appears at all the places occupied by the spot. It would seem that the spot 'transmits' energy, since transmission is to be understood in terms of the 'at-at' theory. This is an objection that will have to be answered, as we turn in the following chapter to an attempt to characterise causal processes in terms of properties such as energy and momentum.

In the previous chapter we considered the twin notions of causality as persistence, such as the kind of causality that is related to having identity through time; and causality as transference, where one body affects another. It was argued that an adequate account of causation should account for both modes, and the transference theory was found to be wanting on the score that it rules out the possibility of causality as persistence. At the start of this chapter, in the discussion of Russell's account of causal lines, it was noted that the idea of a causal line captures what was missing in the transference theory – a notion of immanent causality – but that if causal lines are the whole story on causality, then that account loses as much as it gains, for it now ignores transient causality.

Salmon's account of causation in terms of causal processes and interactions allows nicely for both modes of causality. The concept of a causal process fits perfectly with the kind of causation involved in identity over time, or the immanent causality displayed by the spaceship moving by its own inertia. On the other hand, the concept of a causal interaction accounts well for the idea of transference – since two processes are mutually modified in an interaction. On this score, at least, Salmon's theory appears to be superior to the others discussed. Even if the mark theory itself is inadequate, the approach of treating causality as a characteristic of processes and interactions seems to be the right one.

Phil Dowe, Physical Causation, 2000.

5

The Conserved Quantity Theory

It is profitable to distinguish three key questions about causation. The first question is *what are causal processes and interactions?* I follow Salmon in the view that it is advantageous to focus on this question rather than on more traditional questions about causation. As we have seen, the key task in addressing this question is to distinguish causal from pseudo processes. In this chapter I offer an account of causal processes and interactions that, I argue, adequately makes this distinction. I show how this account answers a range of objections, in comparison to other theories, in particular to Salmon's theory and his recent revisions.

The second question – *what is the connection between causes and effects?* – is not addressed in the present chapter. In Chapter 7, however, I discuss the kind of answer one can give to this second question if one accepts the results of the present chapter; and I defend that answer against its rivals. The third question is *what distinguishes a cause from its effect?* In Chapter 8 I discuss the kind of answer one can give to this if one accepts the results of the present chapter; and I defend that answer against its rivals. It is important to emphasise that the account of causal processes and interactions given in this chapter is not intended to address the second and third questions.

In this chapter an outline of a theory of causal processes and interactions is presented. The approach to be taken is to modify Salmon's theory by introducing the concept of a *conserved quantity*. The central idea is that it is the possession of a conserved quantity, rather than the ability to transmit a mark, that makes a process a causal process. Insofar as it links causation to quantities like energy and momentum, this account also bears some resemblance to the transference theory.

We begin this section with an outline of this Conserved Quantity (CQ) theory,[1] which will be followed by some comments expanding on the intended meaning of the terms used, and some examples.

The conserved quantity theory can be expressed in just two propositions:

CQ1. A *causal process* is a world line of an object that possesses a conserved quantity.

CQ2. A *causal interaction* is an intersection of world lines that involves exchange of a conserved quantity.

A *process* is the world line of an object, regardless of whether or not that object possesses conserved quantities. A process can be either causal or noncausal (pseudo). A *world line* is the collection of points on a spacetime (Minkowski) diagram that represents the history of an object. This means that processes are represented by elongated regions, or 'worms,' in spacetime. Such processes, or worms in spacetime, will normally be timelike; that is, every point or time slice on its world line lies in the future lightcone of the process's starting point. However, it is at least conceivable that the world line of an object may sometimes appear on a spacetime diagram as a spacelike worm. One example of this is the short-lived string. Imagine that a 1,000-mile-long string extended roughly in a straight line spontaneously comes into existence, but then is annihilated one millisecond later. This short-lived string will be a worm in spacetime, but it is not extended far in time. But that worm is the world line of an object, so it is a process on the present account. Another example of particular relevance is a case of a pseudo process, like the spot moving along the wall, which can travel faster than the speed of light. In the case where it does in fact travel faster than the speed of light the process is represented by a spacelike worm. This also counts as a process. Thus, on the present account a process is a worm in spacetime, be it timelike or spacelike; just provided that worm is the world line of an object.

