Dispositions, manifestations, and causal structure

Toby Handfield

15 June 2009

1 Introduction Typically – perhaps always – the manifestation of a disposition involves a causal process. That thought is fairly unexceptional. More interesting is the further thought that the process must be of a certain type for it to be a manifestation of a particular disposition. When salt manifests its disposition to dissolve, for instance, it seems not merely to involve an initial state in which the salt is solid, and a final state in which it is dissolved. Rather, it is important that salt undergo a distinctive kind of causal process. If the salt solution came about by a deviant process, we might say the salt was not manifesting its disposition to dissolve. For example, if we had a host of tiny nano-machines that pulled the individual ions out of the crystalline structure of sodium chloride, and then built new water molecules around those ions, the deviancy of this process makes it look like the manifestation of a rather different disposition.[1]

If this claim about the link between the manifestation of dispositions and particular kinds of causal processes is correct, it suggests that causal processes form *natural kinds*.

Natural kinds are typically thought to (*i*) have real essences, and (*ii*) to facilitate scientific investigation – prediction and inference go better when it is prediction and inference about a kind. For my purposes in this paper, I will be exclusively interested in the first claim about kinds: that they have real essences. This is because it is only the first idea about natural kinds that is directly relevant to my broader metaphysical project. In the next section, I will describe that project and how it motivates inquiry into the possibility of natural kinds of causal processes – but it should be stressed that even if you are not interested in pursuing my particular metaphysical project, there might still be good reason to find it of interest whether causal processes can be members of natural kinds.

To appear in the volume *The Metaphysics of Powers*, edited by Anna Marmodoro (London: Routledge). This is an author's pre-publication draft. For citation purposes, please consult the published version.

1. Though, depending upon the details of the case, we might say that the nano-machines are merely catalysing the process by which the salt manifests its ordinary disposition to dissolve.

Ultimately, I will argue that there are grounds for optimism: that there may be natural kinds of causal processes, and that these natural kinds might be used to explain the relation between a disposition and its possible manifestation. However, this optimism will be tempered with caution. The sorts of causal process kinds that would be compatible with current science appear not to have all the features we might have expected, a priori.

2 Motivation: Humean dispositionalism Why am I interested in looking for natural kinds of causal processes? Principally, because I am unimpressed by neo-Humean ontology, and I think that an ontology of causal powers is more attractive in some ways.[2] However, I think that such an ontology of causal powers could turn out to be more congenial to Humeans than most have thought. A crucial element in demonstrating this congeniality is to establish that there there are natural kinds of causal processes, in terms of which causal powers might be *explained*.

Let me explain this idea a little further.

I understand there to be two especially important parts of the neo-Humean position. The first is the idea that there are no necessary connections between distinct existences – in Lewis's framework, this is the principle of *Recombination*. The second is the idea that everything is loose and separate – in other words, that there are *lots* of distinct existences. In Lewis's framework, this is the thesis of *Humean supervenience*: the idea that all the ontologically fundamental properties are properties of spatiotemporal points and the only ontologically fundamental relations are spatiotemporal. All else supervenes on this 'Humean mosaic'.[3]

Humean supervenience is almost certainly false. In light of quantum entanglement, and possibly also in light of fundamentally stochastic processes, it appears that the world is not anywhere near as loose and separate as Humean supervenience requires.[4] But to concede that Humean supervenience is false in no way impugns the truth of Recombination. Indeed, Recombination appears a deeply attractive principle. It might even be thought that Recombination is an analytic truth; definitive of the very meaning of the term "distinct existence".[5] Whatever the case, I think it is hard to provide clear empirical evidence that this principle is false. Any time someone attempts to demonstrate a necessary connection between distinct things, it seems open to us to reinterpret the evidence as a scenario in which the putatively distinct things are not really distinct: the world is less loose and separate than you thought.

2. For a paradigm example of neo-Humean ontology, see the works of David Lewis, perhaps best represented in Lewis 1999.
3. See the Introduction to Lewis 1986 for Lewis's principal statement of this idea.
4. For a recent and forceful expression of this point, see Maudlin 2007: 53–61. See also Oppy 2000.
5. Or perhaps not. For further discussion of plausible formulations of this idea, see Wilson Forthcoming.

So, if you are a naturalistically minded metaphysician, and you want to propose a substantive account of causation, laws of nature, and dispositions, what is one to do? I don't think you should reject both the Humean doctrines simultaneously. We should try a more cautious approach, attempting to retain Recombination, while ditching the far less important – and most likely false – Humean supervenience.[6]

Retaining Recombination, then, we ask: What account can we give of the apparent necessary connections in nature – the connections between cause and effect, between laws and the phenomena they appear to govern, and between disposition, stimulus, and manifestation? These are, after all, apparently necessary connections between distinct existences.

The neo-Humean response is to deny that the connections are necessary. Rather, they are widespread connections; they are connections which it is particularly useful to know about. But for all that, there is no extra metaphysical glue holding these sorts of events together.

In particular, a neo-Humean explains the fact that a particular property confers a particular causal power by pointing to the global pattern of property-instantiations that the property is involved in (the Humean mosaic). If *being solid sodium chloride* is instantiated in a global pattern such that most of its instances, when coinstantiated with *is wet* are followed by instances of *being dissolved sodium chloride*, then this very pattern of property-instances is what explains the power to dissolve that is conferred by being sodium chloride. This is the feature of neo-Humeanism that I am most determined to avoid. It seems to get the order of explanation entirely the wrong way around. It is – I maintain – because of the intrinsic nature of the property *being sodium chloride* that it confers the power to dissolve.

So can we – consistent with Recombination – explain causal powers as *intrinsic* to the natural properties: as things which are conferred by the intrinsic nature of the properties, as opposed to things which are conferred by the global pattern of property-instances?

Note that, by retaining Recombination, I am opposed to traditional dispositional essentialists, who say that the natural properties are essentially such as to confer certain powers on their bearers, *without further explanation*.[7] At least for deterministic causal powers, this position involves a rejection of Recombination. For instance, suppose positive charge is essentially such as to confer the power to attract negatively charged things.

6. Lewis himself only wished to defend the a priori tenability of Humean supervenience (1986). And in that enterprise, it comes down to a show of confidence in one's philosophical intuitions about the above concepts. If one is prepared to bite a few bullets, Humean Supervenience can come out looking quite tenable indeed – perhaps we can say it is a priori surprising, but not a priori untenable (Oppy 2000: 96).

7. I consider this to be a feature of the views of Ellis, Bird, and Mumford, for instance – though there are significant differences in the precise nature of the ontologies these theorists advocate, they all appeal to brute modal facts of some variety.

Then two atoms, one of each charge, *must* be accompanied by a distinct existence – at least in time. Either they must be accompanied by a later change in velocity, or they must be accompanied by some sort of interfering factor that prevents the acceleration.

It is rather more difficult to show that dispositional essentialists must reject Recombination if they employ only probabilistic causal powers. If the stimulus condition for a probabilistic power is met, it can no longer be said that such a circumstance *necessitates* a distinct existence. But clearly there is a non-trivial constraint on what can happen, and this constraint is not obviously susceptible of further analysis. An alternative approach, which gives some *explanatory* account of such modal constraints, compatible with Recombination, is obviously preferable, in my view.

2.1 The Humean dispositionalist program Working with dispositions in particular, I think that the following idea holds out some hope of explaining at least *some* of the modal nature of dispositions:

1. Identify the manifestation of a power as a natural kind of causal process.
2. Investigate the essential structure of such processes.
3. Attempt to identify internal relations between dispositional properties and such process structures.

The sort of internal relation that I have in mind at stage (3) is like that between the property "being hydrogen" and the structural property "being a molecule of H_2O". There is an internal relation between these, such that necessarily, anything which instantiates the latter has two parts which instantiate the former.

Exactly how to account for this necessary connection, I won't try to say here. But it seems that this is the sort of connection that Humeans should not be chary of. Cowardly types will perhaps try to assimilate it to some sort of analytic truth, but I think we may happily admit it as an a posteriori necessity.

