Norwegian University of Life Sciences
School of Economics and Business

Powers, Probability and Statistics

Stephen Mumford and Rani Lill Anjum

Working Papers No. 10 / 2016

ISSN: 2464-1561

Powers, Probability and Statistics

Stephen Mumford

Professor of Metaphysics, Department of Philosophy, University of Nottingham, Nottingham, UK and Professor II, Department of Economics and Resource Management, Norwegian University of Life Sciences (NMBU), Ås, Norway.

Email: Stephen.Mumford@nottingham.ac.uk

and

Rani Lill Anjum

Research Fellow, Department of Economics and Resource Management, Norwegian University of Life Sciences (NMBU), Ås, Norway.

Email: rani.anjum@umb.no

Abstract

Regularity is often taken as the starting point of our causal knowledge. But pure constant conjunctions are not what science finds. Even in randomised controlled trials, we do not discover a regular frequency of occurrence of some effect. The dispositionalist is able to explain the evidence of science in terms of the ontology of real causal powers exhibiting an irreducibly tendential nature: less than necessity but more than pure contingency. Much evidence of this kind has to be understood probabilistically and there is a frequentist interpretation of in which the facts of frequency of occurrence fix all the facts of probability. However, the dispositionalist has a stronger propensity interpretation of probability at their disposal in which the facts of probability are determined by the individual powers of things. The dispositional approach allows us to make sense of large-scale population data in which different individuals within the same sub-group can have different probabilities of being affected by a cause. On this view, individual propensities can compose to make an overall chance of an effect for a group. But from a starting point of general facts of probability for groups, we cannot decompose those chances back to individuals.

KEYWORDS: causation, powers, probability, propensity, frequency, data, statistics, randomised controlled trials, medicine

1. Causation and regularity

David Hume famously connected the notion of cause with regularity (Hume 1739: I, iii, 2, p. 77). His reason for doing so was his rejection of singularism. When he looked at a single instance, all he could see was one 'object' or event followed by another. He couldn't detect that the first event necessitated another or that, without the first, the second would not have happened either. We would never arrive at the idea of the first event causing the second, he concluded, without the experience of repetition. If there are many instances of the first type of event and they are all followed by instances of the second type of event, then we might start to form the idea of the first being a cause of the other. We have experience of many cases of the eating of apples, for example, and they have all been followed by the nourishing of the body.

From this accumulation of repeated instances, we form the idea that apples nourish or, to use explicit causal vocabulary, apples cause nourishment. Crucially, however, we would never have arrived at this conclusion if our experience was limited to just one instance of an apple being eaten, followed by the body being nourished. And nor would we have drawn this conclusion just from the inspection of the apple, for reason alone could not show us that the properties of being round, smooth and green would be followed by nourishment. The only way we acquire the idea of causation, Hume concludes, is regular succession or what is known as constant conjunction.

Regularity has remained deeply associated with our causal thinking. The idea is then that where one factor A causes another B, it should do so on all occasions. Both Humeans and anti-Humeans seem to an extent to agree on this. Humeans believe there is nothing more to causation over and above that pattern of succession of events. If A is always followed by B, then nothing more is required for it to be a fact that A causes B, except for the two other conditions Hume imposes: that A occurs before B and they are spatially contiguous (for a defence of the Humean approach, see Psillos 2002). Anti-Humeans tend to think that these conditions are not enough and that there has to be some necessity that produces the regularity (see for instance Ellis 2001). But, either way, there is agreement that regularity is a part of causation.

This is a view we will challenge, however. The connection between causation and regularity is not quite so clear cut if one accepts the insights of a dispositional theory of causation based on a metaphysics of real powers (Mumford and Anjum 2011a). We will argue that a world of tendencies requires a reappraisal of the traditional view for it opens the possibility of causation without constant conjunction. The world of tendencies may still exhibit a degree of regularity, however: enough for us to get by. But, as the notion of degree suggests, this regularity will be less than perfect. We will argue that this is entirely consistent with typical results when we investigate the world scientifically. Setting aside the debates around whether such a view is metaphysically adequate or not, we will be claiming that such a metaphysics is capable of producing a satisfactory epistemology. The way that we have come to understand correlation in the discovery of causes, for instance, is one that allows for the tendential nature of causation. And where we have irreducible probabilities, these are to be understood in terms of propensities rather than relative frequencies. Importantly, however, such tendencies must retain their irreducibly dispositional nature and should not resolve into necessities, which would again require constant conjunction in causes and effects.

2. Tendencies

The metaphysics of tendencies offers a third option between the traditional Humean and anti-Humean options. We can think of the Humean view as offering us a world of complete contingency in which anything could follow anything else, a view articulated in its most sophisticated modern form by David Lewis (1986). Hume himself characterised his opponent's position as one requiring necessary connections in nature (Hume 1739: 161) and anti-Humeans have indeed endorsed the view given to them (Harré and Madden 1975, Ellis 2001, Bird 2007). Dispositionalists have traditionally thought of necessity as the way in which causes produce their effects and such necessity is certainly anti-Humean. They posit exactly the thing that Hume claimed we couldn't see and had no idea from which to legitimately form the notion.

