Forthcoming in *Real Metaphysics: Essays in Honour of D. H. Mellor*, edited by G. Rodriguez-Pereya and H. Lillehammer, Routledge, 2002.

IS CAUSATION A GENUINE RELATION?

Peter Menzies
Department of Philosophy
Macquarie University
North Ryde NSW 2109
Australia
Email: peter.menzies@mq.edu.au

1. INTRODUCTION

Over a period of more than thirty years Hugh Mellor's writings have illuminated an enormous range of metaphysical issues to do with chance (Mellor 1971), dispositions, laws, properties (Mellor 1991), and time (Mellor 1981). His work has had a salutary influence in encouraging metaphysicians to think about these issues in clear-headed, realist ways.

His work on the metaphysics of causation (Mellor 1995), in particular, is distinguished by its rigour, cogency and originality. The main outlines of his theory of causation are well known. He has argued that causation relates facts primarily, with causation between events deriving from causation between facts; that causation comes in deterministic and probabilistic varieties; that each variety can be explained in terms of closest-world counterfactuals in which single-case chances play a crucial role; that the important connotations of causation are that causes precede, are contiguous with, are evidence for, explain, and are means for bringing about effects; and that these connotations are consistent with, or imply, that causes increase the chances of their effects.

There is much in Mellor's theory of causation that I find congenial. Indeed, I hold many of the same views precisely because he has persuaded me of their truth. But we disagree on one issue that is central to the conceptual analysis and metaphysics of causation. The issue concerns whether causation is a genuine relation. Mellor believes that it is not, whereas I believe that it is. In this paper I scrutinise his criticisms of the view that causation is a relation, and, in passing, consider some related arguments for the same sceptical conclusion advanced by other philosophers. My conclusion that Mellor's scepticism on this matter is misplaced is not too surprising, but some of the arguments I rely on to reach this conclusion highlight a surprising and hitherto-overlooked feature of the concept of causation.

2. CAUSATION AS AN INTRINSIC RELATION

Before detailing Mellor's criticisms of the conception of causation as a real relation, it is worthwhile considering what reasons there might be that favour this conception.

One reason I have given is that our intuitive judgements about cases of pre-emption and over-determination rely on the idea that causation is an intrinsic relation between a cause and its effect. (Menzies 1996) If causation is an intrinsic relation, it must, *a fortiori*, be a relation. Allow me to rehearse briefly the reasons for thinking that it is an intrinsic relation.

Pre-emption and over-determination examples are alike in that there are two or more processes leading to some effect. In pre-emption examples only one of the processes goes to completion and brings about the effect but in doing so it cuts off the other processes. In over-determination examples, all of the processes go to completion, with no pre-emption of one process by another. Consider one familiar kind of pre-emption example.

> Case 1: Assassin A and assassin B, who are both deadly accurate marksmen, have been hired to kill a prominent political figure. They work independently of each other. But, as it happens, they come across their victim at the same time and place. Both assassins take careful aim, their fingers poised to pull their triggers. But assassin A fires first, his bullet hitting its mark. On seeing the victim collapse, assassin B refrains from pulling his trigger. However, if assassin A had not fired, assassin B would certainly have fired and hit his mark.

This kind of case appears to pose a difficulty for any theory such as Mellor's that employs a counterfactual increase-in-chance condition as a criterion for causation. Introducing a notion of counterfactual dependence at this point will make our discussion more precise.

> (1) e *counterfactually depends* on c if and only e's chance of occurring if c had not occurred would be less than its actual chance given that c did occur.

As Mellor's theory analyses causation in terms of what I am calling counterfactual dependence, it encounters difficulties with pre-emption examples like Case 1. For it cannot explain our intuitive judgment that assassin A's action caused the victim's death, since A's action did not increase the chance of the victim's death. For if assassin A had not fired, the chance of the victim's dying would have been the same as its actual chance—that is, fairly close to 100%—given the presence of the very reliable assassin B waiting in the wings.

Mellor has not discussed this kind of pre-emption example in his published work, as far as I know. I suspect that the reason for this is that he believes that his theory has the resources to be able to deal with it.[1] He argues that there is no unmediated action at a distance (Mellor 1995: 229-34). In other words, a cause and effect must be linked by a chain of intermediate contiguous causes and effects. In this case it seems there is a ready solution to the pre-emption problem. There appears to

be a chain of contiguous causes and effects running from assassin A's actions to the victim's death, but no such chain running from assassin B's action. This seems to vindicate our intuitive judgements about the example.

