert Einstein formulated his special theory of relativity. People often say that, according to that theory, nothing can travel faster than light.[16] This is simply wrong, unless the term "thing" is construed very carefully. The correct statement is that no causal process can transpire faster than light. Pseudo-processes can go at arbitrarily high velocities. Recall the

Genoa lighthouse. Imagine that it is surrounded by a circular wall around which the spot of light moves. The spot must traverse the entire circle each time the beacon makes a single complete rotation. The larger the circle, the faster the spot must travel to complete its course in the required time. Here is a much more dramatic example. In the Crab nebula, 6,500 light years from Earth, a neutron star (a pulsar) rotates thirty times per second. It sends out a beam of light just as the lighthouse does. The "spot" of light that zips past us traverses the "circle" whose radius equals our distance from the pulsar in one-thirtieth of a *second*; the time it takes for *light* to cross the diameter of this circle is 13,000 *years*. As the spot passes us, it is traveling at approximately $4 \times 10^{13} \times c$ (the speed of light).

Causal processes – in contrast to pseudo-processes – are the means by which causal influence is transmitted from one place and time to another. If a process can transmit a mark, it can transmit information, just as a radio signal can transmit messages, orders, and music. Causal processes – again, in contrast to pseudo-processes – can also transmit energy, momentum, electric charge, and various other physical quantities. A causal process is a process that persists on its own, without contributions from any outside source. Once a pulse of light is emitted from the neutron star, for example, it travels vast distances without any *external* influence. It is, so to speak, self-propelled. Pseudo-processes, in contrast, depend for their continuing existence upon something supplied from an external source. The spot of light created by the rotating beacon of the lighthouse will vanish almost immediately if the light is turned off.[17]

Causation in the Objects – Causal Interactions

When two processes intersect, there are two possibilities. On the one hand, both processes might be altered in ways that pass beyond the locus of the intersection. If, as in the example of the auto and the stone pillar, your car collides with the pillar, the car will carry scrapes and dents until you get it to the body shop to have it repaired. The pillar might be marked by some paint scraped off your car, or some of the stone might be chipped. This is a classic case of a causal interaction. On the other hand, if you avoid hitting the pillar with the car and only the shadow touches it, the shadow will be distorted (marked) at the place in which the intersection occurs, but this change will not persist beyond that location. This is an example of an intersection that does not qualify as a causal interaction.

As a second example, consider two airplanes that are flying on intersecting courses at different altitudes on a sunny day. Their shadows will intersect on the ground below, but no alteration of the shape will persist beyond the intersection. The shadows are pseudo-processes; they cannot interact with one another. A genuine causal interaction requires causal processes.[18] If the airplanes were traveling at the same altitude, the result would be a mid-air collision, and both planes would be altered – perhaps disastrously – in lasting ways.

