Spoils to the Vector: How to Model Causes If You Are a Realist About Powers

1. Introduction

Causes and powers are closely connected. Powers are plausibly taken to be productive of their manifestations. A vase breaks when struck, for instance, because of its fragility, while sugar dissolves in water because it is soluble. A relatively recent addition to the theories of causation has therefore been one based on an ontology of real powers or dispositions (Harré and Madden 1975, Bhaskar 1975, Cartwright 1989, Molnar 2003 and Martin 2008). Under such a view, every caused event is produced by one or more powers manifesting, which is why dispositions are sometimes called 'causal powers'. No doubt, their role in causation is one of the chief reasons philosophers are interested in powers.

Despite this being well known, the theory of causation based on an ontology of powers has received relatively little development. It is one thing to have the basic insight that all caused events are the manifestations of powers, but it is another thing to provide a detailed theory of causation grounded in that insight. By comparison, Lewis's theory of causation based on counterfactual dependence between events has been worked out in great detail even though its bedrock—there being similarity relations that hold across equally real but spatiotemporally disconnected concrete worlds—is one that many believers in powers would find unpromising if not ridiculous. The successful provision of detail has nevertheless allowed the counterfactual dependence view to take the initiative and in some respects leave the causal powers theory behind. This must be rectified.

What follows cannot provide all that is needed but it makes a start on a more detailed powers-based theory of causation. It concerns how we might represent causal situations if they were based in powers doing their work. This may seem like a relatively small matter but we say that it is not. We follow Hitchcock when he says, with particular reference to causation, that ". . . the way in which we choose to represent some phenomenon can shape the way in which we think about that phenomenon" (2006, 69). Representing the operation of powers in the right way can be crucial, therefore, in getting our power-based theory of causation off on the right footing.

2. Neuron Diagrams and Their Hidden Ontological Commitments

By way of contrast, we start by explaining how a powers ontologist does not want to represent causation; namely, by neuron diagrams. Neuron diagrams have become a common currency among metaphysicians of causation since David Lewis (1973) introduced them into the debate. It is arguable that they have become the main way of representing causal situations. It is important to emphasize, however, that neuron diagrams are neither metaphysically neutral nor theoretically innocent representations. They contain certain tacit ontological commitments about the way that causation works (see Hitchcock 2006). These commitments are likely to be far more acceptable to a Humean or a Lewisian than to an upholder of powers.

A neuron diagram (for example, in figure 1) depicts a series of neurons, *a*, *b* and *c* that may bear stimulatory connections to each other, represented by arrows. When a stimulatory neuron fires, any neuron to which it bears an upstream stimulatory connection also fires. The firing of a neuron corresponds to the occurrence of an event, while the stimulatory connection corresponds to a causal relation holding between events. This is particularly conducive to a Humean view of causation because it suggests that the events *a*, *b* and *c* are entirely discrete and self-contained. Their identity and nature is in no way determined by the causal relations in which they stand. Event *a* is the same event even if it bears no stimulatory connection to *b* and the causal relations involved are, in that respect, contingent. This is unobjectionable to those of a Humean inclination, according to which everything is loose and separate (Hume 1748, 54). But it would not be true in a dispositional theory of causation where dispositions are for certain manifestations essentially. How things stand in relation to each other causally is, in part or in whole, constitutive of what they are.

"Spoils to the Vector: How to Model Causes
If You Are a Realist About Powers" by Stephen Mumford & Rani Lill Anjum,
The Monist, vol. 94, no. 1, pp. 54–80. Copyright © 2011, THE MONIST, Peru, Illinois 61354.

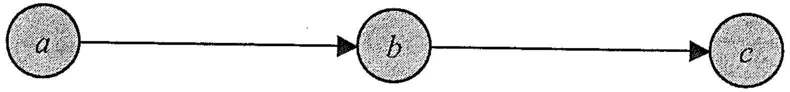

Figure 1: A neuron diagram

Apart from this, there are a number of other problems with neuron diagrams for the dispositionalist. Although neuron diagrams on the one hand seem to promote pure contingency or Humeanism concerning what causes what, they also on the other hand suggest a form of necessitarianism. If *a* fires or occurs, and it bears a stimulatory connection to *b*, then *b* must occur also. A cause is thus depicted as being entirely sufficient for its effect when it succeeds in producing it. Similarly, an inhibitor neuron guarantees that the effect does not occur, even if there is also a stimulatory neuron firing. This is the case where one neuron bears an inhibitory connection to another (see figure 2).

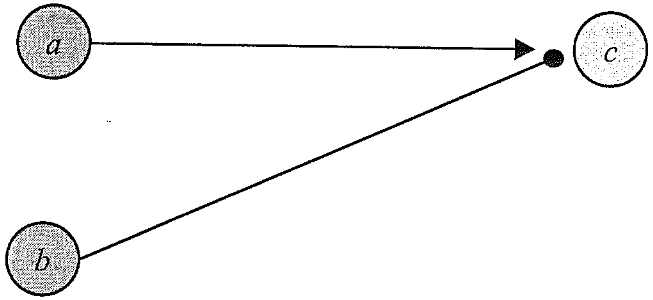

Figure 2: An inhibitory neuron b

On a dispositional theory of causation, however, it is arguable that a cause is something that disposes and no more than disposes towards its effect. Even in cases where the cause succeeds in producing the effect, it cannot do so by necessitation, for something could have prevented the manifestation from occurring even if in fact it did not (Mumford and Anjum 2010). Instead the modality needed for causation would have to be something weaker than pure necessity, yet stronger than pure contingency. The alternative is that the causal modality is irreducibly dispositional: a tendency towards an effect that need not always be manifest (Anjum and Mumford 2011). A tendency is more than mere possibility, because it is directed towards a particular outcome. But it is also weaker than necessity, since it can be prevented or counteracted by other powers. This means that, contrary to how the neuron diagram is set up, a stimulatory connection should not guarantee the effect. It disposes towards an effect, and may succeed in producing it, but if it does so it does not do so through necessitation. Similarly, an inhibitor can dispose away from something happening, without guaranteeing that it doesn't.

A further problem with the neuron model is that it works only for binary cases, where some cause or effect either occurs (fires) or it does not. But what of cases where an effect occurs with some probability, such as when a coin is disposed to land heads with a 50% chance? Another possibility is that an effect occurs to some degree, perhaps as a function of the degree of the cause. Various things may dispose towards the heating of water, for instance, but the kettle does so to a higher degree than does the sunlight. And as well as the degree of the cause, its frequency could also be a factor. A page may break loose from a book on its thousandth turn but it is not just that final turn that causes the break. Each of the preceding turns also made some small contribution.