However, this appearance of a ready solution dissolves on closer examination. It turns out that this kind of pre-emption example, which David Lewis (1986a) has described as *late pre-emption*, is more intractable than first appears.[2] The special problem posed by late pre-emption examples is that they make it hard to establish the existence of a chain of contiguous causes and effects running from the main pre-empting cause to the effect. Consider, for example, the chain of events running from assassin A's pulling the trigger and the victim's death. What is the last link in the chain of contiguous causes and effects? Could A's bullet travelling in mid-air towards the victim's body be the immediate cause of the victim's death? It would appear not, because this event did not increase the chance of the victim's death: even if there had been no bullet in mid-trajectory, the victim would have died anyway from a bullet fired a few seconds later by the back-up assassin B. The same kind of reasoning shows that no event contiguous with the effect satisfies the counterfactual dependence condition for being its immediate cause.

A common strategy for rescuing the counterfactual increase-in-chance condition from cases of late pre-emption is to insist on a very strict criterion of identity for the entities that serve as cause and effect. For example, if one insists that the victim's death could not occur at a different time from its actual time of occurrence, then one might be able to argue that A's bullet in mid-trajectory does in fact satisfy the counterfactual criterion: if A's bullet had not been in mid-trajectory, the victim would not have died at the time he did but a few seconds later, in which case he would have died a different death. Mellor is surprisingly reticent on the question of the criterion of identity for the facts that he takes to be linked by causation. He says in one place that Don's dying is 'his dying roughly then, there and as he does' (Mellor 1995:14). Depending on how rough 'roughly then' is, he might appeal to this strategy of taking facts to be very fragile in order to rescue his account from the problem of late pre-emption. But there is another class of problems from which his theory cannot be rescued so easily.

This class of problems concerns over-determination examples. For example, consider the following variant of Case 1.

> Case 2: Assassins A and B are deadly accurate marksmen, working independently of each other to kill a prominent political figure. They both come across their victim at the same time and place. Both fire bullets into their victim's heart at exactly the same moment.

Neither A's firing nor B's firing satisfies the increase-in-chance condition for being a cause of the victim's death. For if one of them had not fired, the victim would still have died, and indeed died at exactly the same time, from the other's bullet. Once

again the strategy of looking for a chain of contiguous causes and effects does not help due to the persistent problem of establishing the last link in the chain. One cannot say that either A's bullet or B's bullet in mid-trajectory is the immediate contiguous cause of the victim's death because the absence of one or other would leave undiminished the chance of the victim's dying.

Examples of pre-emption and over-determination such as these should, in my opinion, make us very sceptical about the prospects for an analysis of causation in terms of a counterfactual dependence condition. As an alternative to such an analysis, I proposed a functionalist theory of causation as a theoretical entity (Menzies 1996). Adopting a standard treatment of theoretical entities, I argued that the functional role of causation is given by certain crucial platitudes in the folk theory of causation. One crucial platitude is that causation is an intrinsic relation, which means, roughly, a relation determined by the intrinsic properties of its relata and of the process connecting them. In this respect, I think, the commonsense conception of causation simply conflicts with the Humean view of causation as an extrinsic relation depending on large-scale regularities.[3] Another crucial platitude of the folk theory is that the causal relation coincides for the most part with the counterfactual dependence condition, with the notable exceptions of pre-emption and over-determination cases. Consequently, even if the counterfactual dependence condition cannot define causation, it can at least serve as a defeasible marker for the presence of the intrinsic relation that is causation. Combining these crucial platitudes, I offered the following functionalist analysis of causation:

> (2) c is *a cause* of a distinct event e if and only if the intrinsic relation that typically accompanies a counterfactual dependence between events holds between c and e.

This definition offers an *a priori* conceptual analysis of causation in terms of a certain counterfactually specified functional role. As with similar functionalist definitions, it can lead to *a posteriori* identification of the actual occupant of the functional role. Assuming that causation could be defined as an absolute relation in this way, I suggested that the intrinsic relation that occupies the functional role of causation might be the relation of exerting a force, or the relation of transfer of energy or momentum. Given some such identification, it is easy to see how the *a priori* analysis, combined with the *a posteriori* identification, leads to a uniform solution to the problems arising from pre-emption and over-determination. For in each of the problem cases our judgement about which of the potential causes actually caused the effect tracks whether a complete process of a kind that could occupy the functional role defined above connects the potential cause with the effect.