A middle way is possible, however. Geach (1961: 102) attributed to Aquinas a belief in powers having an irreducibly tendential nature. Powers would dispose towards their effects in a way that was stronger than pure contingency but less than full-strength necessity. This idea was mentioned by Harré and Madden but their book was still subtitled *A Theory of Natural Necessity*. Similarly, Bhaskar spoke of tendencies but it looks like he understood these as providing the world with necessity (Bhaskar 1975: 14). In short, where something has a tendency and it is in the right conditions, it *must* manifest that tendency (Bhaskar 1975: 214).

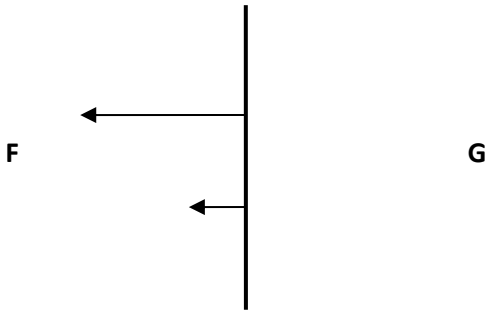

Figure 1: vectors representing two powers with different magnitudes

In recent work, we have invoked a tendential theory that is closer to the one attributed by Geach to Aquinas (Mumford and Anjum 2011a: ch. 8 and 2011b). A key difference between this account and Bhaskar's, for example, is that we assert that when a tendency is in ideal conditions for its manifestation, it still only tends to produce it. The tendential nature of dispositionality is ineliminable. Now such a tendency may be a very strong one. Tendencies come in degrees, just as a wineglass and a car windscreen are both fragile but the former to a higher degree than the latter (fig. 1). Some of those tendencies, especially in ideal conditions, will almost always manifest. Almost always. But, if this is a thorough-going dispositionalism, we should not reduce the notion of tendency to that of necessity, as Bhaskar, and many other dispositionalists, seem to do. If dispositions are understood as necessities, in this sense they are merely playing the same role for which we previously invoked laws of nature. Irreducible tendencies, with a sui generis dispositional

modality, allow for the possibility of a real conceptual advance: a challenge to the paradigm of what we can call modal dualism. Philosophy has had us deal only with pure contingency and necessity but neither, we assert, explain the modal nature of causation adequately.

What is the argument for such a view? We argue (Mumford and Anjum 2011a: ch.3) that powers produce their effects without necessitating them. The argument against necessity comes from the possibility of causal prevention: specifically the variety we call additive interference. A set Σ of causal factors that are mutual manifestation partners for some effect E may on many occasions succeed in producing E. But on at least some occasions, Σ can be accompanied by some φ where φ is capable of preventing E even when Σ. As necessity is understood in philosophy, if A necessitates B, then whenever A is the case, B is the case irrespective of what else occurs with A. With natural causal processes, however, it looks prima facie that additional factors can indeed prevent that effect happening. The modality of causation – what should be thought of as the natural modality – is therefore not necessity as traditionally conceived.

The validity of this test of necessity has been challenged by Lowe (2012). His concern is that this is not in general a way to test the necessity of a conditional. We suggest that it should be possible to add anything to the antecedent and it still come out as true, if the conditional concerns necessity. But then one might add to the antecedent a condition that contradicts what else is in the antecedent, rendering such a conditional still true (trivially), and thereby passing the necessity test whether it expresses necessity or not. But what we propose is not a test of the necessity of conditionals in general but a test of the necessity of causal claims when they are posed in conditional form. And here, it seems, we needn't have Lowe's concern. The reason is that such a conditional could never be trivially true on the grounds Lowe suggests because a contradiction within the antecedent is not something that can exist in nature. Hence, there is no such antecedent that puts the necessity of the claim to the test. All such antecedents that we are testing are ones that can naturally occur and are thus contradiction free.

We thus will retain our commitment to the necessity test for causation. This is not to say that there cannot be other tests of necessity. Perhaps some of those tests are better than ours. But, if so, let the opponents of the dispositional modality propose them and show that causation can reasonably be thought to pass that test. What we wish to warn against is the mere assumption without argument that causes necessitate their effects and, correspondingly for Humeans, the mere assumption that constant conjunction is a necessary condition for causation.