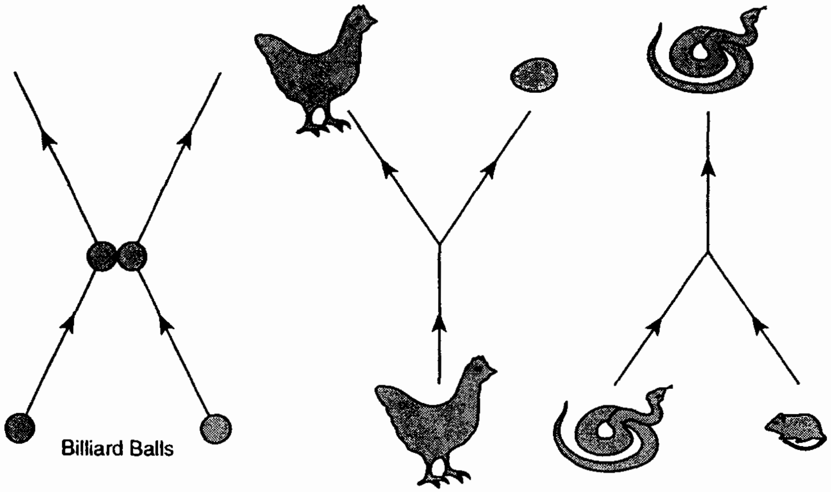

Figure 2.2 Types of causal interaction

The intersections mentioned so far have been cases in which two processes enter and two processes exit. Let's refer to them as *X-type intersections*. Two other basic configurations should be considered. Sometimes a single process splits into two parts. This sort of thing happens when a single-celled organism (e.g., an amoeba) splits into two by fission. We can call this a *Y-type intersection*. A hen laying an egg is another case. At first, we have a single organism, the hen, and later we have two entities, the hen and the separate egg. To be sure, the egg develops inside of the hen, but at some stage two processes exist. For simplicity, I've taken the moment of separation between the hen and the egg as the locus of the intersection.

A mirror image sort of intersection occurs when two processes come together and merge into one. This happens, for example, when a snake ingests a mouse. Following the strategy of the hen–egg example, we can say that the two processes have become one when the mouse is completely inside of the snake. The lower-case Greek letter lambda – which is somewhat similar to an inverted Y – serves as a handy schema. We can call this a *λ-type intersection*.

The criterion of mark transmission that we applied to the X-type intersection doesn't work very well for either the Y-type or the λ-type, so we need a different criterion for distinguishing mere spatial intersections from genuine causal interactions. At this point we have to introduce a tiny bit of basic physics. Don't panic – it's really simple. We've already mentioned energy, momentum, and electric charge; these are familiar examples of *conserved quantities* in physics. Take linear momentum, which is defined as the velocity of a body times its mass. The law of conservation of momentum states that momentum is neither created nor destroyed; for example, when Hume's billiard balls collide, the total momentum of the two balls before the collision is equal to their total momentum after the collision. What does happen in the collision is that momentum is *exchanged*. Prior to the collision, the momentum of the ball initially at rest is zero; in the collision some of the momentum of the moving ball is transferred to the one at rest. After the collision, the two balls retain their new quantities of momentum until some new interaction produces further change.

The point of this example can be generalized: *whenever processes exchange a conserved quantity in an intersection, that intersection qualifies as a causal interaction.* This criterion applies equally to the X, Y, and λ types of intersection. If, in any intersection of processes, the outgoing values of a conserved quantity in these processes differ from the incoming values, then, and only then, does the intersection of processes constitute a *causal interaction*. For example, when a hen lays an egg, the incoming process (the hen) has a different mass from either the hen or the egg after the two have been separated.[19] Likewise, the mouse and the snake each have a different mass from the mass of the snake that has swallowed the mouse.[20]

Note carefully an important philosophical point. In the preceding paragraph, the concept of *causal interaction* is explained entirely in terms of *noncausal* concepts. The key notions are *process* and *intersection*. The distinction between causal processes and pseudo-processes is *not* used. Intersection is essentially a geometrical notion (in four-dimensional space–time). On this approach, *causal interaction* is the most fundamental causal concept. *Very roughly speaking*, if processes intersect, and changes that persist beyond the locus of intersection arise, we have the most fundamental causal phenomenon – a causal interaction. Causal interactions *produce* changes; such changes are *propagated* by causal processes. The remaining question is what constitutes causal propagation.

Causation in the Objects – Causal Transmission

About 2,500 years ago, the Greek philosopher Zeno of Elea asked a simple question that turned out to be exceedingly difficult to answer. In fact, this was one of many paradoxes he posed. How can an arrow travel from the bow of the archer to its target? If it could actually move, then, at any place in its (supposed) path, it would be exactly where it is. It would be occupying a space equal to itself, so there would be no extra space in which to move. Moreover, at any moment or point of time, it is where it is. The moment is indivisible, so the arrow couldn't be at one place in one part of the moment and at another place at another part of the moment. It simply would have no space or time in which to move. At every point in its trajectory it would be at rest; therefore, it can't possibly move. Motion, Zeno concluded, is an illusion.