The crucial feature of this analysis for my discussion here is the reference to causation as an intrinsic relation. In Menzies 1999 I explored several different ways in which the notion of an intrinsic relation might be explained, settling on one explanation in terms of a robustly realist conception of universals, or perfectly

natural properties and relations that carve nature at its joints. Assuming the existence of such properties and relations, I followed Lewis in defining an intrinsic relation as one instantiated by a pair of relata just in virtue of the perfectly natural properties and relations of that pair itself.[4] An intrinsic relation, so defined, supervenes on just the perfectly natural properties and relations of its relata. This definition certainly assumes a very robust conception of causation as a relation. It lays itself open, therefore, to Mellor's criticisms of this conception.

3. MELLOR'S CRITIQUE OF CAUSATION AS A RELATION

Before detailing these criticisms, let me describe some of Mellor's background views on causation (Mellor 1995: 156-62). He argues that the canonical form of causal statements is given by '<u>E</u> because <u>C</u>', where '<u>C</u>' and '<u>E</u>' state facts and 'because' is a sentential connective. Statements of this form certainly appear to state relations, in particular relations between facts. But he argues that this is so only on a broad sense of 'relation', according to which there is a relation corresponding to every relational predicate. It is not so on a narrower ontological sense, according to which relations are universals existing independently of thought and language. Moreover, it is only in a broad sense of 'facts', according to which facts correspond to true statements, that causal statements appear to relate facts. It is not so on a narrower ontological sense in which facts are the ontological grounds or ultimate truthmakers for statements. Mellor reserves the term 'facta' for the truthmakers of statements. In his discussion of whether causation is a real relation, Mellor is concerned with the question whether the facta that are the truthmakers for causal statements like '<u>E</u> because <u>C</u>' consist in a genuine relation between facta.

Mellor advances two arguments against the view that the truthmaker for a causal statement is a relation between facta. One argument is that even if facta exist to act as relata—which, as we shall see, cannot always be taken for granted—a causal statement need not be made true by the existence of a relation between such facta (Mellor 1995: 162-5). He gives an example to illustrate this point. In a golf game, Sue pulls her drive to the left, making her ball bounce off a tree and, by a fluke, giving her a hole-in-one. Mellor's favoured description of the causal relations in this example is that Sue holed out in one *because* she drove her ball but *despite* the fact that she pulled her drive to the left. This accords with his counterfactual dependence criterion of causation. For Sue's driving the ball made it more likely that the ball would fall into the cup for a hole-in-one, while her pulling her drive to the left made it less likely. However, it is not plausible, he argues, to say the truthmaker for the positive causal statement consists in a relation between two facta. Even if we suppose that her holing out in one is a factum, it is not plausible to suppose that the cause of this, her driving the ball, is a factum as well. This is so because there is another fact that entails, but is not entailed by, this fact, which must therefore be a factum, namely her pulling her drive to the left. But the relationship between these facta, Sue's pulling her drive to the left and her holing out in one, is not causal because the first does not increase the chance of the other.

For two reasons I find this argument to be the less compelling of Mellor's arguments. First, it depends on Mellor's very contentious claim that it was Sue's driving the ball, but not her pulling her drive to the left, that caused the ball to fall into the cup for a hole-in-one. As he notes, this example is a variant of some much discussed problem cases initially cited to show that a cause need not increase the chance of its effect. (See, for example, Salmon 1884: 192-202) Those who have advanced these examples would insist—correctly in my view—that Sue holed out in one because she pulled her drive to the left. After all, the ball fell into the cup because it hit the tree and it hit the tree because Sue pulled her drive to the left. Of course, Mellor has independent reasons, stemming from his adherence to the counterfactual dependence criterion of causation, for thinking that his is the right description of the causal facts of the situation. But those of us who reject this criterion will find his description of the causal facts far from compulsory.

My second reason for finding the argument less persuasive is that it depends on a dubious principle about truthmaking facta. More precisely, it depends on the principle that if some fact <u>P</u> is entailed by, but does not entail, some other fact <u>Q</u>, then <u>P</u> cannot be a genuine factum. Is this principle at all plausible? In the particular example, Sue's driving is said to be entailed by, but not to entail, her pulling her drive to the left, which is then assumed to be a factum. But the very same principle that entails that Sue's driving is not a factum would surely imply that her pulling her drive to the left is not a factum either. For there is a fact that entails, but is not entailed by, this fact: viz that she pulls her drive to the left with a minute twist of her wrist. Indeed, the same style of argument would show that this latter fact cannot be a factum either because a still more precise description can be given of Sue's action. Indeed, since any positive event or action can be specified in more and more fine-grained ways, it would seem that Mellor's principle commits him to supposing that there are no facta at all, or ones that are maximally specific with respect to some particular sort of information. But without any account of the rules for determining such maximally specific sorts of information, the notion of a factum constrained by the above principle is useless in addressing the question of the truthmakers for ordinary causal statements. So, in future, I shall understand the notion of a factum in such a way that it need not conform to this principle.