3. Tendential strength and probability

We said that a power can tend towards its manifestation to a greater or lesser degree. A power thus has a strength and this strength is something that can be gauged statistically in frequency of occurrence. We might observe, for instance, that a statistical majority of those who smoke die of a smoking-related disease. A study estimates that of the average 443,000 smoking-attributable deaths per year in the USA, 29% die of lung cancer, 28% of heart disease, 8% of other types of cancer, 4% of stroke and 21% from chronic obstructive lung disease (CDC 2008). From evidence of this kind, we might reasonably conclude that smoking causes all these diseases even though there is no constant

conjunction involved. The tendency towards cancer is a relatively significant one. In the cases of stroke, the tendency is not as strong in a proportional sense but given the seriousness of the effect, even this tendency of smoking can be considered a serious one. There is a stronger tendency towards lung diseases generally but even here it is far short of necessity. The tendential approach shows, therefore, that a smoker lucky enough to avoid all these illnesses in no way proves that smoking does not cause them. We cannot therefore be naïve falsificationists about causal claims.

Should we then treat these, and all other cases, as examples of probabilistic causation? Is the dispositional-tendency view one that simply renders all causation as probabilistic? Such a view, which is one reason why some before us have doubted that regularity is a necessary condition for causation (Armstrong 1983: 29-35), seems to fit prima facie the features of a dispositional account. Causation would indeed involve something stronger than mere contingency and weaker than necessitation. All the tendencies we get from dispositions would be probabilistic, with the strength of the probability matching the magnitude of the disposition.

However, it would be a mistake to immediately infer the ubiquity of probabilistic causation from its dispositional nature. For one thing, we do not see probabilities as always an irreducible feature of reality and there are many cases where such facts of objective chance are produced by the tendencies of the powers and their strengths. Causes may well be probability-raisers, in almost all cases, but we take it that there is a further question of why they raise the chances of their effects and our answer is that they dispose or tend towards them. This is to say that we do not want to explain tendencies in terms of probabilities because we instead explain probabilities in terms of tendencies.

If tendencies are to be the worldly grounds of the truths of probability, then they cannot be probabilistic in nature themselves. But we can see that they are not, on a true tendency account. The strength of a power is not simply the same as the probability of its effect. We measure and state probabilities on a bounded scale between 0 and 1, where there cannot be a higher probability of e that =1 or a lower probability than =0. The strength of a power, however, for at least some cases, is measured on an unbounded scale. This means that there can be more than enough to produce a certain effect. Consider the use of a brick as a paperweight, for example, which should always be more than enough power to stop the papers blowing away. And yet it is still possible to double the extent of the power by adding another brick on top of the first. While the extent of power is thus unbounded, within the theory of tendencies, it is still always the case that the classical probability of an effect is <1, because no effect is necessitated. Hence, we can have as many safety mechanisms in an aeroplane as we like, and we can even back them all up, but we still know that the probability of safely completing the flight is very slightly less than 1. If we think of a propensity as the having of power or tendency to some degree, then, it is clear that it is not the same as a classical probability. That is good news if you want propensities to be the ontological grounds of probability (Mumford and Anjum 2016, forthcoming a)

Nevertheless, as well as the regular, non-probabilistic powers that tend to some degree to manifest, there are also probabilistically constrained powers that involve an objective chance of some outcome or other being realised. To understand the exact nature and usefulness of this distinction, we need to consider different ways in which statistical data can be treated as probabilistic.

4. Interpreting statistical data

It is quite natural and to a degree justified to look for causes using statistical methods. However, we think that the right methods have to be used and their limitations acknowledged. It ought also to be accepted that there are lessons to be learnt from the sort of data that are produced in most scientific trials.

For a start, we should not expect that the pure constant conjunctions, that Hume thought the starting point of causation, are going to be found in scientific trails such as randomised controlled trials (RCTs). That would require that, upon some intervention C, an effect of type E always followed. Even the most 'successful' RCTs will not deliver anything close to that (Howick 2011, Cartwright and Hardie 2012). And we should bear in mind that in the *Enquiry* Hume also invokes the idea that as well as A being followed by B, we should be able to say that without A, B would not follow (Hume 1748, Section VII). Hume is sometimes taken as offering two completely opposed theories of causation here – a constant conjunction theory and a difference-making theory – and mistakenly thinking them to be the same. That would seem quite a blatant error to ascribe to Hume, given his evident power of subtle argumentation employed elsewhere.

Instead, we might think of Hume, having already acknowledged that the idea of cause is complex, as seeing that the requisite regularity took two forms. Certainly A should be followed by B, but that also would not be enough for our idea of cause because if B was ubiquitous – and therefore followed every event – its following of A would not be adequate for us to form the idea that A caused B. What we would need, as well as every A being followed by B, would be the regularity that without A there is no B. Constant conjunction and difference-making elements of regularity make philosophical sense, taken together, therefore, as the signs of causation.

But RCTs, while testing for both elements, do not reveal what Hume is looking for in any absolute way. A typical RCT (see, for example, Broadbent 2013: 4-5 and Howick 2011: 45), we will suggest, may require division of a population into two or more groups: a treatment group that receives the intervention whose effects we wish to test, and at least one control group that receives nothing or a placebo. This type of RCT might be suitable for testing the effect of a new drug, for instance. Suppose the trial drug is hoped to aid recovery from some disease D. Naturally, we will invoke some of the standard requirements on RCT (Rabins 2013: 149). The size of the population must be large enough that we will consider any resultant data significant, and the division into the two groups must be genuinely random, so there are just as likely to be as many men in each group, as many smokers, as many optimists, as many depressives, and so on.

