Brit. J. Phil. Sci. **42** (1991), 397–412 *Printed in Great Britain*

REVIEW ARTICLE[1]

Correlations and Causes

1 INTRODUCTION

Conventional wisdom has it that social surveys are a poor guide to causal connections. An established covariation between the budget deficit and subsequent inflation does not necessarily mean that the former causes the latter. 'Correlation does not imply causation', generations of first-year social science students are told. And the general consensus is that, in areas where active experimentation is often both impractical and immoral, it is well nigh impossible to separate the causal wheat from the correlation chaff.

Even so, working economists and other social scientists have developed a battery of techniques for doing precisely that. Once the social science students have completed the first year, they progress to courses with titles like 'Causal Inferences in Non-Experimental Research'. And there they are taught that, in the right form, correlational premises do imply causal conclusions after all. More specifically, they are shown how to distinguish genuine causal connections from spurious non-causal correlations, by checking whether or not the connections remain when we examine partial correlations conditional on further background factors. For example, if the correlation between the budget deficit and inflation disappears when such further factors as the increase in the money supply are taken into account—if, as it is often put, the increase in the money supply 'screens off' inflation from the budget deficit—then the original correlation is exposed as spurious. But if no such screeners-off exist, then the original correlation indicates a genuine causal connection.

In the 1930s and 40s the econometricians Tjalling Koopmans, Jan

[1] Review of Nancy Cartwright [1989]: *Nature's Capacities and their Measurement*. Oxford: Clarendon Press. $x+268$ pp. ISBN 0-19-824477-0.

Tinbergen, and Trygve Haavelmo showed how to apply this kind of reasoning to complex systems of interconnected variables. Since then their techniques, and the underlying mathematical theory of multivariate regression analysis, have become part of the basic toolkit, not only of econometricians, but also of sociological, medical and psychological researchers. Probably best known under the name of 'path analysis', which is in effect merely a convenient notational variant of multivariate regression analysis, these widely-used techniques have in recent years been facilitated by the availability of standard computational packages which churn out partial correlations for as many as a hundred possibly related variables.

Philosophers have paid surprisingly little attention to these developments. This is a pity, for philosophical analyses of causation have difficulty in accounting for the direction of causation in general, and for the difference between genuine causes and spurious symptoms in particular. A satisfactory philosophical account of the way that econometricians deal with these issues therefore promises to throw some much-needed light on the nature of the causal relation itself.

This line of thought motivates Nancy Cartwright's new book. *Nature's Capacities and their Measurement* begins with the particularities of econometric practice, and ends up with a number of radical general conclusions about causation, including a striking reinterpretation of the causal underpinnings of non-local quantum mechanical correlations.

Cartwright's central theme is that a proper understanding of the inferential methods used by econometricians undermines any generally Humean approach to causation. According to Cartwright, Humeans turn out to be wrong on two counts: first, in holding that singular causal relationships can be reduced to generic causal relationships; and, second, in holding that general causal relationships can be reduced to non-causal laws of association. Cartwright urges that we should turn away from the Humeans' law of association, and instead analyse causation in terms of *capacities*, whose exercise is responsible for laws, in the special circumstances where laws exist, but which cannot be reduced to laws.

Cartwright's rejection of Humean laws of association in favour of capacities is of a piece with the general anti-theoretical stance of her *How the Laws of Physics Lie* [1983]. In particular, it offers a way of making good a central thesis of that earlier work which has puzzled some commentators, namely, that we can know about the unobservable causes of phenomena, even though we know no laws involving those unobservable causes. However, there is no essential interdependence between the two works. The arguments of the new book are presented as quite separate from those in the old. In this review I shall accordingly consider Cartwright's analysis of causation in its own right and independently of her overall philosophical position.

Before proceeding to details, I should like to declare an intellectual interest. I

share Cartwright's conviction that philosophers of causation have much to learn from the multivariate regression techniques of the econometricians. However, I do not agree that the lesson is that we should dismiss Hume. On the contrary, I think that the econometricians' techniques are interesting precisely because they indicate how best to modify and refine Hume, by showing exactly what patterns of lawlike association count as genuine causes.