Mellor's second argument against the view that causation is a relation between facta is more persuasive, I believe. David Lewis has called it the missing relatum objection (Lewis 1999). As Mellor formulates the objection, there are true causal statements involving negative occurrences such as:

(3) Kim has no children because she took contraceptives.
(4) Kim works full time because she has no children.

However, the absence that is said to be a cause in (4) and an effect in (3) is not a genuine factum, since there is nothing in the world to act as truthmaker for the

statement 'Kim has no children'. If the truthmaker for a true causal statement were always a relation between facta, the truth of causal statements such as (3) and (4) would require the existence of a factum corresponding to Kim's lack of children. But since such a factum does not exist, the truthmaker for a causal statement cannot be a genuine relation between facta.

Lewis also endorses this objection against the view that causation is a relation, though he thinks that if it were a relation its relata would be events rather than facts (Lewis 2002). Moreover, he has an additional argument that is targeted directly at the more specific view that causation consists in an intrinsic relation (Lewis 2001). His counterexample to this view involves a case of so-called *double prevention*: a cause prevents something which, had it not been prevented, would have prevented the effect.[5]

> Case 3: A collision between billiard balls 1 and 2 prevents ball 1 from continuing on its way and hitting ball 3. The collision of 1 and 3, had it occurred, would have prevented the subsequent collision of balls 3 and 4. But since in fact the collision of 1 and 3 was prevented, the collision of 3 and 4 was unprevented. Accordingly, the collision of 1 and 2 causes the collision of 3 and 4. Indeed, there is a matching counterfactual dependence: if there had been no collision between 1 and 2, there would have been no chance of a collision between 3 and 4 (Lewis 1999: 13)

There are two problems posed by the example, according to Lewis. First, the counterfactual dependence is an extrinsic matter. Had there been some other obstruction that would have stopped ball 1 from hitting ball 3, the collision of 3 and 4 would not have depended upon the collision of 1 and 2. Secondly, there is no continuous chain of events running from cause to effect. Between the collision of balls 1 and 2 and the collision of balls 3 and 4, nothing much happens. What matters here, Lewis argues, is not what happens, but what does not happen.

It is interesting to note here that such examples of double prevention may also pose a problem for Mellor's theory. Recall that Mellor's theory does not require that an indirect cause be connected to its effect by counterfactual dependence, but merely that it be connected by a chain of contiguous causes and effects, each linked by a counterfactual dependence. It is reasonable to interpret this requirement in the light of his views that facta must be causation's relata if it has any. When interpreted in this way and when applied to Lewis's example of double prevention, the requirement necessitates that there be a causal chain of contiguous facta running from the collision of balls 1 and 2 to the collision of balls 3 and 4. However, since nothing happens in the spatial region between these collisions, there are no facta to form this causal chain.

In summary, then, Mellor's missing relatum objection and Lewis's objection from double prevention seem to raise genuine difficulties for the view that causation

is an intrinsic relation. I am faced with a dilemma at this point. On the one hand, examples of pre-emption and over-determination highlight the plausibility of this view. On the other hand, the fact that causes and effects can be absences, which are not real things, seems to lead into the kinds of difficulties Mellor and Lewis raise. How is this dilemma to be resolved?

4. INTRINSIC RELATIONS RECONSIDERED

We can begin to resolve this dilemma, I believe, if we refocus our attention on the notion of an intrinsic relation. So far I have relied on Lewis's explication of intrinsic relations in terms of perfectly natural properties and relations. On this understanding, an intrinsic relation is one that holds just in virtue of the perfectly natural properties and relations holding of its relata. There is, however, an independent reason to be dissatisfied with this explication. The notion of intrinsicality that it explicates has its most natural application in the actual world in fundamental physics. Lewis claims that the perfectly natural properties and relations are coextensive, in the actual world at least, with the fundamental physical properties and relations. However, we need to be able to explain the notion of intrinsicality as it applies to causal processes studied outside fundamental physics. For example, the commonsense causal claim that the terrorist attacks in the United States in September 2001 caused an immediate dramatic fall on Wall Street is made true, I claim, by an intrinsic process. We are justified in believing in the existence of such a process and seeing it as the truthmaker for this causal claim without having the faintest idea how it might be analysed in terms of the properties and relations of fundamental physics.