This last point also brings to light what will be a crucial factor for us. The causes of an effect are often very complex, involving many different powers of many different things. In neuron diagrams there seems to be only one cause of an effect. Neuron diagrams have to be this way because they were developed as part of the counterfactual dependence view and, of course, if there were two stimulatory neurons for the same effect, each sufficient for it, then the effect would counterfactually depend on neither. Complexity is thus a problem for the theory. Lewis's response is to say that his is a theory only of what it is to be "... one indispensable part, not the whole, of the total situation that is followed by the effect ..." (Lewis 1973, 159). The cause in question is therefore not the notion of total cause, associated with Mill (1843, III, v, 3), but rather *a* cause among others. Neuron diagrams do not rule out complexity of causes, therefore, but for them to permit a counterfactual reading of the stimulatory connection, what they must rule out is that causes can be overdetermined. It is not obvious to us that such overdetermination should be precluded automatically, as we will explain later. But it should be remarked additionally that by focusing on just one cause—the notion of *a* cause rather than *the many* causes—the account of causation misses out on some important factors. How powers work with each other, and sometimes against each other, we will see is a vital matter for the dispositional theory and consideration of it will lead us to some key features of causation. What we need, therefore, is a better way of representing causal situations: one that is more sympathetic to an ontology of powers and what they would bring to the theory of causation. We need a causal model for a truly interconnected, rather than 'loose and separate' world.

There is, of course, another way of modelling causes in addition to neuron diagrams, namely the structural equations developed by Pearl (2000). This kind of causal modelling, however, concerns primarily causal reasoning. It is not an attempt to explicate the nature of causation itself but only to show what the causal structures are. Such understanding can then be used for prediction because the causal graphs produced entail a host of counterfactuals. For that reason, such causal graphs can be seen as an extension of the counterfactual dependence view (Hitchcock 2009, 310). The objectionable features of neuron diagrams transfer over. What they do not show is how one factor can dispose towards another without this being in terms of necessitation: either by being entirely sufficient for the effect or, in the case of probabilistic causation, sufficient for the probability of an effect.

Our aim here is not to destroy neuron diagrams or causal graphs. That would be far too big a task and, in any case, these are merely representative conventions. In theory, even neuron diagrams might be amended in some way as to make them more conducive to a dispositionalist ontology. But we will not try to do that. Instead, we will offer a new way of representing powers that we think will make clearer their causal role and the nature of causation in a world of powers.

3. *The Vector Space and Powers*

Causation involves properties or qualities and frequently concerns change in properties, such as an increase in temperature. To model causation, therefore, it will be useful if we start with a quality space on which to plot such changes. This notion is taken from Lombard (1986) and gives us a blank canvas on which we can depict causal situations. We can start with the simplest cases, which will be one-dimensional. Here we have a single quality that permits a determinable range of changes with two end points or termini, though in theory a determinate range could be without end, if there is no maximum temperature, for instance. In very rough terms, at one end could be the property of being cold and at the other the property of being hot. For a caused change in colour, such as in the ripening of a tomato, we could have green at one end and red at the other. For the moment, we are being quite loose and unscientific about the properties involved. We can label the properties F and G (or F and ¬F if the properties are opposites) at either extreme of the quality space. The quality space will have a starting point, which represents the initial position from which causation will operate. Suppose we are interested in a particular room temperature, for instance, and whether, from its current position, it is caused to increase or decrease. The quality space that concerns us can be represented as a horizontal dimension with a central vertical line as our starting point (see figure 3): the temperature of the room in question, for instance, or the current colour of the tomato. Any changes from that point would then involve a movement left (towards F: getting warmer, for instance) or right (towards G: getting colder).

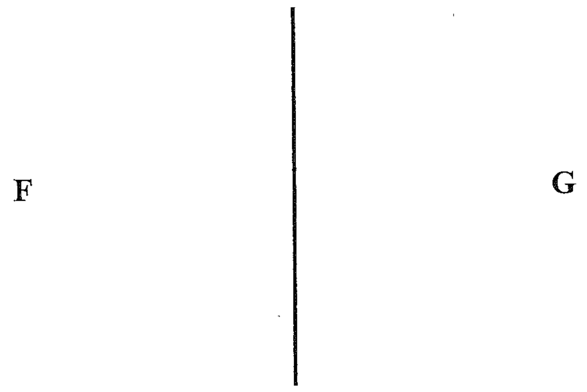

Figure 3: A one-dimensional quality space

Powers are the drivers of change on this view of causation. On to the quality space, therefore, we want to plot powers. These will be powers that act on the relevant quality space. To continue the example of the room temperature, we are interested in any power that operates on the temperature: disposing towards increasing it, decreasing it, or maintaining it. Suppose we turn a heater on in the room. It will dispose towards it warming. We propose that we plot this as a vector operating on the quality space in question, as depicted in figure 4.

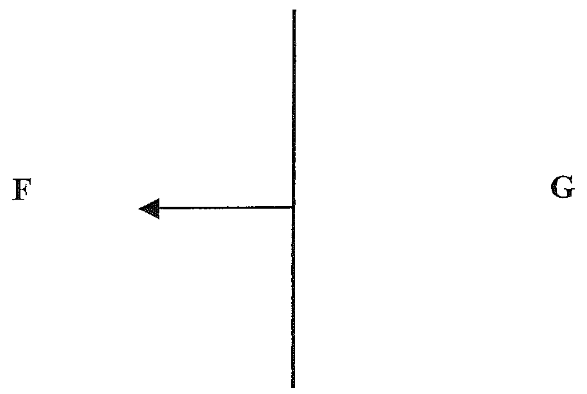

Figure 4: A vector representing a power

We are taking the idea of a vector from physics though our vectors need not have all the features of vectors as used in physics. Vectors, however, have two essential features on which the theory of powers may trade. In the first place, a vector has a direction. Powers have something closely akin. A power is for some outcome or manifestation, and essentially so. If one follows Molnar (2003, ch. 3), for instance, the power truly is directed towards this manifestation (though we do not subscribe to the physical intentionality account in any literal sense). The power that the heater has is essentially a heating one (no doubt capable of some more technical specification, suitable for the needs of science). We plot this power as a vector directed towards F, therefore, as it is something that disposes towards increasing the temperature, within the context of the room at its starting point. Many other powers that the heater has will not be causally relevant to the situation modelled in the diagram, though they may be relevant to some other change in another quality space. Only those powers that dispose towards F or G will show up as directed vectors on the F-G quality space. The second essential feature of vectors is that they can have an intensity or magnitude, conventionally represented by the length of the vector. Powers can also have an intensity, a point that is frequently overlooked (see Manley and Wasserman 2007), which we can indicate using the same convention. Some heaters are better than others and some can be put on a higher or lower setting. We can represent these differences by the length of our vectors.