If you've studied even a little bit of differential calculus, you're likely to realize that, in calculus, it's easy to make a distinction between being in motion at a point

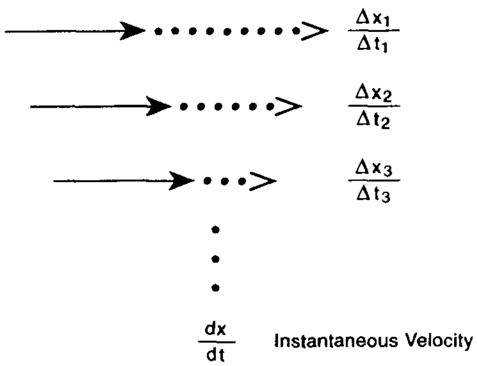

Figure 2.3 Zeno's Arrow

and being at rest at a point. (If you haven't studied calculus, don't worry; the point will become clear very soon.) Pick one point in the arrow, say its center of mass. Then, the motion of the arrow is defined by the position of that point at each moment of time. The calculus defines instantaneous velocity as the derivative of space with respect to time, i.e., dx/dt. The value of the derivative at time t is the instantaneous velocity at that time. If $dx/dt = 0$, the arrow is not moving at t.

Early in the twentieth century, Bertrand Russell showed that, because of the definition of the derivative, this answer to Zeno begs the question. To define the derivative at time t_0 we consider the distance the object travels in a *finite time span* Δt, which includes t_0. The ratio $\Delta x/\Delta t$ is its *average speed* in the *time span* Δt. The operation is repeated for smaller and smaller values of Δt. The derivative at the moment t_0 is the limit of the ratio $\Delta x/\Delta t$ as Δt goes to zero. Thus, the definition of the derivative requires consideration of motions over finite stretches of space and time, precisely the kind of motion Zeno claimed to be impossible. To define the derivative, we have to assume that the conclusion of Zeno's argument is false. This would only evade the problem, not answer it.

Russell then offered an alternative solution. As we have already noted, the motion of the arrow can be represented by noting the position of its center of mass at each moment in the duration of its flight. Russell proposed an "at–at" theory of motion. To say that the arrow moves from *A* to *B* *means* that it *occupies* each point in its trajectory at each corresponding moment of time. He doesn't say that it zips rapidly through these points. If you consider the arrow's state of motion at just one moment, without taking into account its position at any other time, the instantaneous velocity has no meaning.[21] If you ask how the arrow gets from the beginning of the path *A* to the midpoint *C*, he answers that it is by occupying each point between these two points at the appropriate time. If you ask how

the arrow gets from one point to the next, he reminds us that there is no *next* point – between any two points in its continuous path there are infinitely many others.

I find Russell's solution to the arrow paradox completely satisfactory. In addition, it suggests an analogous approach to the concept of causal transmission. Instead of an arrow, think of a bullet shot from a gun. As readers of mystery stories know, when the bullet leaves the gun, marks are made upon it that enable experts to identify the gun from which the bullet was shot. The moving bullet is a causal process; the marks are transmitted. Once the marks have been imposed by the interaction of the bullet with the gun, they remain on the bullet as it travels. To say that the mark is transmitted *means* that it is *at* the appropriate place in the process *at* the appropriate time. Also, the bullet transmits mass, a conserved quantity in this nonrelativistic context. It possesses a certain mass when it exits from the gun, and it continues to possess that same mass without any further interactions to replenish mass. The mass in question is *at* the appropriate place *at* the appropriate stage in the evolution of this process. Thus, we can adapt Russell's "at–at" theory of motion to an "at–at" theory of causal transmission. Note that the bullet transmits information – the marks identifying the gun from which it was shot. It also transmits causal influence. If the bullet strikes a person, it will produce a wound – possibly a fatal wound.

As an additional example of causal transmission, consider a pulse of white light sent out by a beacon. We noted earlier that, if a red filter is placed in the path of this white pulse, the light becomes red and remains red from that point on without any further interactions. The color red is *at* the appropriate place in this process *at* the appropriate stage of its travel.