It seems best, then, to start from scratch to explicate the notion of intrinsicality in such a way that it applies smoothly to the causal processes studied in the higher-level sciences as well as those studied in physical science. We can make a start on such an explication by noticing an implicit relativity involved in our concept of intrinsicality, and indeed in the concept of causation that is to be explained in terms of it. This relativity reflects the fact that our causal thinking is steeped in abstraction. Within any spatio-temporal region there are many different levels of causation, and within each level many crosscutting and intersecting causal processes. To determine the structure of these processes, we are forced to focus selectively on some aspects of what is going on and to background others. The causal schemas by which we interpret the world are irremediably permeated by abstractions that enable this selective focusing. One form of abstraction involves the identification within a given spatio-temporal region of *a system of a certain kind*.

A particular system of a certain kind consists in a set of constituent objects configured in specific ways. Clearly, the kinds of systems investigated by astronomers and cosmologists are different from the systems investigated by biologists and economists: solar systems and galaxies involve different kinds of constituent objects from economies, markets, species, and populations. However, a

system is not just a set of objects, but a set of objects that have certain properties and relations. And not any old properties and relations are relevant to the identification of a system as being of a certain kind. For example, a set of astronomical bodies can be individuated as a kind of planetary system by way of each body's relation to other bodies in the system, but not by way of their relations to objects outside the system. In short, a system of a given kind is a set of constituent objects internally organised in a distinctive fashion. The properties and relations that configure the objects into a system must be intrinsic to that kind of system.

The concept of intrinsicality at issue here is not the concept of properties and relations intrinsic *tout court*, but those intrinsic to a kind of system. It will suffice for our purposes to explain the intuitive idea behind this concept, rather than to present a full analysis, which is beyond the scope if this paper. Modifying an idea of Jaegwon Kim's (1982) concerning the simple concepts, I shall say the following:

> (5) A property F is *extrinsic to a system of kind K* if and only if, necessarily, a member of a set of objects constituting a system of kind K has F only if some contingent object wholly distinct from the set exists.

For example, the extrinsic properties of an astronomical body that is part of a distant planetary system might include being observed by some human on Earth.

The concept of a property intrinsic to a kind of system is defined in converse fashion:

> (6) A property F is *intrinsic to a system of kind K* if and only if, possibly, a member of the set of objects constituting a system of kind K has F although no contingent object wholly distinct from the set exists.

For example, the intrinsic properties of a planetary system would include the mass and shape of the individual astronomical bodies. But the intrinsic properties of the system need not all be intrinsic properties *simpliciter*. For example, the property of being gravitationally attracted to another member of the planetary system is an intrinsic property of the system, though it is not an intrinsic property *simpliciter*.[6] Notice that this definition does not prohibit negative, conjunctive, or even disjunctive properties from being intrinsic properties of a system.

There is a vast multitude of kinds of systems, but very few are of real interest to us. For the most part, we are interested only in the kinds of systems that evolve in lawful ways. As examples of these kinds of systems, we need only consider the kinds of systems investigated in scientific theories. Typically speaking, a scientific theory provides an abstract description of a certain kind of system in terms of a select set of state variables, and explains the behaviour of systems of the kind in question by showing how these variables change over time in conformity with certain laws. For example, classical mechanics employs the state variables of mass,

position, and momentum; and explains the motion of mechanical bodies, described in terms of these variables, by way of the Newtonian laws. Invariably, the state variables that a theory employs are intrinsic properties and relations of the systems under consideration. To summarise:

> (7) *A lawful kind of system* is a kind of system whose intrinsic properties and relations (state variables) evolve over time in conformity with a common set of laws.

In general, a lawful kind of system supervenes on a set of intrinsic properties and relations that conform to a common set of laws. In other words, any two particular systems with the same intrinsic properties and relations conforming to the same laws must both belong, or fail to belong, to a given lawful kind of system.

In terms of the concepts at hand, we are in a position to explain the notion of an intrinsic process that is going to play a central role in the modified functionalist analysis of causation. I suggest the following definition.

> (8) *An intrinsic process* holding in a lawful kind of system is a temporally ordered sequence of states that instantiate the intrinsic properties and relations that constitute that kind of system.

For example, the intrinsic process that I suppose is the truthmaker for the commonsense causal claim 'The terrorist attack in the United States caused a dramatic fall on Wall Street' might consist in some sequence of states such as: the terrorist attack on United States facilities, vast economic losses to major companies, a loss of confidence among major investors, a delay in the reopening of Wall Street stock market and widespread panic among traders at the reopening. Here I assume that the kind of system that is implicitly being considered in this commonsense causal claim can be specified rather loosely as that of an open market economy. Whichever way the kind of system is precisely specified, it is clear that intrinsic processes of this kind are not identified in terms of the properties and relations of fundamental physics.