If the rate of recovery is higher in the treatment group than the recovery rate in the control group(s), this can be thought of as confirming Hume's second requirement, that without the cause, the effect would not exist either. The control group reveals the difference made by the drug. This is vital, since if the recovery rate of the control group were the same as the treatment group, no such effect has been demonstrated.

Say we find that 40% people in the treatment group recover from D and this is greater than the recovery rate in the placebo group. This gives us some indication of how effective the drug is. But there is no regularity, in the constant conjunction sense, between taking the drug and recovery. Statistical regularity here means regularity over a number of instances where each instance is a full trial. Such regularities might reveal only relatively weak correlation, for instance that 1 in 1,000 women taking a type of contraceptive pill develop thrombosis, and this is a two- to threefold higher risk than in woman who take no such pill (Rosing et al. 2001: 194). Although this is a very weak correlation, it could be a regular one, found in a number of trials; and given the undesirability of thrombosis, it would then be considered a significant correlation despite its weakness. People might risk a 1 in 1,000 chance of getting a headache but not the same chance of death, for example.

Even here, however, in this proportional form, we do not find the constant conjunctions quite imagined by Hume. It is unlikely that upon systematic review, we find that every trial of the same treatment reveals exactly the same recovery rate. One thing we could do is simply add together all the trial results and calculate the proportions by averaging. This might not be methodologically entirely sound, however, as not all trials will have employed exactly the study design. Perhaps some of the trials had a treatment group and placebo group but no third control group, for instance. The systematic review would then have to be slightly less strict, therefore, and willing to overlook differences in study design, between the different trials if they are considered minor enough.

The best Humean regularities that are found in statistics, therefore, are most typically a pattern of raised incidence of E upon C, showing a significant enough 'effect', within a range of variance. Epidemiologists may be trained not to assert a causal connection between C and E even in these cases. They are scientists rather than metaphysicians, after all, and the extent of the scientific claim may stop with there being a correlation between C and E. Another pattern that could also be sought, in more sophisticated trails, could be that one factor, F, correlates with another, G. This means that it could be found, for instance, that increasing the extent of F, through appropriate testing, is found to correlate with an increased extent of G, where F and G are quantities able to vary in magnitude. And a lower extent of F could correlate with a lower extent of G.

Those more philosophically minded might think it reasonable to draw an inference from such statistical evidence that there appears to be a causal *influence* of C upon E or of F upon G (see Lewis 2004 for one account that exploits this feature). But given the limited nature of the regularity, as we have described it, there are still some different interpretations that are possible in respect of probability assessments.

In the case mentioned, there is a recovery rate in the treatment group around 40%, in 4 out of 10. But because this is a study over a population – which is necessary for the results to be considered statistically significant – this is a consistent with a number of different possibilities concerning the effects on the individuals within the treatment group. One possibility is that the drug has no effect whatsoever on 60% of the group and a strong and successful effect on the 40% who show recovery. This interpretation would fit well with the idea that causation is a matter of all or nothing. Either the effect occurs; or it doesn't, in which case no causation has happened.

Using the statistical information to guide a clinical decisions, therefore, one could infer that there is a 0.4 chance of recovery with this treatment. It is important to note, however, that this probability

assignment is a matter of epistemology, evidence and risk; not ontology. There is a 0.4 chance that we are right in assuming that the treatment will cure D for this patient; that is, this is how likely it is that the patient belongs to the right sub-group. This probability assignment is therefore about our confidence, or credence, in the predicted outcome.

An alternative is to say that causation itself is probabilistic, and that the statistics reveals that the drug raises the probability of recovering from D to 0.4. This is an ontological form of probability concerning the way some restricted part of the world is. Next, we will look at two ways to interpret statistical data probabilistically in the ontological sense: those ways being frequentism and propensity theory.

5. Frequentism

One key decision we have to make in interpreting statistical data is whether we should read it as frequentists or propensity theorists. On a frequentist interpretation of probability, the probability of an event is determined by its relative frequency over time. This frequency is found through a large number of repetitions under similar conditions. David Lewis (1994) spoke of naïve frequentism, which would simply be a reading off of probabilities from the distribution of events in known experiments, or frequency of occurrence whether known or unknown. There are more sophisticated versions of frequentism that can be followed within a broader metaphysical framework, however (see Gillies 2000: ch. 5 and Eagle (ed.) 2011: Pt V). If one is a four-dimensionalist, one would think that there is a frequency of distribution over events of a certain type for all times. In the entire history of the universe – past, present and future – for instance, it may be that of those things that are F, 59% of them are also G. Given this omnitemporal fact, then we should say that the chance of any F, picked at random, is G is 59%.