1. As originally given in Dowe (1992c), but including some slight modifications prompted by Salmon's (1994) analysis. Although Brian Skyrms, in his 1980 book *Causal Necessity* (1980: 111), was the first to suggest a conserved quantity theory, the first detailed conserved quantity theory did not appear until 1992 (Dowe 1992a; 1992c).

An *object* is anything found in the ontology of science (such as particles, waves and fields), or common sense (such as chairs, buildings and people). This will include noncausal objects such as spots and shadows. A process is the object's trajectory through time. That a process is the world line of an object presumes that the various time slices of the process each represent the same object, at different times; thus it is required that the object have identity over time. The requirement of identity over time of an object rules out certain worms in spacetime: not every worm in spacetime counts as a process, for not every worm in spacetime is the world line of an object. One type of worm that does not qualify as a process is a timewise gerrymander – an alleged object defined in different ways at different times (see section 5.3). On the present account a timewise gerrymander is not a process, for it is not the world line *of an object*, since objects must exhibit identity over time. Thus Quine's characterisation of a physical object as an intrinsically determinate portion of the spacetime continuum (Quine 1965: 229–231) will not suffice, since it admits as objects timelike gerrymanders.

Worms in spacetime that are not processes I call, borrowing Kitcher's (1989) terminology, 'spatiotemporal junk.' Thus a line on a spacetime diagram represents either a process or a piece of spatiotemporal junk, and a process is either a causal or a pseudo process. In a sense, what counts as an object is unimportant; any old gerrymandered thing qualifies (except timewise gerrymanders). In the case of a causal process, what matters is whether the object possesses the right type of quantity. A shadow, for example, is an object, but it does not possess the right type of conserved quantities; a shadow cannot possess energy or momentum. It has other properties, such as shape, velocity and position, but possesses no conserved quantities.[2]

A *conserved quantity* is any quantity that is governed by a conservation law, and current scientific theory is our best guide as to what these are. For example, we have good reason to believe that mass-energy, linear momentum, and charge are conserved quantities (see section 5.2).

An *intersection* is simply the overlapping in spacetime of two or more processes. The intersection occurs at the location consisting of all

2. The theory could be formulated in terms of *objects*: there are causal objects and pseudo objects. Causal objects are those that possess conserved quantities, pseudo objects are those that do not. Then a causal process is the world line of a causal object.

the spacetime points that are common to both (or all) processes. An *exchange* occurs when at least one incoming, and at least one outgoing process undergoes a change in the value of the conserved quantity, where 'outgoing' and 'incoming' are delineated on the spacetime diagram by the forward and backward light cones, but are essentially interchangeable. The exchange is governed by the conservation law, which guarantees that it is a genuine causal interaction. It follows that an interaction can be of the form X, Y, λ, or of a more complicated form.

'*Possesses*' is to be understood in the sense of 'instantiates.' An object possessing a conserved quantity is an instance of a particular instantiating of a property. We suppose that an object possesses energy if science attributes that quantity to that body. It does not matter whether that process transmits the quantity or not, nor whether the object keeps a constant amount of the quantity. It must simply be that the quantity may be truly predicated of the object.[3]

As expressed in the two propositions just given, the CQ theory aims to provide an answer to the first question, viz., what are causal processes and interactions? In particular, it aims to distinguish causal from pseudo processes, and it does this by distinguishing objects that possess conserved quantities from those that don't. As in Salmon's theory, causality is treated fundamentally as a property of processes and interactions.

We may also include a broader sense of 'causal process,' where a series of causal processes and interactions form a unified sequence. Sound waves and water waves will qualify as causal processes in this sense.

We now turn to some examples.

Example 1. Consider a transmutation reaction where a nitrogen atom ($^{14}_{7}N$) is hit by an alpha particle ($^{4}_{2}He$), producing an oxygen atom ($^{17}_{8}O$) and a proton ($^{1}_{1}H$). The nuclear equation (with Q representing the extra energy needed for the interaction) is given by

$$^{4}_{2}He + ^{14}_{7}N + Q \rightarrow ^{17}_{8}O + ^{1}_{1}H$$

3. In previous formulations (Dowe 1992a: 126; 1992b: 184; 1992c: 210) the word 'manifests' was used in place of 'possesses,' but, as D. M. Armstrong has pointed out, this gave the misleading impression that the quantity had to be experienced by human observers (personal communication). This improvement is also suggested by Salmon (1994).