Finally, we are in a position to consider how the functionalist analysis of causation should be modified to accommodate this relativised understanding of an intrinsic process. First, we would expect that there should be a matching relativisation in the causal concept. Elsewhere (Menzies 2002) I have argued that the causal concept must be understood as relativised to the contextual parameter of a lawful kind of system. We shall consider some evidence in support of this context-relativity in the next section. Secondly, we would expect that the analysis of causation should encompass all causal statements, whether they concern positive or negative occurrences. In order to be as neutral as possible over the contentious issue of the nature of the causal relata, I shall simply talk of them as property-instances.[7] Whichever way such property instances are to be understood, they are to include

instances of negative as well as positive properties. With these two preliminary remarks, let me state the modified functionalist analysis:

> (9) If c is an instance of property F and e an instance of property G in a lawful system of kind K, then c and e are *causally related* if and only if (i) there is a kind of intrinsic process that typically holds in systems of kind K when a G-instance is counterfactually dependent on an F-instance; and (ii) a process of this kind holds in the particular system of kind K that includes c and e.

This analysis is meant to apply uniformly to all causal statements, whether they concern positive or negative occurrences.

This analysis works as well as the old one when it comes to explaining our intuitions about pre-emption and over-determination. For example, suppose that we see Case 1 as exemplifying the kind of system that consists of a sole assassin shooting with a rifle at an unprotected person. There are clearly counterfactual dependences between shootings and deaths in this kind of system; and furthermore there is an obvious kind of intrinsic process that typically accompanies these dependences. This kind of process involves an assassin pulling the trigger of his rifle, a bullet being released by the rifle, the bullet travelling through the air and hitting the body of the person, followed by the person's death. In Cases 1 and 2 we look for a process of this kind to discriminate the actual from the potential causes.

But it is important to recognise that the analysis applies just as readily to the examples in which the cause and effect are absences, that is, instances of negative properties. Take Mellor's example in which Kim's use of the contraceptive pill causes her not to have children. If we suppose that the relevant kind of system involved in this causal statement is that of human female's body functioning according to the laws of human anatomy, then we will find counterfactual dependences holding in systems of this kind between the use of contraceptives and the absence of children. Moreover, we will find that there is a kind of intrinsic process that typically accompanies such counterfactual dependences, a process consisting of ingestion of oestregen, disruption of ovulation, absence of fertilisation, and absence of foetus formation. (Note that all these properties count as intrinsic properties of this kind of system.) Moreover, if this kind of intrinsic process obtains in Kim's case, then there is a truthmaker for the claim that Kim does not have children because she uses the contraceptive pill.

Similarly, the analysis applies straightforwardly to the example of double prevention that Lewis discusses. Let us suppose that the kind of system involved in the relevant causal claim is one consisting of four billiard balls with momenta of the same magnitude and direction as those in the original example. Once more, we can expect to find a kind of intrinsic process that typically accompanies a counterfactual dependence between collisions in this kind of system. The process will consist of balls 1 and 2 colliding, with each moving in a different direction from its initial

direction, and then balls 3 and 4 colliding without interference from the other balls. Even though there may be a spatial gap between the collisions, this process is nonetheless temporally continuous. Since such a process obtains in the actual situation under consideration, there is a truthmaker for the causal claim that the collision of balls 1 and 2 caused the collision of balls 3 and 4.

Accordingly, I suggest that causal situations involving absences or double prevention present no difficulty for the modified account. This account allows me to hold on to the view that causal statements are made true by intrinsic processes, while maintaining that there is but one causal concept whose analysis applies to all causal statements, regardless of whether they relate positive or negative occurrences.

Of course, I can maintain both views because I now construe the notion of an intrinsic process in a broad and flexible way. It might be thought that the notion of an intrinsic process has been made so broad and flexible as to be theoretically useless. But this is not so, I would argue. The notion of an intrinsic process still plays an essential role in explaining our intuitions about causation in cases of pre-emption. Consider the following modification of Case 3, Lewis's example of double prevention:

> Case 4. The set-up is the same as in case 3, but there is a billiard ball 5 that is on a collision course with ball 1. If ball 2 had not first collided with ball 1, then ball 5 would have a bit later on, so that one or other collision would have prevented ball 1 from colliding with ball 3. So ball 5 is back-up preventer of the collision between balls 1 and 3, which would have prevented the collision of balls 3 and 4.