The sort of structure I have in mind at stage (2), though, is not the structure of a molecule – it is the structure of a causal process. So I am hoping for internal relations between a property like *is a process of salt dissolving in water* and *is sodium chloride*. An internal relation like that would go some way towards *explaining* why being sodium chloride confers the power to dissolve in water. It would, moreover, explain why being sodium chloride is *essentially* such as to confer that power. Given the nature of sodium chloride, and given the nature of the process-kind, they would not be the things that they are if they did not stand in the internal relation that explains the power.

Moreover, if we start to understand the structure of natural kinds of causal processes, I speculate that we may see why not every conceivable process is possible. For instance, there presumably is no natural kind of process whereby oil dissolves in water. It is because there is no such possible process that *being oil* does not confer upon its bearers the power to dissolve in water.

So in effect, this would amount to an explanation of what has been called the 'Meinongian' feature of causal powers or dispositions: that they seem to 'point' to their potential manifestations, even where these are mere possibilia.[8] The explanation is that this 'pointing' is no more sinister than the fact that the property *being hydrogen* is essentially such that it is compatible with the structure of water.

Another reason to think that this approach to an ontology of powers is attractive is because it appears to handle Neil Williams's "problem of fit" without requiring one to commit to a holistic account of powers.[9]

The problem of fit is to explain how it is that the many and diverse powers can be intrinsically powerful, and yet have mutually interrelated powers. A glass, for example, has a power to shatter when hit by a rock; and a rock has a complementary power to shatter the glass. If these properties are *intrinsically* such as to confer these powers, then it seems a remarkable fluke that they happen to confer these mutually complementary powers. But by the same token, it seems metaphysically impossible that they could fail to confer mutually complementary powers.

If there are natural kinds of causal process however, which can be identified with power-manifestations, then it is possible to give a straightforward explanation of this. The explanation for why the rock and the glass have these two mutually complementary powers is that there is *one* manifestation kind – the shattering of a rock by a glass – in which both *being a rock* and *being a glass* are constituent properties.[10]

That, at any rate, is the program. I have called it "Humean dispositionalism".[11] Note that it by no means promises to explain *all* of the modal phenomena associated with dispositional properties.[12] For instance, it does nothing to explain why some manifestations might be *more probable* than others. All it appears to do is to explain why some manifestations are possible and others are not.

8. Armstrong 1997: 79.

9. Williams, this volume. Moreover, the concern that a powers-ontology is afflicted by some sort of vicious regress seems to be motivated by holism (see, e.g. the contributions of Lowe, Marmodoro, and McKitrick to this volume). So by avoiding holism, my preferred approach promises to avoid any such regress.

10. This is not to say that there can be *only* one manifestation kind in which these properties are constituents. Being a glass confers the power to emit a 'ping' noise when struck gently by a rock. This power, and the complementary power of the rock, are explained by the existence of a different manifestation kind.

11. Handfield 2008a.

12. So my program is less ambitious than Brian Ellis's 'Scientific Essentialism' (2001), which tries to explain all natural necessity in terms of essences of natural kinds. That said, Ellis's ideas are obviously closely related to mine, and I was moved to think more about the structure of causal processes by his paper with Caroline Lierse (Ellis and Lierse 1994), where he explicitly claims that there is a necessary connection between a disposition and the causal process in which it is manifested.

Still, I believe it is a promising start. And with that project as background, in this paper I examine the key empirical claim that the project needs to sustain in order to be viable: the claim that manifestations of dispositions are causal processes that form natural kinds.

3 Natural kinds with causal structure

3.1 Visualising causal structure Suppose that manifestations of a disposition indeed form natural kinds. As I indicated above, the crucial feature of natural kinds that is of interest for my purposes is that they possess some sort of essential structure. For natural kinds of *processes*, such structure is presumably, in some sense, *causal structure*. It is something to do with the causal roles of the various factors in play, and how they interact.

In order to work out whether or not there exist causal processes with such causal structure, I could simply enumerate some examples of causal processes of manifestation, and attempt to determine by armchair reflection whether or not there is any essential causal structure in such processes. But this is not very easy, because it remains quite unclear what such causal structure would consist in.

I suggest that part of our difficulty in clarifying our ideas about causal structure is that we find it hard to *visualise* processes. We don't have any obvious mental picture that goes with the idea of a process, in the way that a Lewis diagram is such a convenient mental image when thinking about the structure of a chemical compound. That said, we do use visual representations – in both science and philosophy – to depict processes. This suggests an indirect way of determining whether or not the desired causal structure exists in the world. We can ask: do any of these representations capture the sort of structure that could plausibly be thought to be the essential structure of a natural kind of process?

Accordingly, in the remainder of this paper I shall consider a variety of ways we actually represent causal processes, and discuss their prospects for representing the "causal structure" – if there is such a thing – of those processes. By seeing what sorts of representations are suited to depicting causal structure, we shall get a better grasp on what such causal structure might be like, and whether such causal structure is likely to exist. Although the strategy is relatively indirect, I believe it will prove to have clear heuristic value. When I tell someone that the truth of an interesting philosophical claim hangs on the question of whether or not the manifestation of a typical disposition has something like an essential causal structure, I have come to expect a blank look. But if I can say, pointing at a diagram, that the truth of the philosophical claim hangs on whether or not this sort of diagram is a good representation of typical disposition manifestations, then there is a much better prospect of understanding.

Employing this approach, I will argue that the prospects are promising for Humean dispositionalism. That is, there are methods of representing causal processes that seem to capture – or at least help us to develop a grasp of – what sort of structure might be essential to those process-kinds. Moreover, these types of representation are drawn directly from areas of empirical science which we have good reason to regard as roughly true. So we have some reason to be optimistic that the empirical nature of the world is at least compatible with the existence of natural kinds of causal processes.

3.2 Criteria for kindhood Before examining those methods of representation, I wish to prescribe some conditions on what the hoped-for kinds must be like. These criteria are not intended to be platitudes about natural kinds, but are crafted with an eye to the Humean dispositionalist project, of grounding facts about causal powers in internal relations between process kinds and natural properties.

First, the sorts of processes that get identified as members of these kinds must at least roughly match our intuitive ideas about what is involved in the manifestation of a disposition. Call this the *identity requirement*. If a candidate account of causal processes tells us that every time a match is ignited, it actually involves a process entirely located more than 100 years in the future, then this is a good sign that the account is inadequate. Similarly, if a candidate account entails wildly implausible identity conditions for the kinds of process, then this is reason to reject it. Whether or not Obama is president of the US at the time a match is lit, for instance, should not make any difference to what kind of lighting-process is instantiated.

Second (and related to the identity requirement), I suggest an *intrinsicness requirement*: natural kinds of causal process must be characterised exclusively in terms of intrinsic structure. This is not because I think it inconceivable for there to be natural kinds with essential extrinsic structure, but simply because I suspect it will prove to be conceptually convenient.[13] To illustrate: consider a process of some salt dissolving in water. This process has many extrinsic properties. Some of them seem obviously irrelevant to the causal process. That the dissolution is being watched by a class of high-school students, for instance, is surely not very interesting. There are other properties, however, which might appear extrinsic and yet also be relevant. Imagine that the dissolution is taking place in an extremely strong magnetic field, and further that this might make a difference to how rapidly the dissolution occurs. In such cases, it seems to me that we can accept the requirement that the crucial structure of the process is intrinsic – we merely need to liberalise our conception of what is *included* in the process. A region of the magnetic field, for instance, seems to be part of the process of dissolution, in such a case.

13. Thanks to Alexander Bird and Jennifer McKitrick for pressing me to clarify my justification on this point.

(Generally, this second requirement will mesh with the first, identity requirement. But it is possible that, in light of empirical science, attempting to meet this requirement will lead us to realise that many processes are much "bigger" than we thought they were. And that will pull *against* the identity requirement – that the processes identified roughly fit our intuitive ideas about process-identity.)