In the case of figure 4, there is just one power operating on the quality of temperature. As we remarked above, however, most cases of causation will involve many powers working together, or sometimes against each other. Causation is complex, by which we mean that the events typically caused in the world we inhabit are polygenous. Polygeny (a term adopted from biology by Molnar 2003, 194–98) means that effects have many causes. On the powers account, this means that any actual change, such as the change of temperature, will be the result of many powers at work. As well as the heater in the room, other things could also dispose towards it warming: the body heat of people in the room, for instance, and the sun shining in through the window. But, at the same time, various powers may be disposing towards the cooling of the room. There is a draft from the window, for instance, and poor insulation in the roof. The temperature to which the room moves, if it moves at all, will thus be caused by all these powers acting together. All sorts of factors could make a contribution, even if it is a very small one. A known causal factor placed into different contexts can result in different outcomes, and these contextual factors we take to be themselves powerful contributors to those outcomes, even if individually their contribution is relatively small. The change in temperature will be a single event, but clearly a polygenic one. We can plot all these different powers on the vector space as in figure 5, with powers *a*, *b* and *c* disposing towards warming the room and powers *d*, *e* and *f* disposing towards it cooling.

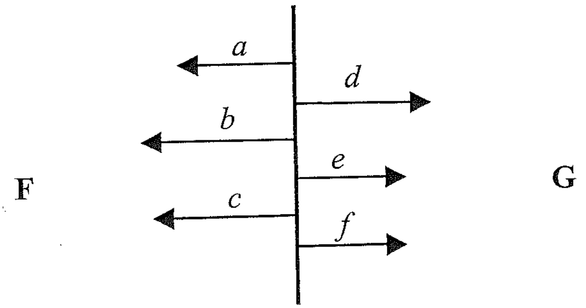

Figure 5: Multiple powers at work

What exactly does a vector in our diagrams represent? Is it just a power? Should we say it's a stimulated power? Or perhaps the manifestation of a power? We want to say that a vector represents a power that is operating: it is active, doing its work, it is exercising and thus causally relevant. It is, therefore, what some people will think of as a power that has been stimulated. But this issue is far more complicated than is traditionally understood, even by those working on dispositions. It is sometimes suggested that powers need stimuli in order to do their work, as if they will sit, passively, doing nothing until they are made to do so. But this is a form of passivism (see Ellis 2001, 7), a view that real powers should count against. A realist about powers, especially if also a pandispositionalist (see Shoemaker 1980 and Mumford 2008), sees the stimulus itself as just one further power in the whole complex of causes. Clearly, stimuli must themselves be powerful. And instead of passive powers being stimulated, Martin's idea of mutual manifestation partners looks more conducive to the powers view (Martin 2008). Hence, if a power, represented by one vector, meets its partner, represented by another, then they begin to do their work together, neither of them needing any additional stimulus. We have a partial rejection of passivism, therefore. Certainly some powers can be inactive or inoperative, doing no work at a particular time. It is, after all, a hallmark of powers that they can exist unmanifested. But the kind of causal situation typically represented in a vector diagram is one in which mutual partnerships have already been made and hence in which the powers are acting.

There is one further idea that can be taken from vectors as used in physics, the idea of vector addition. Contrary to claims that have been made by Cartwright (1983, 59), we take it that powers can add. Cartwright objects to this on a number of grounds. One of those is that addition is merely a metaphor that doesn't apply to the physical world. We perform addition, in our heads or on paper, but powers can't do that. In contrast, we take it that the metaphor of addition can, in at least some cases, apply to powers. Multiple powers can operate on the very same quality space and the eventual movement within that space will be a product or function of all the operating powers and their extents. To take a very simple tug-of-war case, for instance, each team member adds their own power to the contest. Some pull west, and some pull east, and these can be understood as dispositions east and dispositions west. In this case, such dispositions can be explicated in terms of forces. Any actual movement of the rope is the product of the total force west, minus the total force east, with each contestant adding their own power to the total. As we will see later, there are many cases in which simple addition is not complex enough to explain this phenomenon, which is basically what Mill (1843, III, vi, 3) called the composition of causes. But in at least some instances, vector addition will be an appropriate comparison. Figure 6 shows such a case, where we add all the powers disposing towards F and subtract all the powers disposing towards G. We are then left with a resultant vector R (indicated by a thicker line) that shows us overall in which direction and to what intensity the complete situation is disposed, with all the relevant powers in operation.

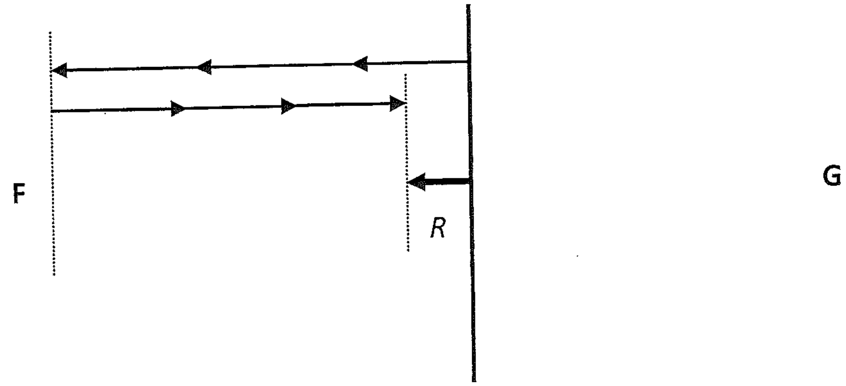

Figure 6: Vector addition for calculation of the resultant R

4. Causation of Absence

One case we might want to consider is that in which nothing conspicuous happens, but where the fact that nothing does is still nevertheless a case of causation. This could be classed as what Dowe (2001) calls causation *of* absence, not to be confused with causation *by* absence, which will be discussed later. Some theories of causation have difficulty with causation of absence and, indeed, neuron diagrams seem unable to represent it. If something is caused, it must happen. In other words, if a stimulatory neuron fires, the stimulated neuron also fires. But what we want is a case where the effect is that nothing happens—things remain unchanged—which might nevertheless be a genuine case of causation. In a neuron diagram, perhaps this could be accommodated by treating the non-occurrence of something as itself a kind of occurrence, but we would take this to be metaphysically misleading (though there are some conceptions of events, such as the property exemplification view, in which non-occurrences do count as events). Two books might lean against each other; each propping the other up at an angle, and the effect is that neither falls, which they would if the other were not there. A magnet might sit motionless on a fridge, with nothing happening; specifically, it doesn't fall. And sometimes there is stability within a movement, such as the orbit of the Earth maintaining a stable distance from the Sun. Causation of absence is easy to understand and represent using the vector model. We have some powers at work disposing towards F and other powers at work disposing towards G but they cancel each other out. The overall situation is a zero resultant vector (represented by a dot). If we return to the tug-of-war case, for instance, we can see that when teams are equally balanced, the rope goes nowhere. But this is nevertheless causation, according to the powers theory, because powers at are work. There is a world of difference between the case where the two teams pull but are equally balanced and the case where neither team has started to pull, even though, in respect of the movement of the rope, the situation is the same: it goes nowhere. The former case is causation but what can be called an equilibrium case, with a zero resultant vector (figure 7). The latter is not causation at all (in respect of the pertinent quality space). But the difference can be represented. Where there is no causation, the situation is merely that as in figure 3: just a quality space without any vectors.