Complete Causal Structures

Now that we have the three fundamental concepts – causal interaction, causal transmission, and causal process – at our disposal, we can answer the question about causation "in the objects." If we want to give an objective causal account of any spatiotemporal region of the universe, we must take account of all of the causal processes in that region and all of the interactions among them. This includes, of course, all of the causal processes entering and leaving this region. Recall the speeding bullet. As it travels from the gun to the target it collides with a huge number of molecules in the air. Light waves also strike it. This is an extremely complex set of processes and interactions. For most practical purposes much of this can be ignored, but if we want the full causal story, none can be left out. This is a pretty simple case. If we enlarge the sphere of discussion to include the operation of the gun and the effects on the target, the story becomes much more complex.

The complexity is even more evident in an example given earlier, namely, the starting of your car. Light rays impinging on your eyes let you see the ignition switch. A complicated set of nerve stimuli and muscular motions enables you to insert the key and turn it. An electrical contact is made that allows electricity to flow to the starter and to the spark plugs that ignite the fuel in the cylinders. The fuel injector must be activated, so that fuel enters the cylinders in the proper order; the timing mechanism has to coordinate the injection of fuel with the spark. This simple everyday situation is exceedingly complex when we consider the complete causal structure. Fortunately, seldom, if ever, do we need to appeal to all of the complexities of the complete causal structure. One manifestation of the possibility of omitting details of the complete causal structure is the fact that for some purposes a given process may be taken as simple, while for other purposes it is a complicated structure involving many processes and interactions. For a traffic engineer, the motion of a car along a street might be taken as a simple process. In contrast, as we have just seen, for an automotive engineer, the car is an exceedingly complicated set of processes and interactions.

To be quite clear on the status of the complete causal structure, let us recapitulate its construction. Philosophically speaking, the first step is the definition of *causal interaction*. As already mentioned, this concept was introduced without the aid of any other causal concept. We used the notion of a process, without making a distinction between causal processes and pseudo-processes. We used the entirely geometrical concept of an intersection, without distinguishing causal interactions from mere noncausal intersections. We referred to changes in properties, without presupposing that such changes are causal. We then defined causal interactions as intersections of processes in which changes occur that persist beyond the locus of intersection. When we think in terms of marks, we can say that interacting processes produce marks in one another that persist beyond the marking location. When we think in terms of conserved quantities – which, following Dowe, I consider preferable to marks – we say that causal interactions are intersections in which conserved quantities are exchanged.

The next causal concept is *transmission*; *causal interaction* is the only causal concept we use to introduce it. We say that a mark is transmitted by a process if it is present in the process beyond the point of introduction by an interaction *without any additional interactions*. We say that a conserved quantity is transmitted by a process over a finite interval if that process possesses a certain amount of that quantity over that span *in the absence of any additional interactions within that interval*.

Our third basic causal concept is *causal process*, the only causal concept required is *causal transmission*. A process is causal if it is capable of transmitting a mark, or if it actually transmits a conserved quantity. Since you might think that *capability* is a further causal concept, I much prefer the conserved quantity alternative.

Using these three causal concepts, introduced in the manner just described, we have everything necessary to define a *complete causal structure*. The result, in my view, is a genuine characterization of causation *in the objects*.

Causes and Effects

You may have noticed that, in sections 3–6, the terms "cause" and "effect" were hardly used at all. Instead, we referred to processes, interactions, and transmission. The complete causal structure is characterized entirely in terms of these latter concepts. However, the "cause–effect" terminology occurs frequently in everyday life and in science. It is used mainly to select those parts and aspects of the complete causal structure that are relevant to a given situation. This means that cause–effect relations are context dependent, and for that reason, they are not independent of human knowledge, interests, and desires. Thus, while processes, interactions, and transmission are "in the objects," cause–effect relations are not entirely so because of their context-dependence.