In consequence, much of my analysis of Cartwright's arguments will be negative in tone, explaining why I do not think her anti-Humean conclusions follow. So I would like to emphasise at this stage that I nevertheless think this is an extremely important and worthwhile book. Cartwright has ventured into exciting but largely unknown philosophical terrain. Even if some of the conclusions she brings back are flawed, all philosophers of causation will profit greatly from her explorations. Certainly in my own case Cartwright's arguments have made me see far more clearly than before exactly what a Humean can and cannot say.

2 REGRESSION AND ERRORS

In her first chapter Cartwright focuses on the standard 'recursive equation' format of multivariate regression analysis:

$$x_1 = u_1$$
$$x_2 = a_{21}x_1 + u_2$$
$$x_n = a_{n1}x_1 + a_{n2}x_2 + \ldots a_{nn-1}x_{n-1} + u_n$$

In this set of equations the variables on the left hand side are supposed to represent the *effects* ('dependent variables', 'endogenous variables'), and those on the right hand side the *causes* ('independent', 'exogenous'). The *u*s represent 'error terms', that is, the remaining unknown or unobservable influences on the dependent variables. Each error term is supposed to be uncorrelated with the other exogenous variables in the equation in which it appears, and also with the other error terms.

If the correlation between some x_i and x_j is spurious, then the relevant a_{ij} will be zero. This is because each a_{ij} is proportional to the partial correlation coefficient of x_i on x_j, given the other *x*s in the equation for x_i. So if a_{ec} is zero for some effect x_e and some putative cause x_c, then this means that the initial correlation between x_e and x_c disappears when we hold the other causes of x_e constant, thus indicating that x_c isn't itself a direct cause of x_e after all, but at best an indirect cause of x_e (if the other causes of x_e include effects of x_c) or a co-descendant of a common cause of x_e (if the other causes of x_e include causes of x_c).

So far, so easy. The initial obstacle, however, to a philosophical explication of these standard econometric techniques is that the causal conclusions they deliver seem to be an artefact of the way the equations are written. Any set of

simultaneous equations can be trivially transformed into a format which is algebraically equivalent, but in which different variables appear on the left hand side. For example, given the conventions of regression analysis, the equations:

$$x = s$$
$$y = x + s$$

represent y as an effect of x. But we could just as well write these equations as:

$$y = 2s$$
$$x = y - s$$

which can then apparently represent x as an effect of y.

This point has persuaded many philosophers of causation that regression techniques simply build their causal conclusions into the way the equations are written. But this is too quick. As Cartwright explains, the priority-reversing transformed equations will normally violate one of the basic requirements of regression techniques, namely, that the error terms be probabilistically independent of each other and the other causes. For example, take the first two equations in the normal triangular array:

$$x_1 = u_1$$
$$x_2 = a_{21}x_1 + u_2 \qquad (1)$$

Now, these two equations are indeed algebraically equivalent to:

$$x_2 = a_{21}u_1 + u_2$$
$$x_1 = 1/a_{21}x_2 - 1/a_{21}u_2 \qquad (2)$$

which represent x_2 as the independent variable and x_1 as dependent on x_2. However, if the error terms in the original equations are probabilistically independent, as required, then the 'error terms' in the rewritten equations won't be: the 'error term' in the second equation, $-1/a_{21}u_2$, will be negatively correlated with the 'error term' in the first, $a_{21}u_1 + u_2$, and also with the other exogenous variable in the second equation, x_2.

So the requirement of independent error terms is a real constraint, which ensures that the ordering of variables in a set of regression equations isn't just an arbitrary importation of prior causal assumptions. From an epistemological perspective, of course, we often won't know directly whether the error terms are uncorrelated, since they normally represent just those factors that are omitted from explicit investigation. And so in practice we will often have to infer their probabilistic independence from prior assumptions about the causal ordering of variables, derived from the temporal order of those variables, perhaps, or from other background knowledge. But, still, from a metaphysical point of view, there is nothing to stop us regarding the probabilistic independence of the error terms as a basic and objective fact, from which the causal ordering derives.

Having explained all this, Cartwright accordingly turns to the question of exactly *why* the probabilistic independence of error terms should be important for causal structure. This is an issue which has been much muddied by the econometric literature, since there are a number of other reasons why the independence of error terms matters to regression analysis. Cartwright is good on separating out these other reasons, which involve questions of estimation and of the identifiability of coefficients, from what she calls the specifically 'Humean problem' of explaining how we can draw causal conclusions from correlational facts.