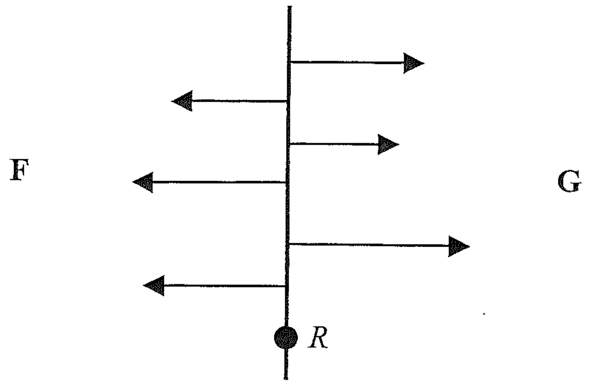

Figure 7: 'Equilibrium', or a zero resultant vector

5. Causal Production and Interference

It is useful to employ an additional notion to account for what happens in some cases of causation. Often there are involved what we may call threshold phenomena where some novel or interesting change occurs. A simple example is the dramatic change that occurs where water reaches 100°C and, at that point, undergoes a transformation. This may be a threshold we endeavour to have met in some circumstance, for instance; when we boil a kettle. Other thresholds may be less dramatic. Our interest may be in heating a room temperature to 20°C simply because that is what we prefer and we then seek to put things in place that will produce that. We represent this by adding a broken vertical line T to our vector diagram, which shows the place on the one-dimensional quality space at which a particular state or determinate quality is achieved (see figure 8).

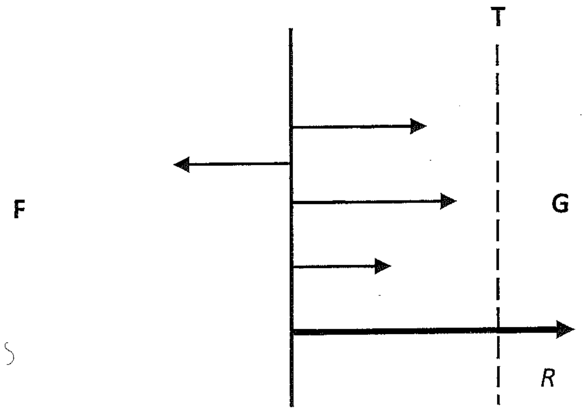

Figure 8: A resultant vector passes threshold T

It is worth noting that in considering the reaching or passing of a threshold, we should take into account not only those powers that dispose towards it. Whether or not some assemblage of powers produces the effect in question—the passing of threshold T—will be dependent not just on those powers disposing towards it but also the powers that dispose away from it and their extent. This is because some assemblage of powers that in one situation would be enough for the reaching of a threshold could in another context not be enough. We could have the very same powers, *a*, *b* and *c* that dispose towards G and in one situation reach T, a threshold in the G-direction, but in another situation do not. This is because a further power *p*, that disposed away from G, could be added such that the resultant, even though *a*, *b* and *c* are included just as before, no longer reaches T. The issue of prevention or interference has long been acknowledged as one that impinges on dispositions. There are always some circumstances or conditions under which a disposition can be prevented from manifesting. This has been regarded as a problem in particular for those who seek to defend the conditional analysis of dispositions. The problem is that for any set of circumstances outlined in the antecedent of the conditional, it is always possible to add something further that renders the conditional false. Defenders of the conditional analysis have therefore had to seek some way of excluding such extra circumstances. One would then have to exclude the possibility of adding extra powers that might interfere with the causal situation. We can represent such interference in figure 9 where an additional power *p* can be seen to prevent the situation overall reaching the threshold T. The powers other than *p* are the same as those appearing in figure 8, in which T is passed.

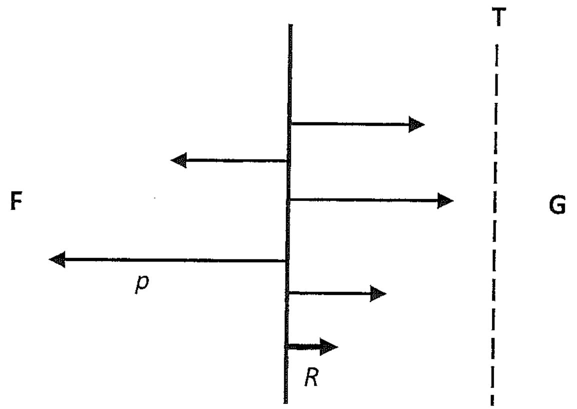

Figure 9: Additive interference

The case depicted in figure 9, which we call additive interference, has an important difference from a superficially similar case that we call subtractive interference. In this second case, an effect is reduced or even prevented by removing of one of the contributing powers. We could stop a match from lighting, for instance, by removing one of the factors that contributes to it. In theory, we could remove oxygen from the room and thereby the combustible power it contributes. The match then no longer lights. This is a case, therefore, where something is subtracted (figure 10) from the assemblage of powers. In both cases of interference—additive and subtractive—it may not be that either F or G is prevented from occurring but simply that the degree to which F or G occurs could be affected. A heated room may still get warmer if we turn on the heater, but an open window can limit the degree to which it warms. The term 'preventer' could then be used for a power of which the addition or removal stops something happening, some particular becoming G, for instance. And the term 'interferer' could be reserved for something the addition or removal of which affects merely the degree or extent to which something happens, such as the degree to which the particular becomes G.

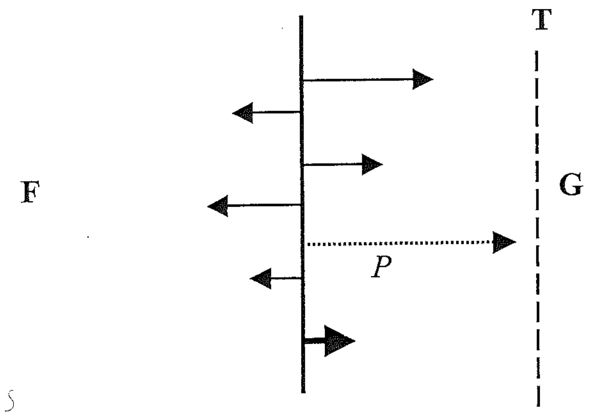

Figure 10: Subtractive interference

6. *Dispositional Modality*

The reason it is important to distinguish the two cases of additive and subtractive interference is that the former (but not the latter) case helps us to see the dispositional nature of causation: specifically that it does not involve any form of necessity or sufficiency. It might be thought by anti-Humeans that the way a cause produces its effect is by necessitating it; that is, where factors *a*, *b* and *c* cause an effect E, they must be causally sufficient for it. But we can see easily in the case of additive interference that this cannot be sufficiency in any standard sense in which philosophers use that term. Similarly, we cannot say that the cause necessitates the effect, as necessity is standardly used in philosophy. If *a*, *b* and *c* necessitated E, then whenever they occurred, E must also occur. C would then guarantee E. But, as shown in figure 8, we could indeed have *a*, *b* and *c*, which are capable of producing an effect E, but fail to do so if accompanied by the additional power *p*. So causes cannot necessitate their effects by guaranteeing the effect. What we have, therefore, is an antecedent strengthening argument against the necessitation of an effect by its cause. If C necessitated E, then C, plus anything added, should also produce E. But this cannot be the case for causation because of the clear possibility of additive interference.