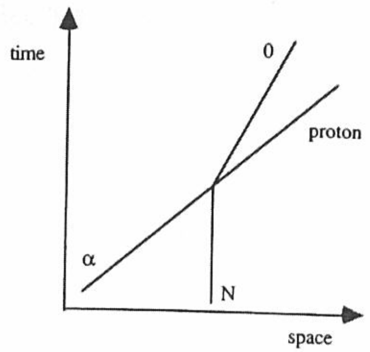

Figure 5.1. Transmutation.

The subscript is the atomic number – the number of protons in the atom, which determines the element type. The atomic number also is the charge of the atom, since the number of protons is equal to the number of electrons. The superscript is the atomic mass, which can vary for an element. The spacetime diagram for this is shown in Figure 5.1.

By definition CQ2, this reaction is a causal interaction, because we have the intersection of world lines where charge, represented by the subscripts, is exchanged. Amongst other things one unit of charge is transferred from the α particle to the nitrogen atom, changing it in the process. So each of the processes involved is a causal process by definition CQ1, because they each possess charge. Note that the N-atom possesses a net charge of zero. On the present view, however, this still counts as possessing a conserved quantity. There is a difference between possessing a zero sum of a quantity and being the sort of object that does not possess conserved quantities. (More on this later.)

Example 2. An example of a Y-type interaction is the decay of radium-226 to radon mentioned earlier, as shown in Figure 5.2:

$$^{226}_{88}Ra \rightarrow ^{222}_{86}Rn + ^{4}_{2}He$$

This qualifies as a causal interaction by CQ2 because there is an exchange of charge, where the charge of the incoming process is divided between the two outgoing processes. The three processes involved are all causal by CQ1 because they each possess charge.

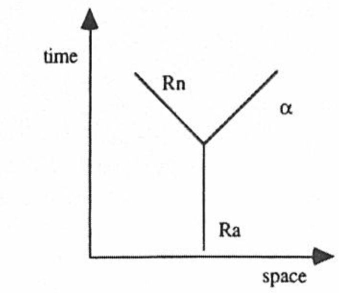

Figure 5.2. Radium-226 decay.

Example 3. As an example of a pseudo process, consider the phase velocity of water ripples. When a stone is dropped into water, the individual waves may travel faster (phase velocity) than the total group of waves (group velocity). Then at the leading edge waves will disappear, while waves will appear at the rear. It is physically possible for phase velocities to travel faster than light, but they cannot be used to convey signals. On our theory these types of phase velocities are not causal processes because they do not possess energy or momentum, or any conserved quantity. The energy, momentum and power of the wave travel at the speed of the group velocity. To generalise, pseudo processes do not possess the type of physical quantities that are governed by conservation laws. Shadows, intersections of rulers and so on do not possess conserved quantities.

V.2 CONSERVED QUANTITIES

A conserved quantity is any quantity that is governed by a conservation law, and current scientific theory is our best guide as to what these are: quantities such as mass-energy, linear momentum, and charge. The idea that the quantities associated with causation are conserved quantities is a suggestion that I present here simply as a plausible conjecture. I have argued in Chapter 3 against Aronson's idea that velocity and certain other physical quantities are the right quantity, and in Chapter 4 against Salmon's idea that an ability to transmit a mark is the right property, but I have no real quarrel with Fair's position that it is energy/momentum. I simply offer the conjecture that other conserved quantities, such as charge, may also serve the function.

Conservation laws play the role of identifying which quantities are significant for causation. The claim is not that certain quantities are locally conserved in an interaction or by the process in the absence of interactions, although that will follow. Rather, the account focuses on those quantities that are globally or universally conserved, and connects causality simply to the possession of those quantities.

It is important that conserved quantities be understood in a way that does not appeal to causation, or else circularity threatens. It is common to define conservation in terms of constancy within a closed system. Now if a closed system is simply one with no external causal interactions, that is, a system causally isolated from all others, then we face an immediate circularity. The idea is fine as a rule of thumb – that is, it is true – but it cannot work as an analysis. Instead, we need to explicate the notion of a closed system in terms only of the quantities concerned. For example, energy is conserved in chemical reactions, on the assumption that there is no net flow of energy into or out of the system.