Needless to say, within science we do not have access to every F that is, was and will be. Nevertheless, part of the rationale for RCTs might be that if one chooses test subjects in a properly random way, the sample of those things that are F and are chosen for the RCT ought to be more or less representative of the group of all Fs as a whole. Of course, one may get an unrepresentative sample but that would be the reason for repeating a number of studies and then reviewing them systematically. If a pattern emerges over a number of trials, then it can be thought of as indicating the approximate proportion of Fs followed by Gs.

The distinctively philosophical commitment of frequentism is that probability is constituted by the distributions of outcomes. For the naïve frequentist, this might simply be the distribution over a sequence of trials; for the more sophisticated version it would include every distribution of Gs within F. This means that probability is generated from statistics, whether known or all possible. Statistics reveals probability for the frequentist because probability is nothing more than the frequency with which the outcome occurs.

We can grant that frequencies are indeed a way in which we can come to learn about probabilities. We might see that a frequency with which a die lands 6-up is more than 1 in 6 times, for instance, over a reasonably large number of trials. This could lead us to conclude that it is a loaded die, of

course, but also something about the degree to which it is loaded. Perhaps after 100 rolls, there have been 50 sixes, and then we conclude that the chance of it landing 6 is 50:50. Here, one simply records the outcomes, calculates relative frequency and draws a conclusion about probability. One has begun from a position of ignorance but this might be just like our initial ignorance over the chance of a drug causing recovery. And for the frequentist, there is no further metaphysical fact grounding probability other than such distributions.

Frequentism faces challenges, however. Some of them concern whether facts of chance can be fixed when there is only a low number of occurrences of a phenomenon (Mumford 2004: 47f). Another problem is that the probability generated from frequency can only tell us something about a statistical average of a sample, which might not be representative for any of the individual instances from which the sample is gathered. Indeed, frequentism is too blunt to distinguish some of the different possibilities there could be when on average 40% of Fs are Gs, such as: all Fs having a 40% chance of being G, half of them having a 50% chance of being G and half having a 30% chance, or each having a separate chance within a range of 0% to 100%. Prima facie, it looks like the theory has no way of distinguishing these conceptually and ontologically distinct possibilities.

Now it might be alleged that the frequentist can make the requisite distinction by being adequately fine-grained about group membership. It might be said that of those things that are F, they divide into two groups, A and B, and frequency of being G in group A is 50%, while frequency of being G in group B is 30%, and again such facts of frequency are all the facts of probability. But this is not a move likely to lead to succeed. For a start, it seems arbitrary where to stop. Every individual could belong to ever more fine-grained sub-groups. And if one resolves the group down to the one that most fits an individual person, it is the N-of-1 group of which this individual is the only member. Probability claims are redundant in N-of-1, within the frequentist framework, given that this individual is simply either G or not. But such a terminus seems avoidable only if one takes some notion of a smallest relevant sub-grouping of F, and that seems an arbitrary matter.

If one avoids going the N-of-1 route, another objection seems to stand that one individual can belong to many different groups and sub-groups at the same time. The facts of possibility do not then looked fixed as an ontological fact but are, instead, merely relevant to our interests. The same individual may have a 40% chance of recovery from a disease *qua* man, but only a 20% chance of recovery *qua* male smoker, and yet a 30% chance of recovery *qua* male vegetarian smoker. Clearly, the same person can belong to a number of different sub-groups with the chance of recovery being different according to how finely grained one considers their grouping (it is the job of actuaries for life assurance companies to calculate chances based on all such factors).

6. Propensities

An alternative to frequentism is propensity theory. We have said that the ontological commitment of the dispositionalist is towards irreducible tendencies rather than regularities. What this will mean is that within any sub-group, no matter how small, there will be a tendency, and no more than a tendency, towards exhibiting some effect. It is important for the anti-Humean to say this. If they did not, then the Humean could easily claim that there is after all a pure, constant conjunction, of the

kind they originally claimed: but it was just that the constant conjunction got masked in large, varied populations with multiple factors at work. While only 40% of the treatment group recovered in our RCT, for instance, the Humean could allege that if a person received the treatment and satisfied a number of other conditions – perhaps they are non-smoking vegetarians who exercise – then they always recover. Here is the pure, Humean regularity. But this is what we deny as dispositionalists.

Now in a relatively small sample, it may be that all those who have this specific set of factors indeed recover. And the tendency towards recovery when all these factors are present may indeed be a very strong one. But perhaps if there were 10,000 people subject to all these causal influences, a handful of them would not respond in the same way. Even a strong tendency could fail to manifest itself, given enough trails. The most likely situation going back to our original recovery rate of 40%, as far as a dispositionalist is concerned, is that the varied population within the treatment group will have differing tendencies towards recovery. Some of the population will have a strong tendency to recovery, some will have little or none, and the statistic of 40% is what that whole bunch of different degrees of tendency produce taken together.