This example of so-called *pre-emptive prevention* is an interesting test case for the modified functionalist analysis.[8] The analysis should be able to discriminate the actual from the potential preventer of the collision between balls 1 and 3. Notice that a pure counterfactual analysis cannot do this. For the absence of a collision between balls 1 and 3 does not depend counterfactually either on the motion of ball 2 or on the motion of ball 5: if one of these events had occurred without the other, the collision between balls 1 and 3 would still have been prevented. But this example is readily handled by the modified functionalist analysis. If we consider the situation under consideration as an instance of the lawful kind of system in which one billiard ball collision prevents a later one, we can see that the kind of intrinsic process that underlies this prevention obtains in this particular case. An essential part of such an intrinsic process is a collision of two balls that disrupts one of them from its collision course with a third ball. Clearly, we can see that the motion of ball 1, but not the motion of ball 5, initiates a process of this kind; and so we are able to discriminate the actual from the potential preventer in this example.

5. THE RELATIVITY TO A KIND OF SYSTEM

It might be thought that a defect of the present account of causation is that it makes the concept of causation context-sensitive by making it relative to a lawful kind of system. However, I think, contrary to this line of thought, that this apparent weakness is one of the great strengths of the analysis. Our concept of causation is marked by a certain degree of indeterminacy and vagueness: about certain kinds of situation we display ambivalence in our causal judgements. By understanding the causal concept as involving a contextual parameter that can be set in various ways in different contexts, one can explain this indeterminacy. Let me illustrate this with just one type of example that is germane to our discussion here.

Cases of pre-emptive prevention have been much discussed of late as interesting test cases for theories of causation. It appears that an indeterminacy affects our judgment about them. Consider the following example that Michael McDermott has described.

> Case 5: You reach out and catch a passing cricket ball. The next thing along in the ball's direction of motion was a solid brick wall. Beyond that was a window. Did your action prevent the ball hitting the window? (Did it cause the ball to *not* hit the window?) (McDermott 1995: p.525)

People express conflicting intuitions about this example. When it is pointed out that the presence of the brick wall means that the window was never in any danger of being broken, people are inclined to say that your catch did not prevent the ball hitting the window. On the other hand, when it is pointed out that something must have prevented the ball hitting the window, they agree that it must have been your catch that did the preventing.

All of this makes sense in terms of the theory of causation presented above. There are different ways of modelling the causal structure of the situation depending on which kind of system one sees it as instantiating. Suppose one thinks of the relevant kind of system as one that includes you, the ball, the window and the brick wall with their given spatio-temporal arrangements. There is no counterfactual dependence between your catch and the ball's not hitting the window in this kind of system and so *a fortiori* no intrinsic process accompanying such a dependence. On the other hand, suppose one thinks of the situation as instantiating the kind of system that abstracts away from the presence of the brick wall — a kind of system that includes you, the ball and the window but excludes the brick wall as an object extrinsic to the system. Then there are counterfactual dependences between your catch and the ball's not hitting the window, and indeed these dependences pick out an intrinsic process of a certain kind. Moreover, a process of this kind holds in the particular situation under consideration, so supporting the judgement that your catch prevented the ball from hitting the window. In this way, the indeterminacy in our causal judgements can be traced to the multiple ways in which the contextual parameter of a kind of system can be fixed.

This account predicts that our readiness to accept the causal judgement that your catch prevented the ball's hitting the window goes hand-in-hand with our readiness to see the situation in terms of a kind of system that abstracts away from the presence of the brick wall. In McDermott's example our readiness to do this wavers somewhat. But now consider a modification of the example introduced by John Collins:

> Case 6: You reach out and catch a passing cricket ball. The next thing along in the ball's direction of motion was my hand. (I leapt up to catch the ball, but because of your faster reaction you caught the ball just in front of the point at which my hand was raised.) Beyond our outstretched hands is a window. Did your action prevent the ball hitting the window? (Collins 2001: 223)

Collins detects some indeterminacy in our causal judgement about whether your catch prevented the window from being broken. But he claims, correctly I think, that we are more inclined to accept it in this example than in McDermott's original example. He explains this in terms of how far-fetched it is to entertain the absence of the back-up preventer. It is easy to entertain the absence of my hand ready to take the catch: one simply imagines that I get my timing wrong so that when I leap I do so not at the right moment to be ready to take the catch. It is more far-fetched, on the other hand, to suppose that the brick wall is absent or that the ball would miraculously pass straight through it (Collins 2001: 227-9).