In answer to this specifically 'Humean problem', Cartwright first considers the straightforward thought that if an error term *were* correlated with one of the other causes, then this would indicate some kind of hidden causal connection which would invalidate the causal order postulated by the system of equations. But she quickly dismisses this suggestion, on the grounds that it assumes just the kind of connection between correlations and causes that a solution to the 'Humean problem' ought to explain.

I found this response somewhat puzzling. Cartwright is addressing a Humean opponent who wants somehow to *reduce* causes to correlations. But such an opponent *does* offer an explanation of the connection between correlations and causes, namely, that the former reduce the latter. Consider, for instance, the neo-Humean reduction of causation offered by Patrick Suppes in his *A Probabilistic Theory of Causation* [1970]: an earlier A is a cause of a later B iff A is correlated with B and there is no yet earlier C such that the partial correlation of A with B given C is zero. Now, there are various objections which can be made to Suppes' account, not least that it cannot as it stands deal with the kind of complexity present in many econometric models. But it surely isn't a good objection to complain that Suppes simply assumes a connection between correlations and causes, when he ought to be explaining it. According to Suppes, certain complexes of correlational facts simply *are* causal facts, and that's why causes are related to correlations.

So I don't see why a Humean addressing multivariate regression techniques shouldn't argue, elaborating on Suppes' approach, that our notion of causation is constituted by the principles that (a) uncorrelated quantities are causally unconnected, and (b) two correlated quantities are directly causally connected, unless (c) their partial correlation given certain further factors is zero, in which case those further factors are either common causes of both original quantities or else causally intermediate between them. These assumptions will then suffice to impose the standard causal ordering on a set of regression equations with uncorrelated error terms. And so the Humean will have solved Cartwright's 'Humean problem' by appeal to the reduction of causation given by (a)–(c).

Cartwright prefers a rather different solution. She argues that, if any attempt to represent x_c as a cause of x_e also represents x_e as an effect of something *we*

independently know isn't a cause of x_e, then x_c can't be a cause of x_e either. This is an interesting and powerful idea, and Cartwright gives it an independent motivation by showing how it can be used in non-probabilistic contexts to explain why certain 'cooked up' INUS conditions, like Russell's Manchester factory hooters, aren't real causes. In connection with regression analysis, the point is that any re-ordering of equations which makes an earlier x_i dependent on some later x_j, as x_1 was made dependent on x_2 in (2) above, will also make x_i dependent on some error term, as x_1 was also made dependent on u_2 in (2), which we already know can't be a cause of x_i.

For Cartwright the significance of this analysis is that it shows how the link between correlations and causes hinges on prior negative information to the effect that certain factors *aren't* causes of certain others, thus blocking any Humean attempt to reduce causes to correlations without residue. But an obvious retort seems open to the Humean. Maybe in Cartwright's non-probabilistic INUS conditions examples the prior availability of negative causal information is essential. But in the regression analysis case, the information that, say, u_2 can't be a cause of x_1, doesn't have to be imported from outside, but can simply be inferred from the fact that u_2 is uncorrelated with x_1. By the time Cartwright offers her solution, we seem to have lost sight of the question she posed earlier: namely, why is the *probabilistic independence* of the error terms important for causal ordering? And in general there seems no reason why a Humean should not view Cartwright's solution to the 'Humean problem' as a simple consequence of the kind of probabilistic reduction proposed above in (a)–(c).

3 PROBABILITIES

In her second chapter Cartwright approaches the same issues from a somewhat different angle, appealing to more familiar philosophical work on probabilistic causation to reinforce her anti-Humean conclusion of 'no causes in, no causes out'. She considers the following proposal, which she calls 'CC':

> C causes E iff
> $Pr(E/C+/-F_1 \ldots +/-F_n) > Pr(E/-C+/-F_1 \ldots F_n)$,
> where $\{C, F_1, \ldots, F_n\}$ is a complete causal set for E.