A third requirement – also derived from the first – is that it should be possible to identify the difference between a process of non-manifestation of a disposition as opposed to the *masking* of a disposition.[14] A rat is given some poison. Having ingested the poison, it is disposed to die in a couple of hours. But the rat then eats some rye-grass, that happens to act as an antidote to the poison. The rat comes through alive because the disposition to die was interfered with – or masked – by the antidote. This process is importantly different from the process by which the rat's sibling – which never consumed poison in the first place – survives over the same period of time. The sibling never had the same disposition to die in the relevant period of time. If a natural kinds account of causal processes is going to be plausible, then, it should enable us to distinguish these two types of scenario as involving distinct kinds of process.

This *masking requirement*, however, is defeasible. It may not be necessary, for all sorts of physical system, to distinguish masking cases from other cases. This might be because no sensible physical distinction between masking and non-masking can be drawn. As I will discuss below, there are indeed some physical systems which appear to be like this.

With these requirements in place, I now turn to examine some methods of depicting processes.

4 Neuron diagrams and causal models Neuron diagrams have been used for some time in philosophical discussion about causation. The conventions of a neuron diagram are appealingly straightforward:

— Circles ("neurons") represent possible events. When shaded, they represent events that occur. When empty, they represent events that do not occur.

— Arrows represent "stimulatory signals" between neurons. A stimulatory signal is sent if and only if the neuron at the base of the arrow is shaded.

— Arrows with blobs on the end represent "inhibitory signals" between neurons. Similarly to stimulatory signals, an inhibitory signal is sent if and only if the neuron at the base of the arrow is shaded.

— The temporal order of events is from left-to-right.

14. See Bird 1998; Choi 2003; Fara 2001; Johnston 1992 for discussion of masks and antidotes.

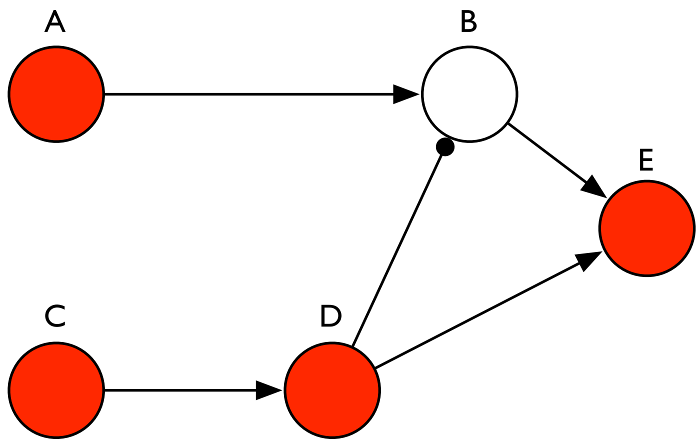

Figure 1: *This sort of neuron diagram represents a scenario where one cause pre-empts another potential cause. C, rather than A, is the cause of E. But had C not occurred, E still would have occurred.*

In a diagram such as Figure 1, it is fairly clear what causes what. It is partly for this reason that neuron diagrams are so useful in the now rather baroque literature attempting to provide an adequate analysis of token causation.

An alternative approach to representing causal phenomena that has found recent favour is to use causal models.[15] Causal models consist of a set of structural equations. Such equations encode in a fairly straightforward way the relevant counterfactual dependencies in play.

I will omit discussing the technical details of how causal models are used to analyse causation because, for my purposes, it suffices to note that they are similar to neuron diagrams: both methods of causal representation relatively directly encode information about counterfactual relations between possible events.

Do either of these approaches represent causal processes in a fashion compatible with the three constraints identified above? Apparently not: they fail the intrinsicness requirement. The reason for this is that both these methods of representing causal structure are concerned with relations of counterfactual dependence between "nodes" in the structure. And relations of counterfactual dependence are not intrinsic to their relata.[16] Consequently, the structure that is represented by neuron diagrams appears not to be intrinsic structure.

To illustrate, suppose that Figure 1 represents the notorious causal scenario by which Suzy throws a rock (*C*) which then strikes a bottle (*D*), causing it to break (*E*). Billy throws a rock (*A*) at the same time as Suzy, and his rock would have hit (*B*) if it hadn't been for

15. Pearl 2000. For a relatively accessible discussion of the philosophical application of causal models see Hitchcock 2001.

16. This point has been made by Lewis (2004) and also Peter Menzies (1999).

Suzy's throw – Suzy throws faster.

Now suppose we modify the scenario somewhat. All the same events happen, and all the *local*, intrinsic features of the scenario are the same.[17] Billy throws, Suzy throws. Suzy's rock hits. Billy's throw does not. Bottle smashes. What is *different* is that there is a robot watching Suzy's throw. If Suzy does not throw, the robot will fire a laser (*R*), destroying Billy's rock. So now, if Suzy does not throw, the bottle won't break. This necessitates a change in the diagram (see Figure 2). We need to insert additional neurons, to represent the robot, and the way in which the robot acts as a further inhibitor on the possible event of Billy's rock hitting the bottle (*B*). Even if we wanted to leave out the robot from the diagram, we would at least need to change the relations of counterfactual dependence between the neurons that we already have. The original diagram, as a representation of this new scenario, falsely implies that the bottle would have smashed if Suzy had not thrown. To rectify this, the arrow between *A* and *B*, in particular, should be removed or the description of *B* should otherwise be changed.

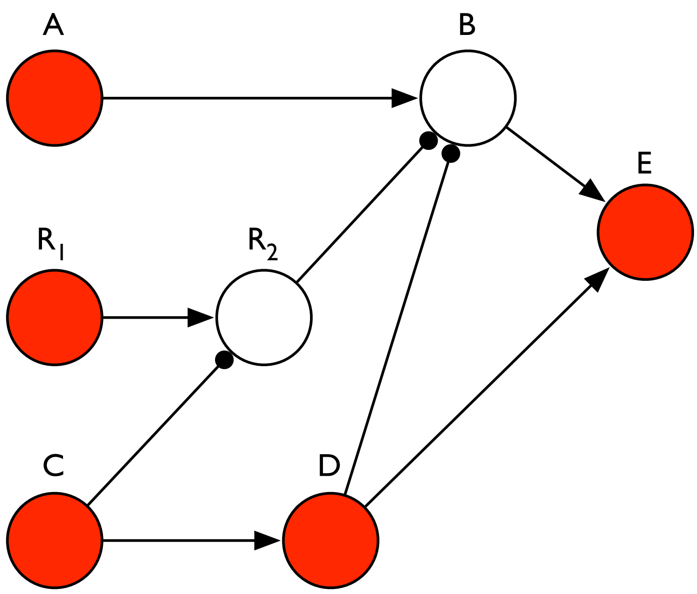

Figure 2: *The modified Billy-Suzy scenario, with a watchful robot, waiting to zap Billy's rock (R), in case Suzy does not throw.*

Notice, though, that this change in the causal structure of the scenario does not involve any change in anything like the "intrinsic structure" of the original set-up.

17. In this context, I am merely appealing to an intuitive sense of location: to a folk understanding of where the process is located. This is appropriate, given my stipulation of an intuitive *identity* requirement. I'll return to the idea of locality in Section 6.

In response, we could expand our understanding of the *location* of the process, so as to include the presence or absence of the robot explicitly. But the resulting conception of processes seems to get the identity conditions for processes wrong, thus failing the identity requirement. It is very strange to think that *whether or not the robot is present* changes the kind of process. The robot does not *do* anything! Admittedly, in a very limited way, the robot's presence makes a difference because it would exert minute gravitational and radiation effects on the flying rocks, thus changing their trajectories very slightly. But it is not for that sort of reason that neuron diagrams push us to pay heed to the robot. Rather, the robot is included because it poses a merely counterfactual threat to the process by which the bottle smashes. This is not the sort of factor that warrants a distinct process-kind for our purposes. Our aim is to identify process-kinds that are essential to the manifestation of dispositions. And although it makes some difference to precisely what occurs, the robot's presence apparently makes no difference to what bottle-smashing dispositions are manifested.