Two sorts of configurations are commonly regarded as cause–effect relations. The first involves simply the interaction between two (or more) causal processes. Hume's example of the colliding billiard balls is a good illustration – in Hume's words, "As perfect an instance of the relation of cause and effect as any which we know either by sensation or reflection" ([1748] 1955, p. 186). Each ball is a causal process; the collision is a causal interaction. The motion of each ball is changed in the intersection, and these motions persist after the collision. Notice how much is left out of Hume's description. The interactions between the air molecules and the balls are ignored. The friction of the balls moving on the billiard table is also ignored. Even the spin on the moving ball prior to the collision is left out. Hume tells us that the second ball is initially at rest, but afterward is moving, but he tells us nothing about the motion of the first ball after the collision. That depends sensitively on its spin before the collision. Hume has told us all we need to know to understand his example and the point he is trying to make.[22]

The second common configuration involves two events connected by a causal process. Suppose that some children are playing baseball in a vacant lot. The child at bat hits a ball that crashes through a window of a neighboring house. One event is the collision of the bat with the ball. The ball, traveling from the bat to the window, is a causal process. The second event is the ball striking and breaking the window. We can invoke Hume's remark again. "This is as perfect an instance of the [second type] of the relation of cause and effect as any which we know. . . ." Note again how much is left out. The collisions with air molecules are omitted; so are the positions of the glass shards on the floor of the house.[23] Suppose the batter, observing the direction of the ball, shouts "Oh, no!" (or words to that effect). The sound waves might reach the window at the same time as the ball; however, they would not be considered relevant.

As Mackie pointed out in connection with his INUS conditions, we often make causal judgments in terms of what is usual or unusual in a given case. If dropping cigarettes on the floor were customary in some particular barn and dry straw on the floor very unusual, we might say that the presence of the straw was the cause of the fire. Mackie strongly emphasizes the context dependence of such causal judgments.

Having discussed the context dependence of cause–effect language, I must now emphasize the fact that, ***given a particular context***, cause–effect relations may be entirely objective. If water comes through a hole in your roof and damages some of your books, the roofer needs to find the hole and fix it. You need to find out how much it will cost to replace the damaged books. Given the context, each of these questions has an objective answer.

I have no intention of suggesting even for one moment that cause–effect terminology is defective or should be banished; quite the contrary. My principal aim has been to establish the objective nature of causation, and to investigate the ways in which our causal claims can legitimately be applied in the contexts in which they are important. Obviously, we want objective answers to questions about causes of airplane crashes and the onset of diseases, as well as a plethora of other phenomena. Causal explanation sheds light on questions of this sort.

Causal Explanation

It is pretty obvious that causal knowledge is sought both for intellectual understanding and for practical control. As already mentioned, most scientists investigating the extinction of dinosaurs now believe that the collision of a massive body – a comet or an asteroid – with the Earth produced atmospheric conditions under which the dinosaurs and many other species could not survive. This is a causal explanation of the extinction, and one I find extremely interesting, but I don't see any practical application of this piece of causal knowledge. Nevertheless, it is satisfying to understand the history of our planet and the forms of life that have inhabited it.

In a more practical vein, airplanes crash, leading to death and destruction. We want to discover the causes of such accidents in order to prevent them in the future. If a crash is a result of pilot error, we seek to understand the error in order to show other pilots how to avoid it. If a crash is caused by wind sheers, we seek better ways of detecting them, and we issue orders that places where they are occurring should be avoided. There is nothing we can do about the wind sheer; that is beyond our control, but measures can be taken to lessen their danger to life and property. If the cause is a mechanical failure, we try to ascertain its precise nature in order to make modifications that will prevent such failures in the future.

Where diseases are concerned, causal understanding may enable us to prevent or cure them. Smallpox has been eliminated from the human population by means of an effective vaccine. Antibiotics can cure many kinds of infections. The knowledge that diseases are caused by germs, not mysterious vapors, has been enormously beneficial to humans and other animals.

Causal information plays a crucial role in assigning legal or moral responsibility. If, for example, someone is injured as a result of falling down a flight of stairs, it's essential to find out whether the victim was pushed by an enemy or tripped on a piece of loose carpet. In the former case, the other person is responsible; in the latter case, the landlord would be at fault. Often, of course, many causes combine to bring about an event. If two cars collide at an intersection, both drivers may be equally responsible because of failure to pay attention to other traffic.