CC says that a probabilistic association amounts to a causal connection provided it doesn't disappear when we condition on other relevant factors. (And so is just a simple qualitative version of the econometric thinking discussed in Cartwright's first chapter.) Cartwright points out that, as formulated, CC requires the Fs to exhaust the other factors which are *causally* relevant to E, and is therefore inadmissible as a reduction of causation. The obvious Humean response, however, is to modify CC and require instead that the Fs include all factors which are in any way *probabilistically* relevant to E,

thus eliminating any mention of causation on the right hand side of CC, and so reinstating it as a possible reduction.

Cartwright considers this idea briefly, but promptly dismisses it on the grounds that any notion of 'true' probability is itself likely to be dependent on causal notions (while subjective notions of probability, and her own favoured notion of probability as 'frequencies that pass various favoured tests for stability', are likely to make too many *F*s relevant to *E*, and thus to disqualify as causes some *C*s which really do cause *E*).

Given the overall structure of the book, it is odd that Cartwright is so quick to rule out an appeal to 'true' probabilities at this point. The metaphysics of objective probability certainly raises a number of difficult philosophical problems, to which nobody has any good answers. But, even so, the idea of a probabilistic association between two quantities is surely prima facie independent of any specifically *causal* notions: saying that two quantities generally and non-accidentally tend to occur together doesn't yet seem to be saying anything about what is causing what. And, in any case, for better or for worse, that is what Cartwright's Humean opponents must assume: for their programme is specifically to develop an illuminating philosophical reduction of causal notions in terms of non-causal laws of probabilistic association.

A few pages earlier, at the end of chapter 1 (pp. 35–6), Cartwright explicitly addresses the point that her Humean opponents appeal to lawlike regularities of probabilistic association, and expresses her unease: 'I do not believe in these regularities . . . I do not see many around . . . [C]apacities . . . eliminate the need for laws . . .' However, at that stage she then admits, 'This is a strong statement of an undefended view, and does not really belong here. But I mention it because it helps put in perspective the arguments of this book.' Given this, it is surprising that something very like this undefended view reappears in the next chapter as a crucial reason for rejecting a natural Humean move.

4 CAUSAL INTERMEDIARIES

In the third chapter Cartwright raises the stakes, and aims to show that, in order to get causal conclusions out of probabilities, we need to put in, not just irreducibly causal premises, but irreducibly *singular* causal premises. Her argument hinges on the existence of causal intermediaries. Consider CC once more. If we allow the *F*s to include *all* factors relevant to result *E*, including factors that are causally intermediate between *C* and *E*, then we won't ever get any *C* causing any *E*, for such intermediate factors will always screen off *E* from *C*. (If *C* causes *E* *via* *I*, then $Pr(E/C\&I)$ will equal $Pr(E/-C\&I)$, and $Pr(E/C\&-I) = Pr(E/-C\&-I)$.)

A standard solution to this problem is to restrict the *F*s to factors which occur no *later* than *C*. The idea behind this move is that, if the correlation

between C and E is really spurious, then some common cause of C and E must obtain *prior* to C; conversely, if the only factors that screen off E from C are temporally intermediate, then C must be a genuine cause of E, albeit one that proceeds via those intermediaries.

However, as Cartwright points out, this strategy cannot cope with certain kinds of fine structure. In particular, it will inevitably be blind to the fact that there might be a number of different causal paths from C to E, along some of which C might even exert a negative influence on E. Cartwright's example (following Hesslow, 1976) is the causal connection between contraceptive pills and thrombosis.