This sort of phenomenon – whereby a network of counterfactual dependence can be disrupted by adding in various "threats" or other potential interveners, is completely general, and in no way relies upon eccentricities of the Billy and Suzy story. So we can generally expect that a counterfactual structure approach will fail either the intrinsicness requirement or the identity requirement. Consequently, I suggest that we reject the neuron diagram and causal model approaches as ways to identify natural kinds of causal processes.

5 Space-time diagrams Another possibility is that causal structure can be represented by space-time diagrams. An obvious benefit to these diagrams, in contrast to neuron diagrams, is that they focus on what actually happens, rather than merely counterfactual possibilities of interaction.

Consider first, a very simple example of how one might identify causal structure in a space-time diagram depiction of a Newtonian world. Take the disposition of two masses to gravitate towards one another. As a space-time diagram, the manifestation of this disposition would look like Figure 3. The mark of causal *interaction*, in such a diagram, is the curvature of trajectories. Inertial motion – unaffected by any causal influence – is represented by a straight trajectory.

So, using examples like this, will it be possible to identify process-structure that is intrinsic, that at least roughly matches our intuitions about the identity conditions for processes, and that accounts for the distinction between masking and non-masking cases?

There are two principal difficulties with space-time diagrams for this purpose. First, a world-line reflects, in its curvature, the *net* forces acting upon the object. So how would one identify the structure of the process of the Earth revolving around the Sun, and distinguish that from the process of the moon revolving around the Earth, for instance?

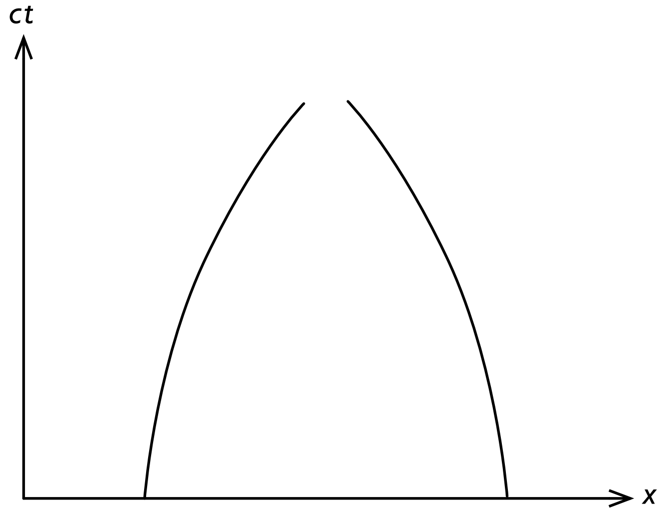

Figure 3: *A space-time diagram of two Newtonian objects accelerating toward each other.*

Perhaps this is a difficulty that we should simply live with, however. Both processes are happening at the same time, and a certain amount of idealisation would be necessary to represent the 'pure' version of each process. Given that the world involves a great deal of causal interaction, it might be naive to hope for a representation that has easily disentanglable components relating to the different causal processes. Moreover, there may be no physical basis for separating the world into any particular number of causal processes. So for now, I propose to gloss over this difficulty. I will return to discuss it in more detail in Section 8.

The second problem with spacetime diagrams, however, is that they will conflate the difference between various types of *masking* and cases of non-interaction. Consider Figure 4. This could be a case of two positively charged particles – *A* and *C* – that are disposed to repel, but this process is being *masked* by the negatively charged particle *B* in between. Or it could be a case of three uncharged particles that have no disposition to repel. A space-time diagram seems incapable of capturing this sort of difference – and this is a crucial difference if we are interested in the distinction between the masking of a disposition versus the sheer absence of the disposition.

For purely gravitational interactions, space-time diagrams do not have this difficulty, because gravitational attractions cannot be masked by other gravitational attractions. So it remains possible that space-time diagrams are appropriate ways of representing the causal structure of some sorts of interaction, but not all. (Indeed, it seems far from coincidental that we do not typically use space-time diagrams to represent electrodynamic phenomena – where masking is possible – but we do use them to represent gravitational

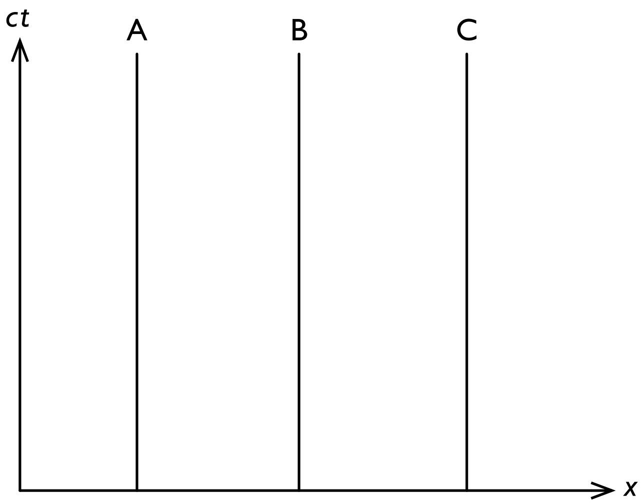

Figure 4: *Three bodies moving through spacetime without undergoing acceleration. Whether this is the result of masking or simply due to non-interaction is not resolved.*

phenomena.[18])

But this simplistic use of space-time diagrams suggests another, richer approach that shares some common roots: the idea that causal processes are space-time trajectories, characterised by the sorts of physical quantities that they instantiate.

6 Conserved quantity approaches Wesley Salmon (1984; 1998) and Phil Dowe (1992; 2000) have both made important contributions to the literature on causation by developing an idea that emerges naturally from the use of space-time diagrams: that causal processes are 'world-lines'; objects which can be represented by space-time trajectories. Not any trajectory will do, however. The space-time trajectory of a cricket ball is just the right sort of object to be a causal process. It is not a gerry-mandered object, and it is clearly capable of effecting causal influences on things (one may readily test this claim by throwing a cricket ball at a window, for instance).

A roving spot of light on a wall, however, is not a causal process, even though it describes a continuous trajectory in space-time. Although the light is interacting with the wall, causing it to radiate energy, it is not the moving spot that brings this about: it is the beam of light coming from the source that is the relevant process. World-lines such as that of the spot are known as 'pseudo-processes'. Below, I review the main attempts that Salmon and Dowe have made to distinguish causal- and pseudo-processes.

In his earlier work, which is perhaps the best known account of causal processes, Salmon claimed that the spot of light on the wall is not a causal process because it is not *able to transmit a mark* (1984: 142). That is, no local modification can be made to the spot

18. Thanks to Alastair Wilson for prompting this thought.

which would endure in subsequent stages of the process. You can't, for instance, paint a polka-dot on the spot and have it propagate to the subsequent stages of the spot's history.

This account, by appealing to an *ability*, seems to invoke a basic sort of disposition (or at least a counterfactual consideration) to characterise the essence of a causal process. This is obviously not congenial to my project of grounding causal powers (which are surely a sub-species of disposition) in causal processes.[19]

Salmon abandoned this account, precisely because it contained unanalysed counterfactual concepts, and instead adopted a proposal of Dowe's (1992), that a causal process is one which transmits *a conserved quantity* (1998: 257–8). For Salmon, a process 'transmits' a conserved quantity between two points *A* and *B*, if and only if the process possesses the quantity at *A*, at *B*, and at every stage in between. Moreover, it must do this *without* interacting with anything else. The roving patch of wall that is illuminated by the spotlight may well possess a quantity of energy that is invariant at every stage of its history, but it does not – and cannot – do this, without interacting with anything else. Rather, it must be engaged in constant interactions with incoming photons from the light source in order to possess this invariant quantity.

Dowe complains that this account has a number of deficiencies, the details of which need not concern us here (see Dowe 2000: 114–19). In turn, Dowe suggests that the roving patch of wall, even though it may possess a conserved quantity, fails to be a causal process because it is a *time-wise gerrymandered* object: it lacks genuine identity through time (98–101).