Causal knowledge is useful, not only for preventing undesirable states of affairs, but also for producing desired results. During World War I, it was found that soldiers whose wounds were infested with maggots had a survival rate greater than those whose similar wounds were free of maggots. Disgusting? Perhaps, but that's not so important where life and death are at issue.[24]

Further Topics

In this article, we have looked at causation at a fundamental metaphysical level. Almost all of the examples have been taken from physical science. Two questions remain open. First, we have omitted consideration of social causation. Is there such a thing as causation where social institutions are involved, and does it differ from the physical causation we have discussed? One point is surely true. There could be no social causation without physical means of communication. This is an important area for application of transmission via causal processes. However, I am not drawing reductive conclusions. I am *not* saying that all social causation can be reduced to physical causation of the sort I've discussed. Since I know of no satisfactory answer to the mind–body problem,[25] I must, in honesty, remain agnostic on this issue. It is a question for philosophers of psychology and the social sciences to confront. Second, in the social and biological sciences, functional explanations play important roles – e.g., elephants have huge ears because they fulfill the function of controlling body temperature. Can such explanations be analyzed in terms of physical causation?[26]

During the closing decades of the twentieth century, philosophers became increasingly aware of the need to analyze *probabilistic* or *indeterministic* forms of causation. This might initially seem like an incoherent or self-contradictory concept, but it is commonly used in science and everyday life. In discussing Mill's method of concomitant variation, for example, I mentioned the taking of aspirin as a preventative measure against heart disease. I pointed out that it does not prevent heart disease in every case, but it lowers the chances of contracting that malady. Similarly, we have known for many years that heavy cigarette smoking causes lung cancer, but not every heavy smoker is a victim. In fact, in the daily papers and TV news, we frequently find reports of some new medication designed to prevent or cure some disease, where it is obvious that, at best, it will be effective for some percentage of people who try it.[27] We have not been able to pursue this topic; there is an extensive and relatively technical body of literature on this subject.[28]

Twentieth-century physics has established beyond reasonable doubt that causation, at least as we naturally think of it, does not hold in the realm of quantum mechanics. My personal opinion is that there are noncausal mechanisms, which we don't really understand, that operate in the quantum domain.[29] Consequently, our analysis of causation does not apply even to all domains of our world, let alone in all possible worlds.

Our aim has been to learn what causality *is* in this world, to whatever extent it operates. Ours has been an exercise in empirical metaphysics, not an analysis of the uses of linguistic expressions. Naturally, we need some understanding of the meanings of the words "cause," "effect," and their cognates as they are actually used in English in order to be confident that we are dealing with the right concept. It must not turn out, for example, that lung cancer causes heavy smoking or that behavior of the barometer causes storms. Given this sort of initial semantic information, we have tried to find out precisely what our causal terms designate. We have sought to understand basic facts about how the world works, not semantic truths about the language used to describe it. I hope that we have discovered answers to the profound questions Hume raised.

Notes

1 Hume's *Enquiry Concerning Human Understanding* is an excellent classic text for beginning students of philosophy. Sec. IV, Pt. 1 and Sec. VII, Pts. 1–2 contain his central thoughts on causation.

2 Logical purists who object to the notion of entailments between descriptions may substitute phrases such as "a statement that an event of a certain description occurs" or "a statement that an object satisfying a certain description exists" wherever it is required in this essay.

3 Descartes, who is usually considered the first modern philosopher, used a priori causal principles in his proof of the existence of God. This comes at an absolutely fundamental point in his philosophy. He maintains that, by "the natural light of reason," we know that a cause must have at least as much reality as its effect. His book *Meditations*, in which that argument is offered, is another excellent classic text for beginning students of philosophy.

4 Of course, relational properties, such as *being the cause of* or *being the effect of* something else are not allowed as part of the description; to put the matter more generally, the properties in the description are confined to those that are observable by the time the event in question occurs.