It is important to note that the reference to current theories does not relativise causation to human knowledge – the point is simply that current theories are our best guide to what the conservation laws are. The reason that we cannot simply define a conserved quantity as one that is universally conserved is that some quantity may be accidentally conserved, and such a quantity should not enter into the analysis of causation. Further, regularities are not by any means the only form of evidence about conservation laws – theoretical considerations are also important.

The identity of 'causal process' with 'the world line of an object that possesses a conserved quantity' is contingent, and not metaphysically necessary. The hypothesis is that in our world, and in close enough worlds, such as most of those that obey our laws, a causal process is the world line of an object that possesses a conserved quantity. We leave aside the question of how far we can stray from actuality before this hypothesis stops making sense. In calling this an empirical analysis (see Chapter 1), we emphasise the priority of the claim that the identity holds in actuality. In calling the analysis a contingent identity, we mean that it is contingent on the laws of nature and perhaps even on boundary conditions.

In particular, the theory does not purport to tell us what happens to the identity in distant merely possible worlds. Suppose $\{q_a, q_b, q_c, q_d\}$ is the complete set of conserved quantities in the actual world W_a, and consider a world W_e where none of this set is in fact conserved, and where a conservation law holds instead for q_e. Is the world line of an object in W_e that possesses q_e but none of the set of quantities conserved in W_a a causal process? Or again in W_e, is the world line of an object that possesses q_a, say, but none of the quantities conserved at W_e, a causal process? The answer in both cases is that the theory does not say.

The theory may tell us about closer worlds – for example, those with the same conservation laws as ours. In a world where q_a is conserved, but there is only one object that possesses q_a, the world line of that object is a causal process. Thus the account is not a (Humean) actual-regularity account.

This raises the question of whether the theory is a singularist account (ontologically, not conceptually). I say the account is singularist in the following sense: a particular causal process is not analysed in terms of laws about that type of processes; rather, that a type of process is causal is a matter of generalisation over the particular instantiations of that process-type. The particular is basic.

Thus whether something is a causal process depends only on local facts about the process, namely, the object's possession of a certain kind of physical quantity. It does not depend on what happens elsewhere in the universe, so in that sense being causal is an intrinsic property of a process.

Is this a supervenience (i.e. nonsingularist) account in the sense (e.g., Tooley) that whether the world line of *a* is a causal process supervenes on whether *a* possesses a quantity *q* such that there is a law governing *q*? No, no such claim has been made. The theory simply says that at this world, just if an object possesses one of the quantities that is actually conserved, then the world line of that object is a causal process. This is a local, particular matter.

Alexander Reuger (1998) has argued that in some general relativistic spacetimes, on the Conserved Quantity theory, it is not a local matter whether a process is causal. Reuger points out that in general relativity, global conservation laws may not hold. In the nonrelativistic case a differential conservation law such as the electrodynamic continuity equation:

$$\text{div}\,\mathbf{j} = -\partial/\partial t\ \rho$$

(where $\mathbf{j}$ is the current density vector [the amount of electric charge moving through a unit volume in a unit time], such that $\mathbf{j} = \rho v$, where

v is the charge velocity and ρ is the charge density) entails, via Gauss's theorem, the integral conservation law:

$$\partial/\partial t \int \rho dV = -\int \mathbf{j}\ \mathbf{n}\ dS$$

for a surface S of a volume of integration V. The differential is the local, the integral the global form of the conservation law.

In the general relativistic context, however, a differential conservation law holds for energy-momentum,

$$\nabla^a T_{ab} = 0$$

for the covariant derivative ∇^a, given Einstein's field equations. But unless spacetime possesses special symmetries, there will be no integral formulation. Reuger concludes that whether conservation laws hold is contingent on the global properties of spacetime, and that the choice is therefore either to insist that causation is intrinsic, and that there are no genuine causal processes, or to abandon the intuition that causation is intrinsic to a process or event.

However, there is a third option, which follows from what I have already said. The Conserved Quantity theory is a contingent hypothesis, contingent on the laws of nature, for example. This means if the laws turned out to be a certain way, the theory would be refuted. This may be the case if it turns out that there actually are no conservation laws.