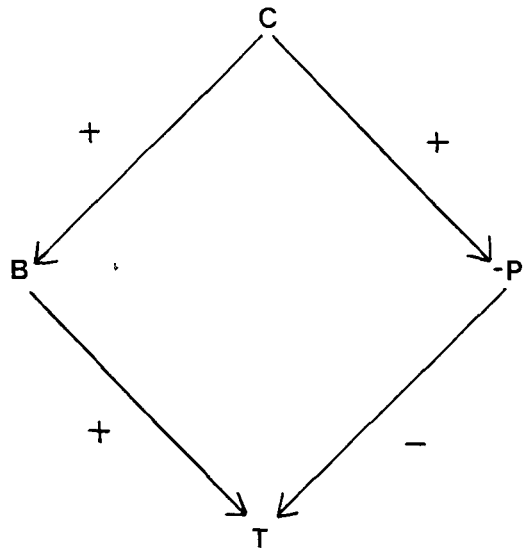

STRUCTURE 1

There are two causal routes between contraceptive pills (C) and thrombosis (T): the pills can produce a blood-clotting chemical (B) and thereby cause thrombosis, but they can also prevent pregnancy ($-$P) and thereby prevent the thromboses which arise in pregnancy. If we control only for factors that obtain prior to taking the pills, then we shall simply average out these two tendencies, and may well conclude that contraceptive pills in general make thrombosis less likely. Cartwright concludes that, in order to discern the path by which contraceptive pills operate as a positive cause of thrombosis, we also need to take into account intermediate events. Not only that, we need to look specifically at the probabilistic difference that the pills make to thrombosis *in those particular cases where the pills cause the blood-clotting chemical.*

This is why Cartwright thinks *singular* causal premises are needed if we are to get causal conclusions out of probabilities. However, as she recognizes, there are various further ploys open to the Humean at this point. After all, the multivariate regression techniques discussed earlier are designed precisely in order to allow the different paths by which one event causes another to be gauged from the probabilities alone. In the case at hand, for example, the multivariate regression technique would be to estimate the path which goes via the blood-clotting agent by considering how strongly the pills remain correlated with thrombosis when we separate our reference class of women into those who are pregnant and those who are not. More technically, our regression model would be:

$$\begin{aligned} x_c &= u_c \\ x_b &= ax_c + u_b \\ x_{-p} &= bx_c + u_c \\ x_t &= dx_b + ex_{-p} + u_t \end{aligned}$$

And we would estimate d, the coefficient which tells us how important the blood-clotting agent is for thrombosis, from the partial correlation coefficient of x_t on x_b given x_{-p}. (In fact, in this example some of the variables—in particular, pregnancy—are essentially qualitative rather than quantitative, so we would need a mixed multivariate regression and analysis of variance model. But this is irrelevant to the general issue of causal structure.)

Cartwright discusses two diffficulties for this strategy. The first relates to the continuity of causation in time. Since even multivariate regression analysis can deal with only a finite number of variables, its models will at best only approximate to continuous causal processes, assuming some discrete 'time chunking', and then drawing conclusions about causal connections between events at those discrete points of time. But, as Cartwright objects, time is not discrete, and so we will always be left with the threat of further significant fine structure within our temporal chunks.

Cartwright does not consider the obvious reply, however. As with all continuous processes, the natural strategy is to start with discrete approximations, and then proceed to the limit. If regression techniques can tell us about the approximations, then a limiting procedure can define the causal truth as that causal structure which is uncovered at some point in the sequence of finer and finer time chunkings, and which is not overturned at any later point. (For versions of this idea, see Papineau [1986], [1989].)

Cartwright's second objection is more fundamental. She points out that the regression strategy of estimating the causal strength of the blood-clotting path by seeing what *extra* difference the pills make, once it is fixed whether or not the women are pregnant, assumes that the pills work *independently* in (i) causing the blood-clotting chemical, and (ii) preventing pregnancy. This is because the regression analysis in effect simply views the overall correlation

as the *sum* of the two paths, and estimates the strength of each path by taking away the other. However, if the chemical and the pregnancy-prevention are interdependent effects of the pills, then this won't work. And indeed, once such interdependence is allowed as an open possibility, then it does seem that any estimation of the strength of the pills–blood-clotting–thrombosis path needs a prior identification of a sample of *particular* cases in which the pills do *cause* the blood-clotting chemical, as Cartwright insists.

5 COMMON CAUSES AND QUANTUM MECHANICS

This last argument cuts deep, for it questions the principle that a common cause will always screen off correlations between its different effects. It is worth being clear about the issue involved here. Cartwright's thought is not simply that the two intermediate effects B and $-$P may be interdependent in the sense, say, that the contraceptive pills may prevent pregnancy *by* producing the blood-clotting chemical. This is simply to hypothesise Structure 2 in place of Structure 1, and the choice between these structures can once more simply be made using standard regression techniques.