A consequence of this is that if the treatment in question becomes a marketed drug, it cannot be inferred with certainty from the outcome of the RCT exactly what effect the treatment will have on an individual considering using it. There is perhaps a 40% chance of recovery for the individual qua normal human being. But it is also quite likely that an individual has a number of other causal factors in operation that perhaps mean that they really have only a 20% chance of recovery. The taker of the drug has no way of knowing whether they are a statistically average kind of person. It is likely that most will deviate from the statistical average to some degree so it is a matter only of how much so. And perhaps there is a point at which a user is so atypical that the drug has little or no effect on them, or is even harmful.

On a dispositionalist view, we understand probability to be a fact of objective chance, distinct from the issue of credence or degree of belief. And what this fact of chance resides in is the individual propensities of things: their dispositions or tending towards certain outcomes to some degree. Such probabilities can be ascribed intrinsically to individuals: such as what the chance is of this candle melting. But they can also be ascribed to whole situations, for instance when we consider the chance of a candle melting that is stood next to a furnace that is just about to be lit.

The distinction between these two types of ascription of probability is important and we should not assume that every attribution of chance is intended in the same way. There might be a probability of a woman getting breast cancer based purely on demography, ultimately based on the sort of evidence we have discussed from trial data. But perhaps given certain environmental factors to which one particular woman has been exposed, such as the presence of carcinogens, she might have a greater or lower chance of getting cancer than is typical for her demography.

However, crucial for the dispositional account is the idea of the intrinsic dispositional tendencies that an individual contributes to the overall mixture of tendencies in play. A sugar cube will have an intrinsic tendency to dissolve in liquid, for instance. Glass will have an intrinsic degree of strength. Now it is more than likely that some strong glass has been broken while some weak glass has not because whether or not any particular pane of glass breaks depends on the other factors in operation, which we understand as mutual manifestation partners. All causally active powers,

operating upon one situation, are understood as mutual manifestation partners jointly producing their effect (Martin 2008: ch. 5 and Mumford and Anjum forthcoming b). Strong glass might, for instance, break with only slight impact if the glass is subject to very low or very high temperatures. Here, the temperature can be seen as a causal contributor. The effect produced is a result of the composition of causes, as outlined by Mumford and Anjum (2011a: chs 2 and 4), which allows for compositional pluralism: the idea that composition need not be according to a linear function. But each particular will bring its own intrinsic propensities into that composition, even though the effect is caused jointly. The individual adds any genetic predispositions for or against cancer, for instance, plus any acquired tendencies towards it. The sugar cube has an intrinsic disposition towards dissolving but one it can exhibit only when it meets a suitable solvent.

Because such tendencies are understood as intrinsic to the individual, rather than individual qua member of a group, they are the sort of thing that might be revealed in one well conducted experiment, in which the experimenter attempts to assemble all the right mutual manifestation partners for that disposition (Cartwright 1989, for example). If one wishes to test magnesium for flammability, for instance, one need not conduct a prolonged series of many trials. One experiment in which magnesium ignites in oxygen might be enough. And any further tests might just be to check the experiment meets the standards of repeatability, or to train budding scientists, rather than be an exercise in acquiring more data.

Statistics can be symptomatic of the intrinsic propensities of things but they do not define or generate them. The reason for this can now be set out explicitly, though it follows from what has been said earlier. Powers compose but they do not decompose. The intrinsic tendencies of all the powers operating on a situation will compose in accordance with some function to effectively become a resultant tendency, which can then be thought of as producing the effect. The resultant power should not be thought of as some new entity. As we say, this is entirely composed by the component powers according to some function. There is no additional element other than the component powers but in their mutual manifestation, or joint action, they can work with or against each other, sometimes in a non-linear way. This means that with any set of individual powers, each with their distinct magnitudes, there is only one resultant power that they compose. But we cannot decompose them, meaning that from any known resultant power, there is any number of component powers from which it could have been compositionally produced.

This finding explains our problem with reading objective chance purely from frequencies. In any population, many different and varying tendencies will go into producing a distribution in a trial, for instance of Fs that are G. But we cannot decompose that distribution and ascribe it back to individual Fs, as the chance that they will be G, in anything other than a purely statistical sense. 40% of Fs being G is consistent with many different degrees of propensity of things that are F being G.

7. Probabilistic tendencies

We can now return to an earlier issue: of whether all powers are effectively being treated as probabilistic in our account. We say not because a distinction can still be drawn between tendencies of a certain strength, that either manifest or don't, and genuinely probabilistic dispositions in which

there is a balanced distribution of dispositions towards various outcomes. A fair coin gives a model of the latter, where the chance of landing heads or tails if tossed is 50:50, which we represent as a single double-headed vector (fig. 2). Another example is a die, which disposes towards six different outcomes simultaneously – equally so if it is fair – and where the chance of all those outcomes must add up to one. As should be clear from what has been said, however, both a regular tendency and a probabilistic disposition can be understood as propensities and are an ontological basis for the assignments of objective chance.