But the fact that there are general relativistic spacetimes in which global conservation laws do not hold does not entail that global conservation laws fail in our world. Whether they do or not depends on the *actual* structure of spacetime, and in particular whether certain symmetries hold. As I understand it, our spacetime does exhibit the right symmetry; global conservation laws do hold in our universe as far as we know. I take it, then, that the conserved quantity theory is not refuted.

I have suggested that the account should probably hold in all physically possible worlds, that is, in all worlds that have the same laws of nature as ours. Has Reuger shown that this is not so? Not at all. To say, for example, that nonsymmetric spacetimes are possible can be misleading. It means simply that it is a solution to the equations of the General Theory of Relativity. But this doesn't mean that such a world is a physically possible world in the sense given here. If such a world violates other laws that hold in the actual world, then that world is not physically possible. This is exactly what we have in these nonsymmetric spacetimes. Symmetries and conservation laws that hold in the actual world break down, so it is not a physically possible world in my sense.

Therefore we need not give up on the Conserved Quantity theory, understood as a contingent hypothesis, nor on the idea that causation is actually intrinsic.

V.3 POSSESSION, TRANSMISSION AND GERRYMANDERED AGGREGATES

In his criticism of the Conserved Quantity theory just presented, Salmon (1994: 308) offers an argument (see also 1984: 145–146) for requiring 'transmits' rather than just 'possesses': Consider a rotating spotlight spot moving around the wall of a large building. This is a classic case of a pseudo process: in theory such a spot could move faster than the speed of light. But the spot manifests energy at each point along its trajectory. Therefore, Salmon's argument goes, we need more than just the regular appearance of energy to characterise causal processes; we need the notion of transmission. In this section I show how the CQ theory avoids this problem without appealing to any notion of transmission.

A spot or moving patch of illumination does not possess conserved quantities. A moving spot has other properties: speed, size, shape and so on; but not *conserved* quantities such as energy or momentum. What possesses the energy that is 'regularly appearing' is not the spot but a series of different patches of the wall. The spot and the patch of wall are *not* the same object. The patch of wall does not move. It *does* possess conserved quantities, *its* world line does constitute a causal process, and *it* is not capable of moving faster than the speed of light. The spot does move, but does *not* possess energy and *is* capable of moving faster than the speed of light. Therefore 'whether or not an object possesses a conserved quantity' is an adequate criterion for distinguishing causal from pseudo processes.

Hitchcock (1995) provides another example of the same objection, where a shadow moves across a charged plate, at every stage manifesting a conserved quantity, charge. The answer to this is the same as for the spot of light.

Salmon (1994: 308) gives an ingenious counterexample to this answer, asking us to consider "the world line of the part of the wall surface that is absorbing energy as a result of being illuminated" (1994: 308). This "gerrymandered" object is the aggregate of all the patches of wall that are sequentially illuminated, taken only for the time that they are being illuminated. Salmon argues that this object does possess energy over the relevant interval, but does not *transmit* energy. The implication is that the world line of this object is not a causal process, yet the object possesses energy; therefore we need to invoke the notion of transmission – possession is not enough.

I think that the objection is misdirected, as I shall now argue. Such gerrymandered aggregates do not qualify as causal processes according to the CQ theory, providing what counts as an object is adequately explicated.

According to the CQ theory, there are causal processes such as billiard balls rolling across tables, and pseudo processes such as shadows and spots of light. Is this exhaustive of all items that may be represented as occupying a spacetime region? The answer is no; there is also "spatiotemporal junk" – items that are not processes at all on the CQ definition. An example is what I earlier called "timewise gerrymanders." A timewise gerrymander is a putative object defined over a time interval where the definition changes over time (the putative object is really different objects at different times). A comparison may be drawn to Goodman's gerrymandered properties (grue, bleen) which are really different properties at different times (Goodman 1955). An example of a timewise gerrymander is the putative object x defined as:

for $t_1 \leq t < t_2$;	x is the coin in my pocket
for $t_2 \leq t < t_3$;	x is the red pen on my desk
for $t_3 \leq t < t_4$;	x is my watch.

Notice that x occupies a determinate spacetime region, and that at any time in the interval t_1 to t_4, x 'possesses' conserved quantities such as momentum (although not strictly speaking, for something must be an object in order to possess a conserved quantity). Clearly, there are innumerable such timewise gerrymanders.