5 My dissertation advisor, Hans Reichenbach, reported an incident from his own experience. He was sitting in a theater in Los Angeles watching a film. Just as a major explosion occurred *in the film*, the theater began to shake. Instinctively, he said, he felt that the explosion depicted on the screen caused the theater to tremble. What happened in fact was that a minor earthquake occurred by chance just when the explosion occurred on the screen.

6 Kant said that Hume had "awakened him from his dogmatic slumbers."

7 In saying that these propositions are self-evident or necessary, Kant did not mean to say that no other logically consistent type of geometry could exist. If he had said that Euclidean geometry is the only consistent geometry, he would have to have concluded that the statements of Euclidean geometry are *analytic*, not synthetic. He was saying, instead, that Euclidean geometry provides the only framework in which we can visualize spatial relations among the objects in our world. Shortly after Kant's death, non-Euclidean geometries were discovered and shown to be consistent (if Euclidean geometry is consistent), but this did not refute Kant's thesis. However, the use of non-Euclidean geometry in Einstein's general theory of relativity to describe the physical space of our universe shows that spatial relations can have a non-Euclidean structure.

8 Smith (1933), p. 218.

9 If you have gas appliances in addition to a gas furnace, your gas consumption will not go to zero even in the warmest months.

10 Mill made this point in the nineteenth century.

11 Mackie requires an analysis of *causal* priority that is independent of *temporal* priority. This additional requirement does not affect our discussion of the "regularity part" of his theory.

12 David Lewis (1973) is the most famous advocate of the counterfactual account.

13 I am referring, of course, to the "regularity part" of Mackie's analysis.

14 Donald Davidson (1967) argued this thesis convincingly, in my opinion.

15 We will use the word "process" as a noun; however, it also occurs in everyday usage as a verb or an adjective. For instance, a processor in a computer *processes* information in some useful way; I'm writing this article using a word-processor. Process cheese is the result of mixing different kinds of cheese together. I've used *Webster's Ninth New Collegiate Dictionary* (1989) as the basis of these remarks about usage, but I've consulted other dictionaries as well.

16 Light travels at different speeds in different media; the maximum speed is the speed of light in a vacuum.

17 I say "almost" because the last bit of light from the beacon takes some time to travel from the source to the wall. As a pretty good approximation, we can say that light travels at one foot per nanosecond. In the case of the pulsar, "almost immediately" would be a gross error because it takes light 6,500 *years* to get to us.

18 However, causal processes can intersect without interacting. If two light rays intersect, they *interfere* in the locus of the intersection, but they continue beyond as if nothing had happened.

19 In this nonrelativistic situation, mass can safely be taken as a conserved quantity. This supposition is retained throughout the remainder of this essay.

20 The conserved-quantity approach to physical causation was introduced by Phil Dowe. His *Physical Causation* (2000) presents his most recent and most fully developed account.

21 Russell claims that these considerations actually vindicate Zeno. However, Russell points out, even if the arrow is at rest at each point of its trajectory, it doesn't follow that it is always in the same place. To say that what applies to each member of a class holds of the class itself is an example of the elementary fallacy of composition.

22 To a serious player of billiards, the behavior of the first ball after the collision is crucial.

23 If we were concerned with the curve ball thrown by the pitcher, the interactions with the air molecules would be highly relevant.

24 The practical applications of causal knowledge are so prominent that some philosophers have attempted to analyze causality in terms of manipulability; Gasking (1955) is the classic source.

25 This problem is eloquently stated by Descartes in his *Meditations*.
26 Wright (1976) defends the affirmative answer.
27 Some TV ads state explicitly, "Individual results may vary."
28 Patrick Suppes (1970) is the classic starting point for current discussions. Salmon (1980) contains a critical survey of various approaches.
29 See d'Espagnat (1979) for a clear and elementary discussion of the quantum situation. Mermin (1981, and more colorfully in 1985) offers another clear exposition.