This has consequences for the modality related to causation. Hume had argued that all causal connections were contingent and that no necessary connection is found between them. Those who believe in real powers have tended instead to claim that causes necessitate their effects. But we can see that this is an overreaction. Causation involves neither pure contingency nor pure necessity and the vector model shows us why. Causal powers do certainly dispose towards some outcome, so we have more than contingency and not everything is loose and separate. But dispositions do not necessitate any outcome either. What occurs will be dependent on the circumstances and what other powers are also operating. So we have less than pure necessity. What better name could we find for the modality involved in causation than 'dispositionality'? Causes dispose or tend towards their effects, even in the cases where they succeed in producing them. This is clear with general causal claims, as when we say that smoking causes cancer even if not everyone who smokes actually gets cancer. But also for singular causal claims we can see that it remains true. If some particular assemblage of powers caused a particular effect, they did not do so by necessitating it. Had some interferer also been present, it might have prevented that effect, or changed it. The cause disposed towards the effect, and indeed succeeded in producing it, but might not have done had a countervailing power been also present.

Most of this should already be understood pre-theoretically as the dispositional theory is not counterintuitive. Aquinas understood it (see Geach 1961, 101–102), though many post-Humean philosophers have forgotten it. But the causal reasoning we all employ is clearly nonmonotonic. We can allow that a match, plus striking, plus oxygen may produce flame, while also allowing that a match, plus striking, plus oxygen, plus high humidity would not tend to produce flame. A theory of causation ought to explain the nonmonotonic nature of causal reasoning, or at least be consistent with it. The dispositional theory explains it through the case of additive interference, clearly illustrated in the vector model in figure 9. And we can see that this antecedent strengthening consideration that counts against the necessitation of its manifestation by a power applies just as much to any resultant power as it does to the individual component powers. The resultant too only disposes towards an outcome for it too is subject to the possibility of additive interference.

We may of course think that, within the model, the final resultant power is necessitated by all its components. It is open to us to accept this because we see the resultant as nothing more than the combination of all the component powers. The components do not cause the resultant: they constitute it. But had there been, counterfactually, additional components then, of course, we could have had a different resultant. Necessity may then be an appearance of the model only. We can stipulate the model as a closed finite system of powers. Models can fail us when applied to the real world because we then have to deal with an open system. Additive preventers that were not assumed in the model can always interfere once we are dealing with a real-case scenario, whether in the lab or out in the world.

7. *Multi-dimensional Cases*

So far we have dealt only with the simplest cases, which concerned causation within a one-dimensional quality space with only two properties,

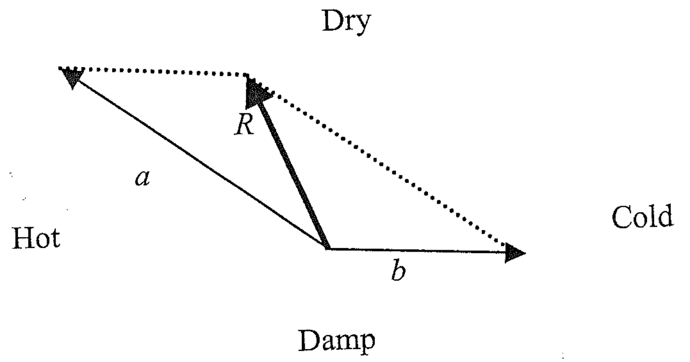

Figure 11: Powers within a two-dimensional quality space

F and G, at the extremes of the dimension. But we can also model more complex causal situations by use of multi-dimensional quality spaces. Suppose, for instance, that a room contains a heater that disposes it towards warm dry air but also an air conditioner that disposes it towards cool damp air. When both devices are switched on we can get a movement within a two-dimensional quality space, as shown in figure 11. A resultant vector can be calculated using the parallelogram law.

It is not clear that we need to be committed to irreducibly multi-dimensional quality spaces, however. Multi-dimensional cases may just break down into conjunctions of different one-dimensional quality spaces. Suppose something is made to become hot and red, for instance, from having been cold and black. We may be able to understand this as one set of powers operating on the cold-hot quality space and another set of powers operating on the black-red quality space. If we have distinct powers operating on distinct spaces then we can explain the causes entirely in terms of one-dimensional spaces and their conjunctions. A reason we might need multi-dimensional cases, however, is if there were to be a power that inseparably worked on two or more dimensions such that it could not decompose into two separate single-dimensionally oriented powers. We will see shortly that in cases of nonlinear and probabilistic cases there might be something like this. Our aim here is to show only that, if there were irreducibly multi-dimensional powers, they can still be modelled within a vector diagram. A three-dimensional case is modelled in figure 12 as a cube with 6 different properties. But in principle one could model a diagram with any number of properties by using multi-sided geometrical spaces such as a tetrahedron, icosahedron or dodecahedron. The most complex one would be a sphere with infinitely many points towards which the vectors can be directed.

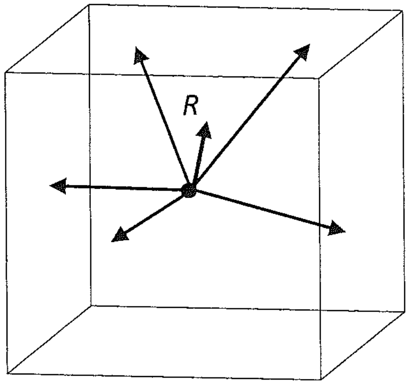

Figure 12: Powers within a three-dimensional quality space

8. *Probabilistic Causation*

It might be thought that there are genuine cases of indeterministic causation where, given a certain cause, there is an irreducible probability Pr of the effect occurring. If there are such irreducibly probabilistic cases of causation, then they can be also represented within the vector model. The simplest example is the toss of a coin. Assuming for the moment that this is a genuinely probabilistic case, it can be illustrated within our model. Suppose that it is a fair coin, which means that it is disposed equally towards both outcomes, head and tails. We take this to be a single disposition but one towards two possible outcomes with equal probability. Consequently, we represent this power as a single, double headed vector (figure 13). There is a reason for this. We should treat this power as a single unity because the probability of one outcome must constrain the probability of the others. The probabilities must add up to one, so we cannot treat them as independently variable. The best explanation is that the coin has a single, multi-directed power. The coin represented in figure 13 is a fair one but of course some coins might not be fair. A loaded coin would be one that disposed more in one direction than another. We can model this by moving our vector over towards one side. This illustrates that, in probabilistically constrained causation, the outcomes might not be equiprobable.