For our purposes, it is not necessary to adjudicate between Salmon's later account and Dowe's account of causal processes. Rather, what is crucial is that we establish that some sort of account of causal processes can be developed, along broadly similar lines, which preserves the idea that causal processes have the sort of structure which meets the three requirements indicated above: the requirements of intrinsicness, identity, and masking.

6.1 Immediate problems in the Dowe–Salmon account There are three evident obstacles to employing a Dowe- or Salmon-style account to the end of identifying causal structure suited for natural kinds.

First, the very idea of a *conserved quantity* might be tacitly dispositional. It is plausible to think that a conserved quantity is not merely one which *happens* to be conserved, but one which *would be* conserved under a variety of counterfactual circumstances. Consequently, conserved quantity accounts of causal processes might be uncongenial in much

19. Though it would admittedly be an interesting result if all the many and varied causal powers are grounded in a single basic sort of disposition: the ability to transmit a mark.

the way that Salmon's earlier, mark-transmission theory, was uncongenial for the purposes of giving an account that grounds facts about dispositions.

But there is no reason to insist that conservation be understood in this way. Conservation might be explained in terms of some other nomic concept, or it might be explained in purely actualist terms, as a brute accident.[20] Even if conservation is a modal phenomenon, and its modal nature has to be appealed to in the explanation of dispositions, it still constitutes significant progress, in my view, to have explained modal features of dispositions in terms of a more narrowly delimited modal phenomenon: the conservation of certain physical quantities.

A second concern arises with respect to Salmon's account. By defining a process in terms of an "absence of interaction", it may appear to breach the intrinsicness requirement. That is, whether or not something is a causal process, for Salmon, depends not only on how that process itself intrinsically is, at every stage of its history, but also on whether or not it is interacting with another process.

This concern, however, seems to be misplaced. A causal interaction – for both Dowe and Salmon – is defined as "an intersection of world-lines which involves exchange of a conserved quantity" (Dowe 2000: 90; Salmon 1998: 253). Interaction is therefore something that cannot happen at a distance. Whether a region is an intersection of world-lines may not be strictly intrinsic to that region: one might need to look just outside the intersection to region to see if there are multiple processes leaving or entering. But even if this concession must be made, it is plausible that this is a *local* property, in the sense that it supervenes on any infinitesimally small neighbourhood of space-time surrounding the region of intersection.[21]

(So I am, at this point, suggesting that the intrinsicness requirement might best be relaxed to a requirement of *locality*. Notice that the putative causal structure that was captured by a neuron diagram failed, not only to be intrinsic, but also failed to meet this requirement of locality.)

A third concern affects Dowe's account. Recall that Dowe invokes the idea of a 'time-wise gerrymandered' object, and explicates this as one which does not have constant identity through time. So an object (call it *Nikita*) defined as the mereological sum of (*i*) the Sphinx during the year 1801; (*ii*) Freud's pipe during the year 1901; and (*iii*) the Sun during the year 2001 – would be a time-wise gerrymander. It is to be distinguished from a space-wise gerrymandered object, such as: the Sphinx, Freud's pipe, and the Sun, all during the year 1901. This latter object, though rather strange and unfamiliar, at least

20. Salmon notes this point of difficulty, and responds by explicitly embracing a mere regularity account of conservation: a quantity is conserved just in case it is, for all time, unchanged. "[I]t makes no difference whether that true generalisation is lawful; only its truth is at stake" (1998: 259). Dowe does not directly address the issue.

21. See Butterfield 2006: 714–6 for some brief discussion of locality and its relation to intrinsicness.

has the same identity conditions through its life, and so is eligible to be a causal process, says Dowe.

One could, of course, create a new predicate – *is ertnog* – which denotes the property of being the Sphinx in 1801, *or* Freud's pipe in 1901, *or* the Sun in 2001, and insist that the time-wise gerrymandered object Nikita does have consistent identity conditions: it is always ertnog. Exactly how Dowe might respond to such moves, I will not address here.[22] What I will address is the broader concern, that *however* we spell out the required concept of identity over time, it is essential that the account not rely upon extrinsic properties, lest the intrinsicness requirement be breached. (The remainder of this subsection may well be skipped, for those who are happy enough to take it on faith that an intrinsicalist account of identity through time can be given.)

Dowe considers there to be two ways in which identity through time might be explicated, consistent with his overall approach. The first is to insist that identity through time is strict numerical identity. In the jargon that has become popular, objects involved in causal processes *endure* through time, by being wholly present at many times, rather than *perduring* by having different temporal parts located at different times. The surprising implication of an endurance account, however, is that any properties of an enduring thing which are possessed only temporarily turn out to be relational, by the following sort of argument:

Suppose, for reductio, that a property P is possessed temporarily by an enduring object, a at t_1. Later, at t_2, a lacks P. So:

(1) Pa at t_1.

(2) $\neg Pa$ at t_2.

(3) a at t_1 = a at t_2.

By the indiscernibility of identicals, a at t_2 must be P. But we have already established that it is not-P. Contradiction.

We can most easily avoid this absurdity by denying that "is P at t_1" denotes having the same property as "is P at t_2". Rather, what is involved in each case is a relation between the object a and the time. The object a is P-wise related to t_1 and it is not P-wise related to t_2. No contradiction arises from this, but it comes at a serious cost: what we thought were intrinsic properties have turned out to be relations to times. In particular, this means our account has profoundly violated the intrinsicness requirement, for properties of causal processes, such as possessing a certain quantity of mass-energy, will turn out not to be intrinsic. Consequently, it would appear that the relevant structure of a causal process is doomed to be extrinsic.

22. See Dowe 2000: 98–101, where Dowe manifests an awareness of these concerns, and elects simply to leave identity through time as a primitive, unanalysed concept in his theory.

However, again, I think that this variety of extrinsicalism is much less threatening than the sort of extrinsicalism which motivated us to reject the neuron diagrams considered earlier. Here, the sorts of relations involved are relations to the space-time region in which the process is located. That is at least a *local* phenomenon, even if it is not strictly intrinsic.[23]

The second approach to analysing identity through time which Dowe identifies as compatible with his theory is a broadly Humean perdurance account in terms of similarity and continuity. On this approach, the different stages of a thing's history are distinct parts of a larger, temporally extended whole. Different temporal parts may have different intrinsic properties, in the same way that different spatial parts of a spatially extended whole can have different properties.

This second approach may be more problematic for the purposes of identifying causal structure, however, because it sometimes relies upon relatively global judgments about similarity and continuity, in order to find the most highly eligible candidate to be the future 'continuant' of the original entity.

To illustrate, consider a version of the famous thought experiment concerning the Parthenon.[24] A stone thief gradually removes the stones of the Parthenon, over a period of several years, until all have been removed. However, whenever a stone is taken, it is soon replaced by the caretakers of the Parthenon, so that, even after all the original stones have been removed, the structure that is on the site is extremely similar to that which was there before the thefts began. Is the structure that is present after all the stones have beeen stolen identical to the Parthenon, or is it merely a fake? Arguably, this depends upon what has happened to the original stones. If they have been shipped by the thief overseas, and reassembled into a Parthenon-like structure in San Francisco, then common-sense ontology would judge that the Parthenon has been moved to San Francisco. If the stones have been broken down to make gravel, typical intuitions are that the Parthenon survives on the original site. Therefore, whether or not the structure on the original site is identical to the original Parthenon depends on facts that are remote from that spatiotemporal region.

The Parthenon, being an artefact, may not be the sort of object that is particularly important for the purposes of a metaphysics of science. So these concerns may be somewhat out of place. For present purposes, I merely wish to register the danger that Dowe's account of causal processes, wedded to a continuity/similarity-style account of identity through time, is likely to run afoul of my project. Other perdurantist accounts may work,

23. In related vein, Bradford Skow (2007) argues that on all viable accounts of space, shape properties – often taken as the very paradigm of intrinsic properties – are in fact extrinsic. However the sort of extrinsicness he identifies is often – to my mind at least – relatively 'unalarming'. It involves relations to the spatial region in which the shape-bearer is located, or to necessarily existent entities, such as numbers.