Bibliography

Davidson, Donald (1967). "Causal Relations," *Journal of Philosophy*, 64, pp. 691–703.
Descartes, *Meditations*.
d'Espagnat, Bernard (1979). "The Quantum Theory and Reality," *Scientific American*, 241, No. 5 (November), pp. 158–81.
Dowe, Phil (2000). *Physical Causation*. Cambridge: Cambridge University Press.
Fair, David (1979). "Causation and the Flow of Energy," *Erkenntnis*, 14, pp. 219–50.
Fales, Evan (1990). *Causation and Universals*. New York: Routledge.
French, Peter A., et al., eds. (1984). *Midwest Studies in Philosophy*, 9. Minneapolis: University of Minnesota Press.
Gasking, Douglas (1955). "Causation and Recipes," *Mind*, 64.
Hume, David [1739–40] (1888). *A Treatise of Human Nature*. Ed., L. A. Selby-Bigge. Oxford: Clarendon Press.
——. An Abstract of *A Treatise of Human Nature*, in Hume [1748] (1955), pp. 183–98.
——[1748] (1955). *An Enquiry Concerning Human Understanding*. Indianapolis: Bobbs-Merrill.
Lewis, David (1973). "Causation," *Journal of Philosophy*, 70, pp. 556–67.
Mackie, John L. (1974). *The Cement of the Universe*. Oxford: Clarendon Press.
Mermin, N. David (1981). "Quantum Mysteries for Anyone," *Journal of Philosophy*, 78, pp. 397–408.
——(1985). "Is the Moon Really There When Nobody Looks? Reality and the Quantum Theory," *Physics Today*, 38 (April), pp. 38–47.
Mill, John Stuart (1843). *A System of Logic*. London: John W. Parker.
Russell, Bertrand (1943). *The Principles of Mathematics*, 2nd ed. New York: W. W. Norton.
——(1948). *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.
Salmon, Wesley C., ed. (1970). *Zeno's Paradoxes*. Indianapolis: Bobbs-Merrill.
——(1980). "Probabilistic Causality," *Pacific Philosophical Quarterly*, 61, pp. 50–74.
——(1998). *Causality and Explanation*. New York: Oxford University Press.
Smith, Norman Kemp (1933). *Immanuel Kant's Critique of Pure Reason*. London: Macmillan.
Sosa, Ernest, and Michael Tooley, eds. (1993). *Causation*. Oxford: Oxford University Press.
Suppes, Patrick (1970). *A Probabilistic Theory of Causality*. Amsterdam: North-Holland.
Wright, Larry (1976). *Teleological Explanations*. Berkeley: University of California Press.

Chapter 3

What Events Are

Jonathan Bennett

1 Introduction

The furniture of the world includes planets and pebbles, hopes and fears, fields and waves, theories and problems, births and deaths. As metaphysicians, we want to understand the basic nature of these and other kinds of item; and my topic is the basic nature of births and deaths – more generally, of *events*. If events are things that happen, what differentiates them from sticks and stones, which are things that exist but do not happen? Do events constitute a fundamental ontological category, or is our event concept just a way of organizing material that could be handled without its aid?

With questions like those in the background, I ask: what sort of things are events? Locke and Leibniz knew the answer to this; then Kim rediscovered it; but his rediscovery did less good than it might have because it was ambushed by an error. I shall explain.

A sparrow falls. That fall of that sparrow is a particular, located in space and time. It occurs where the sparrow is when it falls, and it occurs just then. It is, then, closely linked to the sparrow, and even more closely to the fact that the sparrow falls there and then. Witness the opening of this paragraph, where I said that a sparrow falls, and went straight on to speak of "that fall." That the fall exists (= occurs) is a logical upshot of the fact that the sparrow falls. Every event results logically from some such underlying fact: there was a fight because some animals fought, there was a storm because wind and water moved thus and so. In section 12, I shall discuss the rival view that some animals fought because there was a fight.

What metaphysical categories have a role in the fact that a certain sparrow fell? Can any of them be identified with *the sparrow's fall*? I shall consider five candidates: a fact, a thing, a temporal part of a thing, a property, and a property-instance.