The issue of probabilistic powers ought also to be kept separate from the issue of indeterminism. The probabilistic power of a die, which disposes in six different directions, may well be a perfectly deterministic one, if the side the die will land is fixed by all prior states. And contrariwise, the power of radioactive decay might well be entirely indeterministic, in respect of the exact moment at which the particle decays, even though it can be understood as a regular tendency, disposing towards only one outcome.

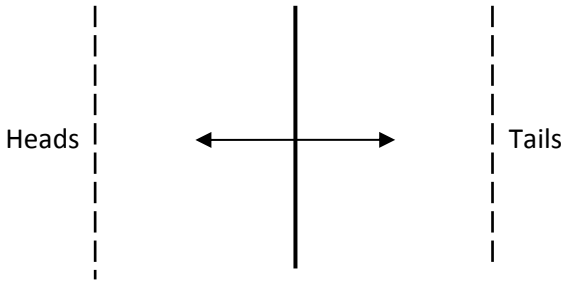

Figure 2: a probabilistic power, disposing equally towards two possible outcomes

What the dispositional modality that we have advocated does allow, however, is a category of genuinely chancy or stochastic causation. When Mellor (1971) and Popper (1990) present their propensity theories, where they too understand propensities in terms of powers, they both give accounts in which chancy causation works by resolving into non-chancy causation. An F may have a 50% propensity to be a G, for instance, but what happens when it becomes a G is that it goes through a temporally extended process in which its chance of being G gets closer and closer to 1. And when that chance reaches 1, that is exactly the point at which it becomes G. These accounts resolve, therefore, into causal necessitation theories in which a cause produces its effect by eventually necessitating it. As already mentioned, this makes causation a matter of all or nothing rather than authentically probabilistic all the way through.

With the dispositional modality, however, there is conceptual space for the possibility that causes produce their effects without necessitating them. In other words, there could be a genuinely chancy case in which there was a 50% propensity towards H and H is then just caused, without any additional factor necessitating it. This goes against the engrained causal necessitarian way of thinking. But if we are to allow stochastic causation at all, it might as well be genuinely stochastic instead of just the same old necessitarian causation, thinly disguised.

8. A case for causal singularism

We have a variety of ways of investigating the world to discover causes and probabilities. It seems plausible from what we have considered that different experimental methods suggest different thinking about the nature of causation and probability. Generalising, experimental methods to be found in chemistry and physics (see for instance Baetu 2011 and McKay Illari and Williamson 2011) are conducive more to causal singularist and propensity theories of probability, for instance, while the use of RCT statistics and population data fits better with relational theories of causation and frequentism about probability. Some call the latter covering-law models of causation. The essential difference between these two models of causation and probability is that the covering-law model takes the general causal and probabilistic facts to be primary, which are then distributed to the particular instances. The singularist takes that particular causal truths and truths of propensity to be primary, which then determine the general causal truths. We have effectively presented an argument in favour of singularism: the argument being that singular causal and probabilistic truths can reliably determine the general ones, but the general causal and probabilistic truths cannot reliably determine the single, particular ones.

Perhaps a more pressing methodological conclusion, however, is that one approach does not automatically transfer over to another (Kerry et al. 2012, Eriksen et al. 2013). When making a clinical decision about an individual treatment in medicine, for instance, it matters whether the patient has a genuine propensity of responding to it or whether there is 'merely' a statistical fact that others have responded. It seems that different methods give evidence of different things and that these might sometimes conflict. We can think of each experimental method as providing different symptomatic evidence of causation and propensity. But as we have shown, these different approaches may be based upon different ideas of what causation consists in and what is the nature of probabilities in nature. The theory of powers suggests reasons for a degree of scepticism about some approaches, which might be thought of as providing a rough and imperfect approximation of causation: its symptoms, at best. A singularist, propensity theory allows better for individual variation, causal complexity and context sensitivity of effects, which we take to be essential features of causation (Mumford and Anjum 2011a: ch. 7).

9. Conclusion

We have argued that regularity, often taken to be the starting point in many accounts of causation, is not something easily found in typical causal connections. Some experimental methods do not seem to be looking for it. And those that do – population studies and randomised controlled trials – also show a more complicated picture. Such results indicate a frequency of distribution, but one that is not necessarily constant through a series of trials, and instead suggest we understand causation in terms of tendencies and probabilities.

A metaphysics of causation built upon real causal powers, manages well to interpret such experimental methods. But it also fits more naturally with a singularist and tendency view of nature

in which the statistical facts are generated by the individual propensities of things. This gives the dispositionalist good cause to reject frequentist interpretations of probability as inadequate both metaphysically and in terms of scientific methodology.