I think that Collins' explanation is on the right track to the extent that it links our willingness to accept the judgement that your catch prevented the ball's hitting the window with our willingness to abstract away from the presence of the back-up preventer. I would go further and explain this linkage in terms of the way we model the causal structure of a given situation in terms of kinds of systems that abstract away from the presence of factors that are viewed as extrinsic to the system. Again, I think there is something to Collins' explanation of our varying degrees of willingess to do this in terms of how far-fetched it is to imagine the absence of the back-up preventer. But I would prefer to see the matter of how far-fetched it is to imagine such things as rooted in fairly objective issues about the features of the situation itself: How permanent a feature of the set-up is the back-up preventer? Is it something that is an external intrusion in an otherwise isolated system? Would the system that abstracts away from its presence fall under wider, more robust laws than the system that retains its presence? Such considerations can yield fairly objective reasons for modelling a situation in terms of one kind of system rather than another.

However, it is crucial in such discussions to keep in mind the implicit relativity of causal judgements to a lawful kind of system. Overlooking this context relativity makes one more liable to fall into conceptual traps. As an illustration of this, consider an argument of Collins' to the effect that examples of pre-emptive prevention falsify any theory that takes causation to consist in an intrinsic process. He observes that your catch prevents the window from breaking when it pre-empts

my catch from preventing the window from breaking, but not when it pre-empts the brick wall from doing so. Yet the only difference between these cases, he says, has to do with features extrinsic to the simple process involving your catch. 'The process that includes the ball's flight, your catch, and the window's not breaking is causal in the case where my hand was poised behind yours to take the catch, but it is *not* causal in the case where a brick wall is there instead of me.' (Collins 2001: 226)

A sufficient counter to this argument starts from the observation that the concept of an intrinsic process, like that of the causal concept, must be seen to be relative to a lawful kind of system. Indeed, as we have seen, intrinsic processes are often widespread features of entire systems, rather than localised parts of the system. Now notice that the two causal judgements that Collins' argument turns on involve quite different kinds of systems. The judgement about the modified example that your catch prevented the ball hitting the window involves the kind of system that *excludes* the back-up preventer of my outstretched hand. The opposite judgement about the original example involves a kind of system that *includes* the back-up preventer of the wall. In order to compare the intrinsic processes that could act as truthmakers for these judgements, we have to consider sequences of states involving all the intrinsic features of these systems. In the first system the intrinsic process will consist of a sequence of states holding true of you, the ball, and the window, whereas in the second system the intrinsic process will consist of a sequence of states holding true of you, the ball, the brick wall and the window. The presence of the brick wall in one system but not the other makes a big difference as to what counts as the intrinsic features of the set-ups. It is false to say, therefore, that the two set-ups agree in intrinsic processes and only differ in matters extrinsic to these processes, namely the presence of the back-up preventer. The difference between the set-ups with respect to the presence of the brick wall makes for a difference in intrinsic processes, a difference that ultimately underlies our readiness to accept a causal judgement about one set-up but not the other.

[1] See Noordhof 1998 for a good discussion of this point.

[2] In *early* pre-emption examples, *the main process* that goes through to completion and brings about the effect cuts short all alternative processes before the effect has occurred, whereas in *late* pre-emption examples the main process goes through to completion, but it is *the effect* itself that cuts short all the other alternative processes after it has occurred.

[3] For further discussion of the conflict between the commonsense conception of causation as an intrinsic relation and the standard Humean position see Menzies 1998.

[4] Lewis's definition of an intrinsic relation appears in Lewis 1983 and 1986. He actually defines two kinds of intrinsic relations: relations intrinsic *to their relata* and relations intrinsic *to their pairs*. The relevant kind of intrinsic relations I consider in connection with causation correspond to relations intrinsic to their pairs.

[5] Similar cases of double prevention are discussed in McDermott 1995 and Hall 1994. Hall explicitly draws out the implications of such cases for the supposed intrinsic character of causation.

[6] A problem infects these definitions parallel to the problem Lewis (1983b) pointed out for Kim's definition of the simple concepts. Modifying some concepts introduced by Lewis, let us say that a system is *accompanied* if and only it coexists with some contingent object wholly distinct from it, and *lonely* if and only if does not so coexist. The definitions I have presented amount to saying that that the extrinsic properties of a system are those implied by the accompaniment of the system and the intrinsic properties of a system are those compatible with its loneliness. The problem is that loneliness of a system is intuitively an extrinsic property of the system (since it can differ between duplicates of the system), but it counts as an intrinsic property by the definition (since it is compatible with itself). One possible remedy to this problem may be to adapt to our purposes the refinement of Kim's original idea to be found in Langton and Lewis 1999. This refinement is supposed to circumvent the defect Lewis detected in Kim's original idea.

[7] For a more detailed discussion of the nature of the causal relata see Menzies 1989, in which I argue that fact-like entities, which I call situations, are the primary relata of causation.

[8] Examples of pre-emptive prevention have been discussed in McDermott 1995, Hall 1994, Lewis 1999, and Collins 2001.