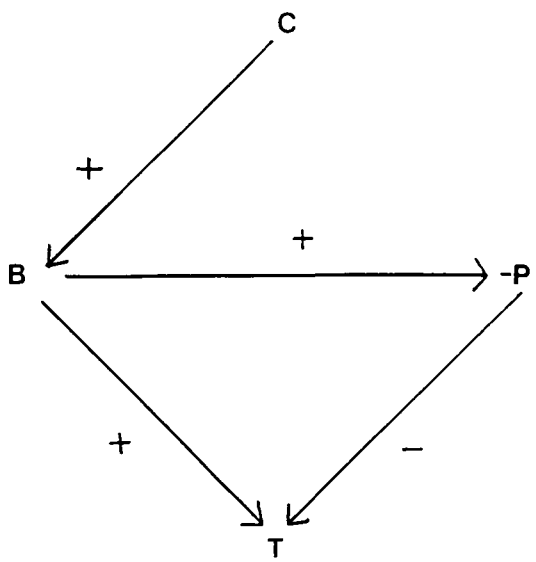

STRUCTURE 2

Rather Cartwright is suggesting that, although B and $-$P may be *causally* independent effects of C, in that neither affects the other, as in Structure 1, they may yet fail to be conditionally *probabilistically* independent, in that they may remain correlated even after we condition on C.

According to Cartwright, there is nothing very surprising about this possibility, and she offers some everyday examples of how common causes can fail to screen off correlated effects in certain kinds of 'constrained' structures. However, I find the idea of such unscreened-off correlations between effects highly counter-intuitive, and I was not persuaded that her everyday examples could bear the necessary argumentative weight: in general such everyday unscreened-off correlations can always be reinterpreted as arising from causal links between the 'independent' effects, or from those effects being different descriptions of the same event. (*Cf.* Butterfield [1989].)

Of course, Cartwright also has in mind the Einstein–Podolsky–Rosen quantum mechanical correlations. These are correlations between spatially separated results which are not screened off by the prior quantum mechanical state (nor indeed, as John Bell showed, can they plausibly be screened off by *any* prior state). But even these EPR correlations do not yield unproblematic support for Cartwright's position. Cartwright wants to say that these correlations manifest common causation without screening off. But this requires, not just that these correlations are not screened off by anything, which is uncontroversial, but also that the prior quantum mechanical state is a *common cause* of the correlated results, which is a highly non-standard view among interpreters of quantum mechanics.

Perhaps it is right to accept that the prior quantum state in some sense *accounts for* the correlated probabilities of the separated effects—after all. nothing else will. But it is a further step to say that the quantum mechanical state is the common *cause* of those results. It seems to me no accident that most philosophers of quantum mechanics, Cartwright apart, are disinclined to make this further step. For it seems likely that the screening-off requirement on common causes is somehow closely related to the *asymmetry* of causation. (Note that the reverse requirement is manifestly false: common *effects* do not generally screen off correlations among their joint *causes*.) By dropping the screening-off requirement on common causes, Cartwright threatens to leave us with a notion of dependency which has lost this crucial asymmetry. Of course, it remains possible that Cartwright might be able to account for the asymmetry of causation in other ways. But she does not address the issue, and so fails to allay the suspicion that she has shorn causation of its central feature.

Even so, Cartwright's comparison of econometrics with quantum mechanics is extremely illuminating. As Cartwright observes, if the independent *x*s in a set of multivariate regression equations each make a *deterministic* contribution to the dependent variables, with the randomness contributed entirely by the 'error terms', as is suggested by the standard way of writing regression equations, then the probabilistic independence of those error terms *entails* the screening-off requirement. (Compare the way in which the assumptions of deterministic hidden variables plus probabilistic independence of measurement settings imply Bell's inequality for the EPR set-up.) By contrast, if the

contribution of the independent variables is itself an essentially stochastic matter—and Cartwright develops an alternative notation to cater for this possibility—then there is room for an independent variable's contributions to two dependent variables to be coordinated in such a way as to violate the screening-off requirement (which is in effect what happens in the quantum mechanical context when the EPR correlations violate Bell's inequality).

Cartwright concludes, as we have seen, that in our indeterministic world common causes need not always be screeners-off. But it seems to me as plausible to run the argument the other way: that is, to retain the screening-off requirement on common causes, for the reasons connected with asymmetry mentioned above, and infer instead that genuinely *causal* connections, as opposed to the non-causal dependencies described by quantum mechanics, are peculiar to those macro-events which enter into effectively deterministic relationships. This is not to deny that our world is permeated by indeterminism. It is simply to hold that, when we focus on that indeterminism, and lose sight of the effectively deterministic connections which characterize much of the macro-world, then we also lose sight of the asymmetries which are essential to our causal notions.