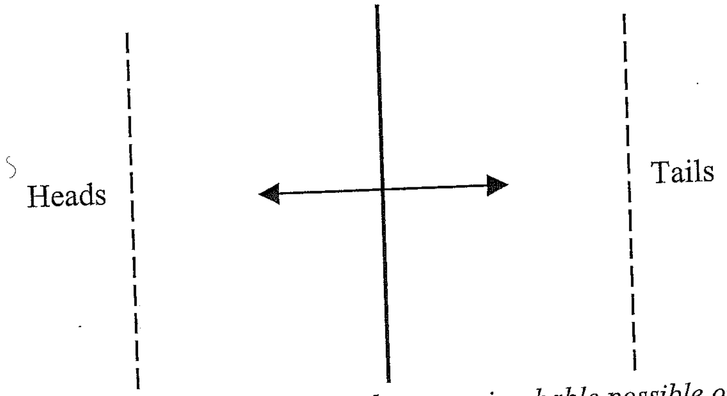

Figure 13: A probabilistic case with two equiprobable possible outcomes

Probabilistic causation need not have just two possible outcomes either, of course. The model of a dice, for instance, has six possible outcomes.

Now let us suppose that this is a genuine case of probabilistic causation. For the same reason given above, this has to be modelled as a single multi-directed vector in which each outcome, for a fair dice, is equiprobable. The faces of the cube used to represent this (figure 14) show the thresholds for the dice landing on each face, and the vector is directed one sixth towards each of them.

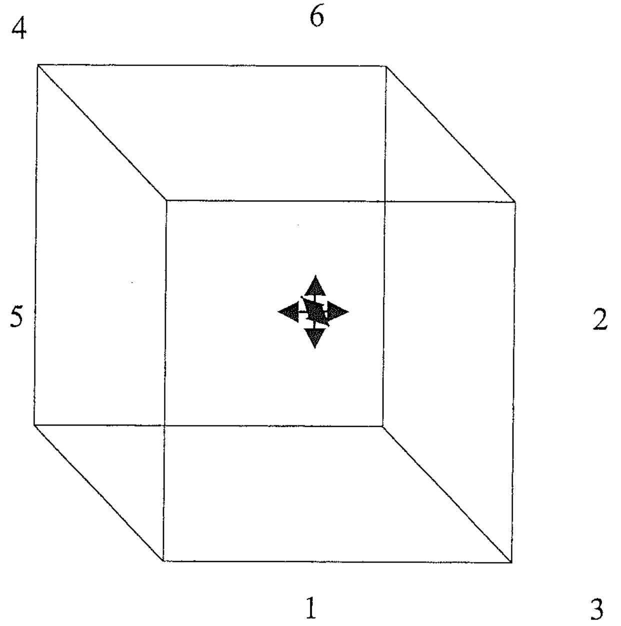

Figure 14: A probabilistic case with six possible outcomes

We would have some difficulty in predicting outcomes in any causal case in which at least one of the operating powers is irreducibly probabilistic. A resultant vector in such a case might reach a certain threshold T and yet, because of the probabilistic element in one of the causes, we have no guarantee that the outcome corresponding to the passing of T actually occurs. But this is only a specific case of what we have already said holds generally. The resultant vector only tells us overall to what the combined powers dispose. Where one of those powers is probabilistic, we can see even more that the resultant power *only* disposes towards an outcome. This

makes perfect sense within the calculus of probability for it is well-known that a probabilistic case is consistent with any outcome. A hundred coin tosses landing heads is at least a possibility, even when the coin is fair. But it is vastly improbable. A probabilistic disposition will tend towards a certain distribution of outcomes over a series of trials, but it never guarantees that distribution as an outcome.

9. *Causation by Absence*

Having dealt earlier with causation of absence, we ought also to offer a judgment of causation *by* absence, such as when lack of water is said to cause a plant's death or lack of insulin causes diabetes. Such a case offers a prima facie difficulty for the powers account, which is an attempt to explain all cases of causation in terms of the manifestation of powers. An absence is nothing at all, so how can it have any causal powers? How, then, can the death of a plant be the manifestation of some power of the absent water?

Fortunately, there is a simple solution to this for the dispositionalist, which is along the lines of that offered by Dowe (2001). The vector model makes it particularly easy to see. The case should be seen as one in which the plant has maintained an equilibrium state, which has allowed it to develop and grow. But it faces the twin perils of death by drowning and death by dehydration (figure 15). Water is to be added to the plant but only enough to allow it to maintain this equilibrium with a zero-resultant vector R. Too much, or too little, and it could slide out of equilibrium.

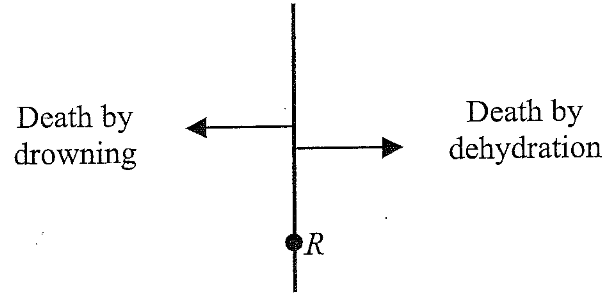

Figure 15: Plant moisture is kept in equilibrium

What then happens when I forget to water my plant? Assuming it dies, what caused it to do so? It is not the absent water, we maintain, whose powers have now been taken away. Rather, it is the remaining powers, as shown in figure 16. The equilibrium had involved the counterbalancing of four powers, let us assume. Of most interest to us is that the addition of water was counterbalancing the dehydrative power of the air around the plant, which sucks the water out. When the water is removed, we get a directed resultant vector towards death by dehydration. But it is the remaining powers that have caused this, and not the absent water.

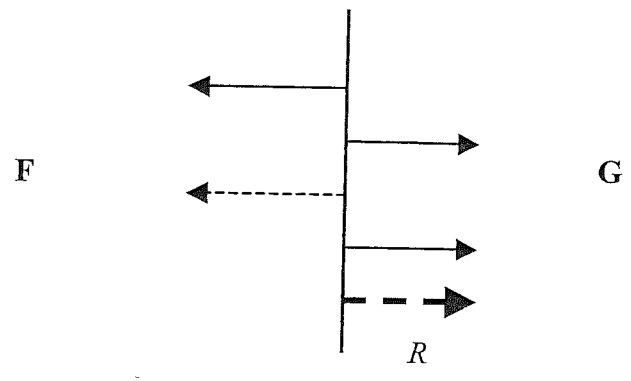

Figure 16: Causation 'by absence'

We say, therefore, that there is no genuine causation by absence. All effects are indeed caused by the present powers, not the absent ones. The powers theory, however, also explains why the absence of water is invoked in explanation of the plant's death. It is, after all, the absence of water we invoke in the death, rather than the absence of, for instance, methane. But the reason it is relevant to look to the absence of water is that it does have the power, in certain circumstances, to sustain a plant's life, which methane does not. This is also our account of counterfactuals: when water is absent, there is thus still a true counterfactual that had the plant received water, it could have lived. The power of water is thus the truthmaker for the counterfactual.