24. See Dauer 1972, responding to an earlier paper by Smart (1972).

but they will have to ground facts about identity through time in considerations intrinsic to – or at least local to – the object in question.

So, depending upon how we characterize the notion of a causal process, it looks like we are likely to have to breach the intrinsicness requirement. But this has led us to a more moderate requirement of *locality*: the structure of a causal process must supervene at least on the spatiotemporal neighbourhood of the process. It appears that, on both Salmon's later account and on Dowe's account, this requirement can be met.

6.2 *Causation by disconnection* The Dowe–Salmon approach to causation is notorious for having difficulty in dealing with certain sorts of causation: causation involving "disconnection" is not going to show up as a continuous world-line, nor as a series of world-lines linked by interactions. Consider the claim that my plant died because I failed to water it. There is no world-line leading from my "non-watering" to the death of the plant. There are world-line processes that occur in the plant that lead to its death, and these processes would not have occurred if watering had taken place. But there is no actual physical connection between non-watering and death.[25]

This is clearly a problem if you think that causal processes must serve the role of underwriting an analysis of token causal claims such as '*A* is a cause of *B*'. But I am not interested in that purpose. That is, I am not attempting to give an analysis of token causal claims – indeed, for such claims, I suspect a counterfactual approach will be necessary. Rather, I want to know if a world-line account of a causal process can be used to capture the causal structure of a disposition's manifestation, and if it can do so in a way compatible with the three requirements I listed above.

6.3 *Causal structure in Dowe–Salmon processes* Having warded off these initial concerns, it seems clear that Dowe–Salmon processes are highly suited to capture the sort of structure that is of interest to us. To briefly review the two key problem cases that we have considered so far:

Billy and Suzy Whether or not the robot is present will make – as we have noted – some very small difference to the causal process of Suzy's flying rock. But if we allow ourselves some sort of abstraction away from minor details like that, there is surely going to be some common causal structure to the case where Suzy smashes the bottle in the presence of the robot and the case where she smashes the bottle in the absence of the robot. In both cases, what is crucial is that the rock-process reach the bottle with a certain momentum, and that it then participate in a causal interaction in which much of that momentum is transferred to the bottle, which then splits into a large number of fragmentary processes. This case does not naturally fit talk of dispositions, but if we

25. See Dowe 2000: Chapter 6 for Dowe's attempt to deal with this problem. See Schaffer 2004 for a spirited rebuttal.

allow ourselves to describe it in such unnatural terms, this seems to be a good approximation of what we might think essential to the manifestation of the rock's "disposition" to smash the bottle.

Masking versus non-interaction The three particles in Figure 4 do not indicate, from their trajectories alone, whether or not they are interacting in such a way as to mask their dispositions to attract or repel, or if they are in fact not disposed to interact at all. A Dowe–Salmon account of the masking case would presumably involve appealing to interactions with the *fields* generated by the charged particles. The non-interaction case would not involve such interactions. So there would be a clear difference in causal structure.

The above discussion of Dowe–Salmon processes is already enough, I suggest, to give us grounds for optimism about the prospects for identifying essential causal structure in the processes by which dispositions are manifested. But the reader might reasonably wonder whether such essential causal structure can be found in quantum phenomena. While lacking the expertise to give a comprehensive answer to this question, I will, in the next section, attempt to show that the use of Feynman diagrams in quantum field theory is highly apposite for the purposes of representing causal structure.

7 Feynman diagrams Feynman diagrams are an appealingly intuitive method of representing interactions in quantum field theory. What is particularly striking about them is that they seem to provide precisely the sort of structure that we have been looking for. They appear to represent interactions as discrete and local events, hence giving rise to high hopes that they will capture some sort of intrinsic or local structure. In momentum space, they do not merely represent inputs and outputs to a causal interaction, they also represent different *ways* that the output can come about. Hence they appear to be able to distinguish between cases of masking as opposed to cases of non-interaction.

As it turns out, however, what I have just said about Feynman diagrams might be a little misleading. What exactly are Feynman diagrams? They are devices to assist in calculating the probability of an outcome, given an initial state, in quantum field theory. Each line and vertex in the diagram is associated with a term or step in the calculation. So they serve a remarkable dual purpose of giving an iconic representation of an interaction, while also facilitating calculation of the probability of that interaction.

A few words on how to interpret Feynman diagrams: First, the temporal order of events is from the bottom to the top. Though, strictly speaking, it is only the events at the *upper boundary* of the diagram that are taken to be in a definite temporal order relative to those at the *lower boundary* of the diagram. Any apparent ordering of events in between these extremes is merely an artefact of the representation. This is because – apart from the upper and lower boundaries – Feynman diagrams are just graphs. They are defined by *topological structure* alone.

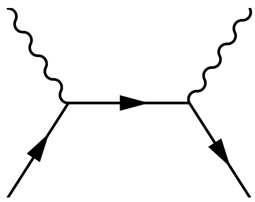

Figure 5: *A Feynman diagram "representing" positron-electron annihilation, resulting in the emission of two photons.*

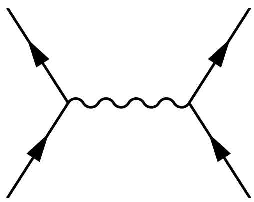

Figure 6: *Two electrons exchange a photon. This is the sort of underlying process that might account for a repulsive interaction between two negatively charged particles.*

Arrows pointing up represent electrons. Arrows pointing down represent positrons. Each straight section of arrow represents an electron moving at constant momentum for a period of time. Wiggly lines represent photons. Junctions of lines (i.e. vertices) represent absorptions/emissions of photons, positrons and electrons. That's it, for our purposes.

So a diagram like Figure 5 represents a positron and an electron annihilating, and two photons being emitted. And in Figure 6, the sort of process described would be the repulsion of two electrons. The change in momentum of the electrons is mediated by the exchange of the photon.

Because of conservation constraints, every emission or absorption of a particle is accompanied by a change in momentum. So the exchange of particles is associated with changes of state. This looks like a strikingly *discrete, local, and intrinsic* account of causal interaction. An extremely attractive thought, then, is that a Feynman diagram could indeed represent the degree of causal structure we are looking for in trying to identify natural kinds of causal processes.

7.1 Interpreting Feynman diagrams So what's the problem? The problem with this thought is that Feynman diagrams do not – at least in standard interpretations of QM – represent "what actually happens". Rather, they are an aid to the calculation of the probability of an outcome, given an initial state. Here is how you use Feynman diagrams in a calculation:

1. Take a given initial state – involving two electrons with given momenta, say.

2. Take a final state, in which they have different momenta.

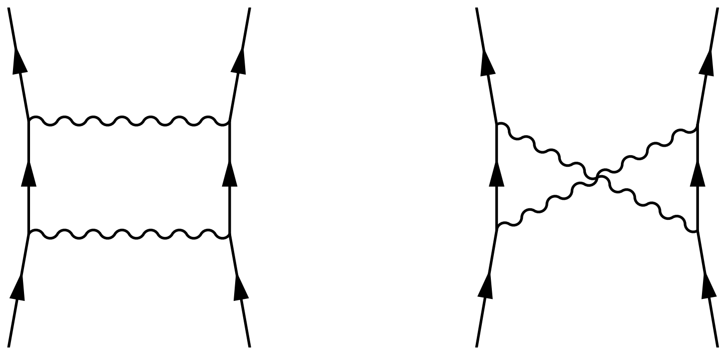

(a) First way (b) Second way

Figure 7: *Two further ways for two electrons to manifest repulsion.*

3. Then draw *all possible* Feynman diagrams which have these initial and final states.

So to calculate the probability that two electrons will repel each other, in addition to a diagram such as that in Figure 6, it would also be necessary to consider diagrams like those in Figure 7. All of these – plus infinitely many more – represent "ways" that the particles can go from the initial state to the final state. And each of them contribute to the calculation of the probability of the final state, given the initial state. (More precisely, each represents a term in a perturbation series, which is used as a means to gain an accurate approximation to the actual calculation. For most purposes of approximation, only the first one or two terms in the series need to be calculated – but in principle, all are required.)