Cartwright's line on the EPR correlations does avoid one difficulty which faces the Humean. In an EPR experiment the result on one wing is correlated with the result on the other wing and nothing screens off this correlation. But nobody wants to say therefore these spacelike separated events cause one another. Yet this is what is implied by the Suppes-style Humean reduction outlined in Section 2 above. However, there are other ways of dealing with this, apart from Cartwright's strategy of interposing a prior quantum mechanical 'common cause' between the two results. For example, it is arguable that the EPR correlations lack a certain 'robustness' characteristic of more familiar everyday correlations (Redhead [1987], pp. 102–7), which then opens the way for Humeans to argue that causes reduce specifically to unscreened-off *robust* correlations.

6 CAPACITIES

Cartwright's defence of capacities as the basic causal reality comes in chapters 4 and 5. She points out that a given causal factor will often make just the *same* probabilistic difference to a given effect, whatever other causal influences it is acting in concert with. Indeed this requirement is built into many familiar causal models; in particular, it is built into multiple regression models, via the assumption that the dependent variables are a linear function of the contributions made by the independent variables; and, more generally, it is built into any causal model which allows us to speak of *the* context-independent amount by which a given cause raises the probability of some effect.

Cartwright takes this kind of unanimity of causal influence across different contexts to provide an argument for causal capacities. She argues that any given causal factor has the capacity to produce a certain effect, which it carries with it from context to context. For example, aspirin *per se* has the capacity to stop headaches, in that, given any specific causal context, aspirin increases the chance of a headache disappearing. Cartwright argues laws of probabilistic association, insofar as they exist, are consequences of capacities, not their basis: headaches disappear with specific frequencies in specific circumstances *because* aspirins and other relevant factors have the capacities that they do, not the other way round.

Cartwright recognizes that these observations will not necessarily convince her Humean opponents. Humeans will argue that Cartwright's 'capacities' are simply a kind of meta-pattern, a kind of recurring structure in first-order patterns of association. Such meta-patterns are certainly noteworthy, and often methodologically important, but this is no reason to conclude, as Cartwright does, that causation would not exist without them. After all, it seems easy to imagine causal factors whose influence does depend on what other causes they are acting in concert with. Indeed, this seems to happen often enough in the actual world, as when we say that one cause 'interacts' or 'interferes' with another.

In response to these Humean points, Cartwright appeals to her earlier argument, in chapter 3, that singular causal premises play an essential role in linking facts of generic causation to laws of association. In the present context, this emerges as the claim that, in order for generic capacities to manifest themselves in laws of association, we need to hold fixed singular facts about the presence or absence of interference or interaction. In effect, Cartwright accepts that the meta-patterns which manifest capacities are not perfect, because of interference and interaction; but instead of dismissing capacities, and settling for the overall Humean mosaic, she appeals to independent premises about interference and interaction, in order to disentangle the capacities from the Humean mosaic.

My major difficulty here, as with Cartwright's conclusions in chapter 3, was with the ontology of her independent causal premises. I understand what it is for anoxia, say, to 'interfere' with aspirin's ability to stop headaches, if this simply means that aspirin together with anoxia is not generally followed by headaches disappearing, whereas aspirins plus most other things is. But of course this kind of general pattern is simply more grist to the Humean mill. Cartwright's extra causal premises about interference and interaction need somehow to transcend further laws of association. Yet she is less than convincing in explaining why this is so:

> In practice one looks for independent evidence that an interaction is occurring, and some account of why it should occur between these variables and not others, or at these levels and not others. The chemistry examples are a good case. One

> does not just say the acid and the base interact because they behave differently together from the way they behave separately; rather we understand already a good deal about how the separate capacities work and why they should interfere with each other in just the way they do. (p. 165.)

From my Humean point of view, I couldn't see anything in this line of argument to distinguish it from the familiar but unsuccessful thought that we can restore some anti-Humean cement to macro-laws of association by explaining them in terms of micro-laws of association.