This does not mean that our account coincides with the counterfactual theory of causation. For one thing, our causal counterfactual is one that deploys only the dispositional modality: had a certain factor been present, it would have *disposed* towards a certain outcome. But for another thing, we have already noted that the counterfactual dependence

theory of causation works only if overdetermination is ruled out. If A and B both are enough to cause E, then E counterfactually depends on neither. We do not see that overdetermination should automatically be ruled out by a theory of causation, for there are some cases that seem intuitively plausible and require much work to dismiss. Why couldn't two separate factors each be adequate to produce E on their own but just happen to occur together? The powers account allows this possibility. Overdetermination can be represented in the vector model as in figure 17. There are other differences between the dispositional and the counterfactual accounts of causation, for there can also be counterfactual dependencies between events without causation, but this has been outlined elsewhere (Mumford and Anjum 2009).

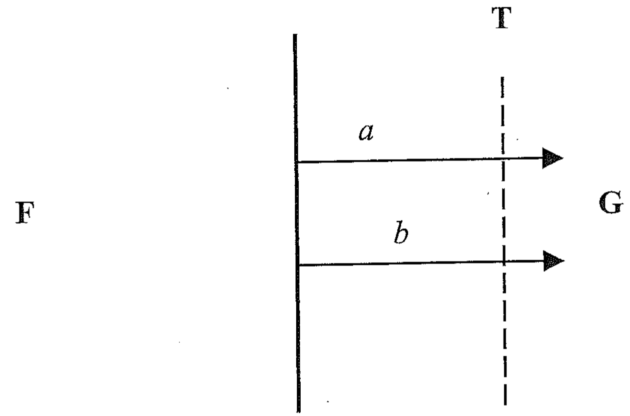

Figure 17: Overdetermination

10. Compositional Pluralism

As indicated earlier, the issue of how powers compose into a resultant power is a complicated one. The model provided so far works along the lines of vector addition but that is too restricted a model. The reason for this is that not all powers compose additively. A principle of additive composition would tell us that where D_1 is the power to M_1 and D_2 is the power to M_2, then D_1 plus D_2 combined would compose into nothing more nor less than the power to $M_1 + M_2$. But there are a number of cases in which we know that this does not occur. One example is taken from medicine. There are two drugs, clonidine and beta-blockers, both of which individually dispose towards the lowering of blood pressure. Surprisingly, however, if they are taken in combination, they tend to raise blood pressure in a significant number of cases (Warren, et. al. 1979). This is a rather extreme antipathetic case. Most drugs that dispose towards the same result would work in sympathy if taken together but here we find that, not only do they not compose additively; taken together their resultant vector R actually disposes in the *opposite* direction (see figure 18).

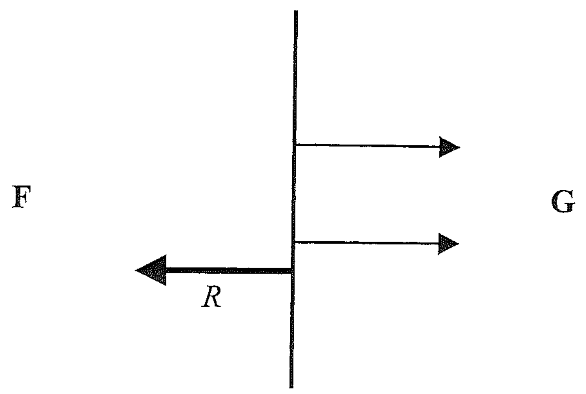

Figure 18: Two powers that jointly dispose away from their separate effects

The problem with the principle of additive composition is that it assumes that composition works on linear principles, but we have enough evidence to suggest that many causal connections work nonlinearly. A chocolate bar disposes towards pleasure when eaten, for instances, but ten chocolate bars eaten together do not produce ten times the pleasure of one. Each additional chocolate bar produces less pleasure. Similarly, the connection between wealth and happiness might also be nonlinear. In a linear function, the extent of the output is directly proportional to the extent of the input such that we would get a straight line when plotted onto a graph. Wealth may well be causally relevant to happiness but in a nonlinear way, hence the graph could for instance be a curve (figure 19). For nonlinear composition, the extent of the output would not be proportional to the extent of the inputs.

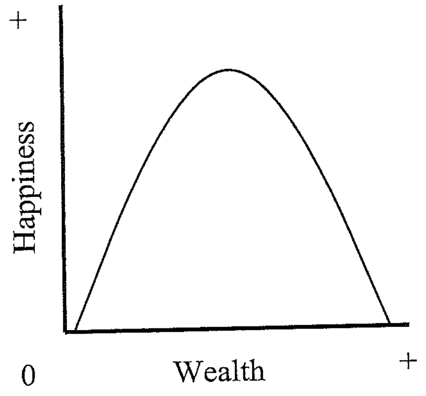

Figure 19: A nonlinear causal connection

A further kind of case worthy of note is the case of chaotic systems, as the weather is supposed to be. Among other features, chaotic systems are notable for their hypersensitivity. This means that we can make the tiniest of changes to the cause and it can result in a massive difference in the effect, as exemplified in the so-called butterfly effect. In the case of the vector model, therefore, we might be able to add one very tiny extra power to a system with chaotic causal behaviour and get an enormous difference in the resultant vector.

There are various responses that we could give to these kinds of cases but the simplest and most accurate thing to say is that we should allow many different functions that describe the composition of powers. Addition is just one function for the composition of powers and suffices for any linear case. But there is no reason why we should not allow many other functions that describe how powers interact. Addition is simple and will no doubt appeal to Humean sensibilities. Composition might be thought by a Humean, for instance, to consist just in mereological relations. But the dispositionalist has reason to doubt this picture in any case for they think the world is complex and integrated with genuine interactions of the causal powers at work. This view is called, therefore, compositional pluralism for it allows a plurality of ways in which powers compose. This, it will be recalled, is a reason why we might want to allow multi-dimensional powers and quality spaces. We may have two powers relating to different quality spaces but such that, when they act together, they interact in a nonlinear way that does not decompose into the individual actions of those powers had they acted alone.

This sacrifices some of the simplicity of the vector model but at the gain of accuracy. It would be nice and easy if all composition were additive, for then we could just use vector addition to predict an outcome. But we know we cannot. Powers can still nevertheless be a basis for prediction, if they and their interactions are understood. To understand certain nonlinear causal interactions, therefore, it is not enough that we know and understand the individual powers that are involved. That would be too simplistic and suggest a reductionist view of the physical world. We need also to understand how those powers relate and interact. This is to understand causal processes holistically and there is no reason why we cannot do so. The functions that describe the nonlinear causal interactions are, after all, as empirically accessible as anything else to do with powers.