Feynman was wont to say – or at least imply – tantalising things like: each diagram "represents a way" that the process can occur (Feynman 1985: 93). But if one is told an event can occur in one of N ways, and the event does occur, one would typically think it safe to infer that the event has occurred in one – *and only one* – of those N ways. But this sort of claim is not vindicated by the apparatus of QM. Rather, what can be said is that the state vector evolves in accord with the Schrödinger equation. The state vector's evolution is entirely deterministic, and there is no question of it evolving one way or another. The evolved state-vector yields a certain probability that we will observe a given outcome.

As far as orthodox QM is concerned, then, the individual diagrams in Figures 6 and 7 do not depict separate processes. They are merely devices we use to calculate the probability of a process that is specified in terms of:

— the initial particles' types and momenta, and

— the final particles' types and momenta.

In other words, processes are defined purely by input and output. Call any such process – specified in the above terms – an "orthodox QM-process".[26]

7.2 Do orthodox QM-processes have causal structure? So: could there be *natural kinds* of orthodox QM-processes? That is: does an orthodox QM-process possess intrinsic structure that might constitute the real essence of a natural kind?

I think we can hazard a positive answer to this question. While it is unsettling that we cannot identify the causal structure of an orthodox QM-process with any single Feynman diagram, as we might have wished, we can at least say that the structure of an orthodox QM-process is approximately represented by a *set* of Feynman diagrams. (The very fact that this seems to undermine traditional ways of representing a process as a single way that things happen is perhaps apt, given the stochastic and/or indeterminate nature of QM phenomena.) Such a set has intrinsic structure. It has intrinsic properties such as 'contains diagram *X* as a member'. Moreover, each of the members have associated amplitudes – so there will be further intrinsic properties of the set, such as: 'has a member with amplitude *n*'; 'contains members whose amplitudes sum to *p*'; etc.

However, although orthodox QM-processes are quite rich in structure in the above sense, they lack the sort of structure required to make the most rigorous sort of distinction between masking and non-interaction. Recall Figure 4, which was ambiguous as to whether it represented three uncharged particles undergoing non-interaction, or three charged particles whose repulsive and attractive forces cancel out. The two sorts of scenarios involved are different not only with respect to the causal interactions, but also in the *nature* of the participating particles. Only one of the cases involves charged particles.

Feynman diagrams – and the apparatus of QM in general – are perfectly capable of distinguishing cases of this sort, for the simple reason that orthodox QM-processes are defined, in part, by the types of particles in the initial and final states. So a process involving charged particles is ipso facto a different type of process from one involving uncharged particles.

But one might have thought it possible for there to be two distinct possible processes involving precisely the same final and initial conditions. In one case, the particles interact, but the interactions cancel out – giving rise to a case of masking. In the other, the particles simply do not interact. Although these two possibilities are incompatible with determinism, one might have thought they were genuine possibilities in light of the apparently indeterministic nature of quantum mechanics.

This distinction, however, is strictly not possible given the nature of orthodox QM-processes. No two distinct QM-processes exist that have identical final and initial states.

26. By 'orthodox' here, I do not mean to invoke thoughts of the Copenhagen interpretation. I merely mean approaches that do not reify what is represented by Feynman diagrams – and this includes a wide variety of relatively standard approaches.

So we cannot use the apparatus of Feynman diagrams to draw this more subtle distinction between masking and non-interaction. But as I noted in Section 3, the requirement to distinguish masking from non-interaction is a defeasible one, and this may be precisely the sort of case where it is apt to be defeated.[27]

8 Demarcating causal processes The above reflections, both upon the use of Feynman diagrams in physics, and on the nature of Dowe–Salmon processes, are grounds for cautious optimism about the possibility that the world contains the sort of causal structure that could constitute the essences of natural kinds of causal processes.

But, as mentioned in Section 5, there is reason to be concerned that, even if the world contains the right sort of structure, the processes that we end up with will still not fit the *identity* requirement, which requires that the kinds have intuitively plausible identity conditions. For more intuitively familiar types of natural kind, it is generally quite clear where each distinct member of a natural kind is located, and consequently, it is clear how many members of the kind there are. One could take a sample of chemicals, for instance, and – worries about vagueness aside – there would be a determinate fact about how many atoms of each chemical kind are present in the sample. Causal processes, in contrast, seem less amenable to determinate facts, both about identity and about location.

With respect to the numerical identity of causal processes, any bit of space-time that contains more than one or two fundamental particles will likely involve an enormous plurality of causal influences and phenomena. There seems to be no non-arbitrary manner of counting up the 'number' of causal processes that are present in that region. With respect to the location of processes, given the possibility of events in the past or future having an effect on a causal interaction, very distant phenomena might need to be included in our understanding of a causal process. Consequently, causal processes might be surprisingly large. If we have to take account of all the past influences, we might need to include the entire backwards light cone of the objects that are paradigmatically thought to be involved in the process.

Putting these two concerns together, it begins to look like the only non-arbitrary way of defining a causal process is simply to take the *entire universe* as a single process, and give up hope of finding smaller processes that are members of genuine natural kinds.

This response would be an overreaction. The identity conditions for process-kinds may not be as transparent to us as they are for other kinds, but that fact is not, in itself, grounds for despair that there could possibly be such conditions. Rather, there are at least four grounds for resisting the sorts of worries cited above.

27. See Handfield 2008b: 304–5 for further discussion of the issue of finking and masking for the fundamental properties. Some dispositionalists have even suggested – on quite independent grounds – that fundamental causal powers may not be susceptible to finks or masks (Bird 2007: §§3.3.2–3.3.3). For such dispositionalists, this convergent argument might be quite welcome.

In the first place, chemical compounds are often thought to be excellent examples of natural kinds, but chemical compounds are themselves subject to vague identity conditions. Consider two molecules of H_2O that are undergoing the process of converting into one molecule of H_3O^+ and one hydroxide ion (OH^-). This involves one of the hydrogen atoms being pulled out of its covalent bond, and forming a hydrogen-bond with one of the unbonded electron pairs on the oxygen atom of the other molecule. When this sequence of events is examined on the micro-scale, it is not the case that there is some single moment at which the hydrogen nucleus is unequivocally 'released' from its original covalent bond. Rather, we are confronted with the same problem we have in working out when we have stepped off a mountain, or when a man becomes bald, or when a collection of grains of sand constitutes a heap. There are numerous possible places at which we could equally plausibly draw the line, and no clear fact of the matter as to which is the correct one.[28]

Secondly, as the example of water illustrates, we are already reasonably familiar with the possibility that distinct kinds can 'overlap'. A molecule of water contains parts – such as protons, electrons, hydrogen atoms, and so on – that are themselves members of other natural kinds. So the fact that the world appears to contain numerous overlapping causal processes is not itself fatal to the possibility of there being distinct kinds involved.

In the case of causal processes, to accommodate such overlap, we cannot merely hope to distinguish the distinct processes by spatio-temporal means. A sphere of metal may be undergoing a process of rotation as well as a process of heating, for instance.[29]

The typical method by which we try to distinguish such intertwined processes is to use a degree of idealisation – we can ask what is the trajectory the relevant objects would take *in the absence of any other processes*?[30]

Where causal influences interact in complex, non-linear ways, we will have less reason to think that the answer to hypothetical questions like this will be of much use in

28. Even more unsettling to our comfortable view of water as involving a clearly defined, stable structure, recent results in physics suggest that, when observed on extremely short time-scales, water behaves like $H_{1.5}O$ (Schewe, Riordon, and Stein 2003)! While the interpretation of these results is still uncertain, it reminds us of what should be a familiar fact: that whatever traits we take to be essential to water, we may, in light of empirical discovery, have to live with various qualifications and hedges, to the effect that these are merely *typical* properties of water.