References

Armstrong, D. (1983) *What is a Law of Nature?*, Cambridge: Cambridge University Press.

Baetu, T. (2011) 'Mechanism Schemas and the Relationship Between Biological Theories' in P. McKay Illari, F. Russo and J. Williamson (eds) *Causality in the Sciences*, Oxford: Oxford University Press, pp. 407-24.

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

Bird, A. (2007) *Nature's Metaphysics*, Oxford: Oxford University Press.

Broadbent, A. (2013) *Philosophy of Epidemiology*, London: Palgrave Macmillan.

Cartwright, N. (1989) *Nature's Capacities and Their Measurement*, Oxford: Oxford University Press.

Cartwright, N. and Hardie, J. (2012) *Evidence-Based Policy: a Guide to Going it Better*, Oxford: Oxford University Press.

CDC - Centers for Disease Control and Prevention (2008) 'Smoking-Attributable Mortality, Years of Potential Life Lost, and Productivity Losses – United States, 2000–2004', *Morbidity and Mortality Weekly Report (MMWR)*, 57: 1226-1228.

Lowe, E. J. (2012) 'Mumford and Anjum on Causal Necessitarianism', *Analysis*, 72: 1-4.

Eagle, A. (ed.) (2011) *Philosophy of Probability*, London: Routledge.

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

Eriksen, T. E., Kerry, R., Mumford, S., Noer Lie, S. A., Anjum, R. L. (2013) 'At the Borders of Medical Reasoning - The Aetiological and Ontological Challenges of Medically Unexplained Symptoms', *Philosophy, Ethics and Humanities in Medicine*, 8: 1-11.

Geach, P. T. (1961) 'Aquinas', in G. E. M. Anscombe and P. T. Geach, *Three Philosophers*, Oxford: Blackwell, 1961, pp. 65-125.

Gillies, D. (2000) *Philosophical Theories of Probability*, London: Routledge.

Guyatt G. H., Sackett D. L., Sinclair J. C., Hayward R., Cook D. J., Cook R. J. (1995) 'Users' Guides to the Medical Literature. IX. A Method for Grading Health Care Recommendations', *The Journal of the American Medical Association* (*JAMA*), 274: 1800-4.

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

Howick, J. (2011) *The Philosophy of Evidence-Based Medicine*, Oxford: Blackwell.

Hume, D. (1739) *A Treatise of Human Nature*, L. A. Selby-Bigge (ed.), Oxford: Clarendon Press, 1888.

Hume (1748) *An Enquiry Concerning Human Understanding*, P. Millican (ed.), Oxford: Oxford University Press, 2007.

Kerry, R., Eriksen, T. E., Mumford, S., Noer Lie, S. A., Anjum, R. L., (2012) 'Causation and Evidence-Based Practice: An Ontological Review', *Journal of Evaluation in Clinical Practice*, 18: 1006-12.

Lewis, D. (1986) *On the Plurality of Worlds*, Oxford: Blackwell.

Lewis, D. (1994) 'Humean Supervenience Debugged', in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, pp. 224-47.

Lewis, D. (2004) 'Causation as Influence', in J. Collins, N. Hall and L. Paul (eds), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 277-90.

Martin, C. B. (2008) *The Mind in Nature*, Oxford: Oxford University Press.

McKay Illari, P. and Williamson, J. (2011) 'Mechanisms are Real and Local', in P. McKay Illari, F. Russo and J. Williamson (eds) *Causality in the Sciences*, Oxford: Oxford University Press, pp. 818-44.

Mellor, D. H. (1971) *The Matter of Chance*, Cambridge: Cambridge University Press.

Mumford, S. (2004) *Laws in Nature*, Abingdon: Routledge.

Mumford, S. and Anjum, R. L. (2011a) *Getting Causes from Powers*, Oxford: Oxford University Press.

Mumford, S. and Anjum, R. L. (2011b) 'Dispositional Modality', in C. F. Gethmann (ed.), *Lebenswelt und Wissenschaft*, Hamburg: Meiner Verlag, pp. 380-94.

Mumford, S. and Anjum, R. L. (2016, forthcoming a) 'Overdisposed', in *What Tends to Be: Essays on the Dispositional Modality*.

Mumford, S. and Anjum, R. L. (2016, forthcoming b) 'Mutual Manifestations and Martin's Two Triangles', in J. Jacobs (ed.) *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*, Oxford University Press, pp. 74-87.

Popper, K. (1990) *A World of Propensities*, Bristol: Thoemmes.

Psillos, S. (2002) *Causation and Explanation*, Chesham: Acumen.

Rabins, P. (2013) *The Why of Things: Causality in Science, Medicine, and Life*, New York: Columbia University Press.

Rosing, J., Curvers, J., Tans, G. (2001) 'Oral Contraceptives, Thrombosis and Haemostasis', *European Journal of Obstetrics & Gynecology and Reproductive Biology*, 95: 193-197.