11. Conclusion

We began with the claim that neuron diagrams were suited to a Humean picture of causation rather than a dispositionalist view. As an alternative, we have shown how a dispositionalist can instead represent powers as vectors. This would allow us to show how causation is complex, how the extent of a power makes a difference, and how countervailing powers make a difference. Coming out of this, we have seen that the powers view can account for causation of absence, causation by absence, and overdetermination, and we have shown how the model can be extended to multi-dimensional, probabilistic and nonlinear cases.

The model offers support to the idea of causation as involving a non-Humean *sui generis* dispositional modality. Hume wrong-footed his opponents when he claimed that the believer in powers accepted necessary connections in nature. Our vector model shows us why a power doesn't necessitate its manifestation, even in a case where it succeeds in producing it. Causal production does not, therefore, mean the same as causal necessitation. If we accept this view, then powers return to the centre of our philosophy of nature. We can say 'return' because this is after all an Aristotelian and Aquinian view in which a cause is seen as something tending towards an effect. The metaphysics that best explains this, we maintain, is one of real causal powers. We have not, in this paper, offered any direct

support of that metaphysics but we hope to have shown some of the use to which is could be put.[1]

Stephen Mumford

University of Nottingham

Rani Lill Anjum

UMB—Norwegian University of Life Sciences

NOTE

[1] We thank audiences at Granada, Hertfordshire, Oxford, Aarhus, Saskatoon, Lausanne, Köln, Buffalo, Tromsø, Warsaw, Athens, Oslo, and Melbourne where parts of this material was presented. We thank in particular the Nottingham Dispositions Group—Charlotte Matheson, Markus Schrenk, Matthew Tugby—who have commented on many drafts of this material over the past couple of years. We are grateful to Jonathan Tallant for assistance with the title.

REFERENCES

Bhaskar, R. 1975. *A Realist Theory of Science*, Leeds: Leeds Books.

Cartwright, N. 1983. *How the Laws of Physics Lie*, Oxford: Oxford University Press.

———. 1989. *Nature's Capacities and Their Measurement*, Oxford: Oxford University Press.

Dowe, P. 2001. "A Counterfactual Theory of Prevention and 'Causation' by Absence," *Australasian Journal of Philosophy* 79: 216–26.

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

Geach, P.T. 1961. "Aquinas," in G.E.M. Anscombe and P.T. Geach, *Three Philosophers*, Oxford: Blackwell, 56–125.

Harré, R. and E.H. Madden, 1975. *Causal Powers: A Theory of Natural Necessity*, Oxford: Blackwell.

Hitchcock, C. 2006. "What's Wrong with Neuron Diagrams?" in J.K. Campbell, M. O'Rourke and H. Silverstein (eds), *Causation and Explanation*, Cambridge, MA: MIT Press, 69–92.

———. 2009. "Causal Modelling," in H. Beebee, C. Hitchcock and P. Menzies (eds), *The Oxford Handbook of Causation*, Oxford: Oxford University Press, 299–314.

Lewis, D. 1973. "Causation," *Journal of Philosophy* 70: 556–67.

Lombard, L. 1986. *Events: A Metaphysical Study*, London: Routledge and Kegan Paul.

Manley, D. and R. Wasserman 2007. "A Gradable Approach to Dispositions," *The Philosophical Quarterly* 57: 68–75.

Martin, C.B. 2008. *The Mind in Nature*, Oxford, Oxford University Press.
Mill, J.S. 1843. *A System of Logic*, London: Parker.
Molnar, G. 2003. *Powers: A Study in Metaphysics*, S. Mumford (ed.), Oxford: Oxford University Press.
Mumford, S. 2008. "Powers, Dispositions, Properties, Or: a Causal Realist Manifesto," in R. Groff (ed.) *Revitalizing Causality: Realism about Causality in Philosophy and Social Science*, London: Routledge, 139–51.
Mumford, S. and Anjum, R.L. 2009. "Double Prevention and Powers," *Journal of Critical Realism* 8, 277–93.
———. 2010. "A Powerful Theory of Causation," in A. Marmadoro (ed.), *The Metaphysics of Powers: their Groundings and their Realizations*, London: Routledge.
———. 2011. "Dispositional Modality," in C. F. Gethmann (ed.), *Lebenswelt und Wissenschaft, Deutsches Jahrbuch Philosophie* 2, Hamburg, Meines Verlag, 380–94.
Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.
Shoemaker, S. 1980. "Causality and Properties," in *Identity Cause and Mind*, expanded ed., Oxford: Oxford University Press 2003, 206–33.
Warren, S.E., E. Ebert, A.-H. Swerdlin, S.M. Steinberg, and R. Stone, 1979. "Clonidine and Propranolol Paradoxical Hypertension," *Archives of Internal Medicine* 139, 253.

POWERFUL QUALITIES, NOT PURE POWERS

1. Introduction

According to the neo-Humean metaphysic, defended most prominently by David Lewis,[1] the world is a vast collection of particular, local matters of fact—it's just one little thing after another. Necessary connections between distinct existences are, on this view, anathema. "[A]nything can coexist with anything else . . . Likewise, anything can fail to coexist with anything else" (Lewis 1986, 88). The world is a sort of mosaic of facts, and the connections between them—nomic, causal, or modal—supervene on the patterns in that mosaic. Properties, the colors of the mosaic, are purely qualitative. In themselves they are impotent, devoid of any intrinsic nomic, causal, or modal character. They get connected to other properties only by way of the laws of nature, which are themselves contingent patterns in the mosaic of qualities. Call the resulting view of properties categoricalism.

To contemporary essentialists, the happenings in a neo-Humean world are radically contingent in a way that threatens to make the world wholly unintelligible. In such a world, events lack any genuine unity. It is tempting to identify categoricalism about properties as the culprit. But what competing theory of properties does the essentialist have to offer? The available alternatives are surprisingly underdeveloped. Properties are, according to the essentialist, intrinsically powerful, packing their own nomic, causal, and modal character. But this is far from an account of properties. What *are* properties?

I explore two accounts. The first is the pure powers view, defended by Alexander Bird and Stephen Mumford,[2] on which properties are powers and nothing but powers. The second view is the powerful qualities view, defended by C.B. Martin and John Heil,[3] on which properties are *both* powerful and qualitative. Both views are, unfortunately, easily misunderstood. In section 2, I clarify the pure powers view, and raise worries for it. In section

"Powerful Qualities, Not Pure Powers" by Jonathan D. Jacobs,
The Monist, vol. 94, no. 1, pp. 81–102. Copyright © 2011, THE MONIST, Peru, Illinois 61354.