29. The example is due to Davidson, who was discussing the ontology of events (1980: 178).

30. In defending his account of causal processes, Wesley Salmon had to tackle much the same issue, because his definition of a causal process requires the absence of interaction.

> "You'd want to say that the speeding bullet transmits energy-momentum from the gun to the victim, but what about its incessant, negligible interactions with ambient air and radiation?" Of course. In this and many similar sorts of situations, we would simply ignore such interactions because the energy-momentum exchanges are too small to matter. Pragmatic considerations determine whether a given 'process' is to be regarded as a single process or a complex network of processes and interactions. (Salmon 1998: 258)

identifying genuinely distinct processes. But in many cases it does seem that the multiple influences combine in a linear fashion, or in some other fashion that makes their independent reality seem plausible. Accordingly, we have some reason to think that there could be genuine kinds of causal process at play.[31]

Thirdly, anyone who wishes to believe in natural kinds will surely think that fundamental particles such as electrons are among the very best candidates for such status. But arguably we are beset with even greater difficulty in providing determinate identity conditions and determinate locations for electrons. The indeterminacy of an electron's location is a well-known phenomenon of quantum mechanics. Less well known is that there are some parts of quantum field theory in which the very number of electrons (and other particles) is indeterminate.[32] So if we are to have any natural kinds at all at the quantum level, it appears we must put up with a degree of indeterminacy in those kinds. That should make us more comfortable, in turn, with some more limited indeterminacy with respect to natural kinds of processes.

Finally, and most importantly, by adopting approaches such as the Dowe–Salmon account of a process, rather than a neuron diagram style approach, then we can analyse the causal structure of processes in terms of *actual influence* rather than merely counterfactual potential-to-influence. Consequently, we can legitimately restrict our attention to the influences *in a given space-time region*. In doing that, we are not required to take into account all of the possible threats and hypothetical interveners that could have influenced what happened in that space-time region. I suspect it is only when we assume that causal structure must include merely counterfactual influences that we will be forced to blow out the location of a causal process, with the eventual result being that we must include the entire backwards light-cone of an event in order to include all the causal structure that goes into the process. The conserved quantity approach forestalls this difficulty.

9 Conclusion It is uncontroversial that the world contains causal processes. What is more doubtful is that those processes contain a sort of structure which might be suitable to serve the purpose of constituting *natural kinds* of causal processes. In particular, it seems doubtful that causal processes contain a suitably intrinsic structure.

However, if we relax our intuitive pre-conceptions about the nature of natural kinds, we can see that there may well be natural kinds of causal processes. In Dowe–Salmon accounts of causal processes, we have a relatively well-developed and empirically informed account that is compatible at least with highly localised causal structure, even if it is not strictly intrinsic. Moreover, in current physics, the use of Feynman diagrams is an exam-

31. For a recent discussion of the relation between the reality of component causal influences and the analytic method in science, see Corry (2009).

32. See Teller 1995: 110–13.

ple of a method of representation that, despite incorporating a fair degree of quantum ontological strangeness, is still sufficiently rich that it appears to represent sufficient causal structure to constitute natural kinds of causal process.

I have done nothing to address those who are sceptical that the world contains natural kinds of any sort whatever. In particular, I have done nothing to address scepticism with regard to the de re essential properties that are characteristic of such kinds. But on the basis of the above discussion of causal structure, I conclude that, if the world contains any natural kinds at all, it is quite plausible that it contains natural kinds of causal processes.

School of Philosophy and Bioethics
Monash University

Acknowledgments I wish to thank Stephen Barker, John Bigelow, Alexander Bird, Richard Corry, Phil Dowe, Antony Eagle, Harry Perlman, Dean Rickles, Neil Williams, and Alastair Wilson for their valuable comments on earlier drafts of this paper. Thanks also to the audience at the conference 'Powers: their grounding and manifestation', held at Lady Margaret Hall College, Oxford 2008.

REFERENCES

Armstrong, D. M. 1997. *A World of States of Affairs*. Cambridge: Cambridge University Press.

Bird, Alexander. 1998. "Dispositions and Antidotes". *The Philosophical Quarterly* 48: 227–34.

———. 2007. *Nature's Metaphysics*. Oxford: Clarendon Press.

Butterfield, Jeremy. 2006. "Against Pointillisme about Mechanics". *British Journal for the Philosophy of Science* 57: 709–53.

Choi, Sungho. 2003. "Improving Bird's Antidotes". *Australasian Journal of Philosophy* 81: 573–80.

Corry, Richard. 2009. "How is Scientific Analysis Possible?" In *Dispositions and Causes*, edited by Toby Handfield. Oxford: Oxford University Press.

Dauer, Francis W. 1972. "How Not to Reidentify the Parthenon". *Analysis* 33: 63–4.

Davidson, Donald. 1980. *Essays on Actions and Events*. New York: Oxford University Press.

Dowe, Phil. 1992. "Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory". *Philosophy of Science* 59: 195–216.

———. 2000. *Physical Causation*. Cambridge: Cambridge University Press.

Ellis, Brian. 2001. *Scientific Essentialism*. Cambridge: Cambridge University Press.

Ellis, Brian, and Caroline Lierse. 1994. "Dispositional Essentialism". *Australasian Journal of Philosophy* 72: 27–45.

Fara, Michael. 2001. "Dispositions and Their Ascriptions". Doctoral thesis. Princeton University.

Feynman, Richard P. 1985. *QED: The Strange Theory of Light and Matter*. Princeton: Princeton University Press. Published in paperback by Penguin, 1990.

Handfield, Toby. 2008a. "Humean Dispositionalism". *Australasian Journal of Philosophy* 86: 113–26.

———. 2008b. "Unfinkable Dispositions". *Synthese* 160: 297–308.

Hitchcock, Christopher. 2001. "The Intransitivity of Causation Revealed in Equations and Graphs". *The Journal of Philosophy* 98: 273–99.

Johnston, Mark. 1992. "How to Speak of the Colours". *Philosophical Studies* 68: 221–63.

Lewis, David. 1986. *Philosophical Papers*, vol. 2. New York: Oxford University Press.

———. 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.

———. 2004. "Void and Object". In *Causation and Counterfactuals*, edited by John Collins, Ned Hall, and L. A. Paul. Cambridge, Ma.: M.I.T. Press.

Maudlin, Tim. 2007. *The Metaphysics Within Physics*. Oxford: Oxford University Press.

Menzies, Peter. 1999. "Intrinsic versus Extrinsic Conceptions of Causation". In *Causation and Laws of Nature*, edited by Howard Sankey. Dordrecht: Reidel, 313–29.

Oppy, Graham. 2000. "*Humean* Supervenience?" *Philosophical Studies* 101: 77–105.

Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Salmon, Wesley C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

———. 1998. "Causation without Counterfactuals". In *Causality and Explanation*. New York: Oxford University Press, 248–60.

Schaffer, Jonathan. 2004. "Causes Need Not be Physically Connected to their Effects: The Case for Negative Causation". In *Contemporary Debates in Philosophy of Science*, edited by Christopher Hitchcock. Blackwell, 197–216.

Schewe, Phil, James Riordon, and Ben Stein. 2003. "A Water Molecule's Chemical Formula is Really Not H_2O". *Physics News Update* 648.

Skow, Bradford. 2007. "Are shapes intrinsic?" *Philosophical Studies* 133: 111–30.

Smart, Brian. 1972. "How to Reidentify the Ship of Theseus". *Analysis* 32: 145–8.

Teller, Paul. 1995. *An Interpretive Introduction to Quantum Field Theory*. Princeton: Princeton University Press.

Wilson, Jessica. Forthcoming. "What is Hume's Dictum, and Why Believe It?" *Philosophy and Phenomenological Research* .