Timewise gerrymanders are to be distinguished from spacewise gerrymanders. An example of a spacewise gerrymander is the putative object y consisting of my watch plus the red pen on my desk plus the coin in my pocket. The spacetime representation of x consists of three vertical lines that do not coexist at any time, whereas the spacetime representation of y consists of three vertical lines that coexist over the entire interval.

Timewise gerrymanders are sometimes defined just by a single formula. For example, in a billiards game, take x to be 'the closest ball to the black ball,' in the sense that, say;

for $t_1 \leq t < t_2$;	x is the pink ball
for $t_2 \leq t < t_3$;	x is the red ball
for $t_3 \leq t < t_4$;	x is the white ball.

Here x is a timewise gerrymander, occupying a spacetime region and 'possessing' conserved quantities. A similar case would be 'the president of the USSR' taken to refer to a single object consisting of the mereological sum of each of the presidents from Lenin to Gorbachev, each taken only for the time they were in office. Other single-formula timewise gerrymanders include 'the object currently in the centre of my field of vision' or 'the object nearest to my car,' provided these are taken in the gerrymandering sense and not in the usual sense as referring not to objects of which the description may be true at times other than the present, but to objects of which the description is currently, if not always, true.

Timewise gerrymanders sometimes display consistency of some feature. For example, in a box of molecules, take x to be whatever molecule has momentum p_x, taken just for the time that it has that momentum. Again, x is a timewise gerrymander, occupying a spacetime region and 'possessing' conserved quantities, but here x also has a *stable* momentum. The object has a consistency of some property over its entire history. Finally, timewise gerrymanders sometimes display spatiotemporal continuity. For example, consider a line of ten contiguous stationary billiard balls. Let x be the mereological sum of the first ball during the first time interval, plus the second during an immediately subsequent time interval, and so on for the ten balls. Then x is a timewise gerrymander, represented on a spacetime diagram by a diagonal line roughly one billiard ball wide. Note that x is not to be confused with the object y consisting of the whole line of ten balls, which *is* a genuine object, and which is represented by a vertical block ten balls wide.

Since the CQ theory has it that the world line of an object possessing a conserved quantity qualifies as a causal process, doesn't x qualify as a causal process? The answer is no, because x does not qualify as an object. As we have seen in section 5.1, there is implicit in the CQ theory a restriction on what counts as an object – it must display identity over time. A timewise gerrymander is a collection of different objects at different times. So, timewise gerrymanders are not objects.

So, turning to Salmon's example of the aggregate of the patches of wall sequentially illuminated, we can see that although it is generated by a single description, involves some uniformity, and displays spatiotemporal continuity, it nevertheless is a timewise gerrymander, and not in fact an object on my definition, since it does not display identity over time. The spot itself is an object (although not causal), and the entire patch of wall is an object (like the ten billiard balls), but the timewise gerrymander is not. It therefore is not a process of any sort, let alone a causal process; it qualifies on my account as spatiotemporal junk.

We need to be careful of an equivocation. The world line of 'the patch illuminated at time t_1' is a genuine process – it is an object that is temporally illuminated. However, the object made up of 'the patch of wall illuminated at time t_1' plus 'the patch of wall illuminated at time t_2' plus . . . is a timewise gerrymander.

Therefore we may conclude that Salmon's example, being a timewise gerrymander, does not prove that the conserved quantity definition is too inclusive, since such timewise gerrymanders do not qualify as objects. Therefore, the example does not force us to supplement the notion of possession of the relevant quantities. Possession may be only nine tenths of the law, but it is the full story on causal processes.

Max Kistler (1998: 16–17) argues that this account is circular in the sense that the requirement that an object display identity through time rules out timewise gerrymanders only if you already know that their temporal parts are not parts of the same object. This is true, but it misses the point of my analysis. I take it that it is intuitively clear that the temporal stages of certain timewise gerrymanders are not temporal parts of a single object. Once that is recognised, it becomes clear that Salmon's series of spots is also a timewise gerrymander. It's up to an account of identity through time to explain why the temporal stages of a timewise gerrymander are not parts of a single genuine object. We discuss such accounts in the next section, although it is not the burden of this book to give a theory of identity.

V.4 IDENTITY THROUGH TIME

The notion of a process, as explicated in the CQ theory, involves the idea of identity through time. A process is the world line of an object, so fundamentally, to constitute a process, an object must persist over time. This analysis presupposes a notion of identity through time; since