Cartwright has an additional argument for capacities, based on an analysis of Galilean idealizations. Scientists often consider what would happen in ideal but never-actualized situations, such as in a perfect vacuum. Cartwright argues that such deliberations can only yield substantial information if capacities are real. Her thought here is that, if massive bodies have the capacity to attract each other, say, then their behaviour in a perfect vacuum will directly reveal a power which also operates, in concert with other powers, outside vacuums; but if there are no capacities, then knowledge of the ideal situation will yield nothing but empty claims about never-instantiated generalizations.

This seems to me an interesting and valid argument. However, as a Humean, I deny the premise that the analysis of ideal situations yields substantial information about the actual world. Instead I take Galilean idealizations to be primarily of heuristic interest: they are an aid to investigating the asymptotic behaviour of certain effects as certain causes tend to zero. In my view, the real import of an ideal analysis lies in its heuristic suggestions about what kind of variation to expect as we move away from the limit, heuristic suggestions that can be confirmed and refined by further empirical investigation. From this perspective, ideal analysis remains important for science, but not because properties have causal capacities which they carry with them into unactualized contexts, but simply because idealizations provide scientists with bearings to help explore the overall Humean mosaic. (For further discussion of ideal analysis, see Papineau [1976].)

7 CONCLUSION

Cartwright's discussion of capacities illustrates both an overall strength and an overall shortcoming in her book. Cartwright is a scientists' philosopher of science. Her strength lies in her ability to illuminate important features of scientific thought. The shortcoming is her readiness to draw metaphysical morals from methodological insights.

For example, Cartwright is quite right to insist that the compositionality of causal influences is a crucial and little-recognized precondition of most scientific theorizing. And she is extremely illuminating on the way in which different sciences deal with this compositionality requirement, and on how

they admit different kinds of exceptions to their meta-patterns. But her arguments lose force when she attempts to reify this meta-patterning into a new ontological category, and unconverted readers are likely to judge at this point that she is making a metaphysical necessity out of a practical prerequisite.

Similarly, Cartwright is quite right to urge the importance of Galilean idealizations for scientific theorizing. And she is incisive in distinguishing such idealizations from various other kinds of theoretical abstraction. But when she assumes that, since idealizations are methodologically important, they must yield genuine information (namely, about capacities), she seems once more to be moving too quickly from premises about scientific practice to a conclusion in metaphysics.

Perhaps this tendency to conflate methodology and metaphysics can also be discerned in Cartwright's central discussion of causal ordering. In Section 5 above I criticized her treatment of this issue on the grounds that it threatens to undermine the asymmetry of causation. However, in Cartwright's own discussion, this threat is not as obvious as it might be, since she takes it as given that, whatever else is true of causal ordering, causes never succeed their effects in time. (In particular, as Cartwright allows at the end of her book, this temporal assumption plays a crucial role in identifying the prior quantum mechanical state in the EPR set-up as a common cause.)

Now, it is certainly true that, as a matter of practical methodology, few working scientists ever give the possibility of causes which temporally succeed their effects a second thought. But it is equally true that nearly all philosophers interested in the metaphysics of causal direction think that there are good philosophical arguments against assuming temporal order in explaining causal order (and that, if anything, the explanation should go the other way). Because of this, Cartwright's arguments are again likely to seem more persuasive to readers with a methodological focus than to those whose interest in causation is fundamentally metaphysical.

In suggesting that Cartwright's metaphysical conclusions seem on occasion to outrun her methological premises, I do not want to belittle the substantial philosophical merits of this book. As I warned at the beginning of this review, I have focused on Cartwright's professed aim of discrediting Humeanism, and I have stressed the ways in which she seems to me to fail in this aim. But I trust my discussion has also made it clear that she has introduced a number of important new strands to the theory of causation, which deserve fuller discussion than I have been able to give them here, and that her analysis of the causal significance of correlations penetrates much further than any previous attempt.

I should also mention that the book deals with a number of further issues and case-studies which I have not had space to discuss at all. This wealth of detail gives the book a depth of purpose which is rare in the philosophy of

science. Such further material includes: the Lamb dip in gas lasers, the recent history of economic forecasting, the Einstein–De Haas magnetism experiment, J. S. Mill on tendencies, Glymour, Scheines, Spirtes and Kelly on causal structure, propagational rationales for screening off, semantics for generics and habituals, functional abstraction, and the Stanford Relativity Gyroscope Experiment.

D. PAPINEAU
Department of History and Philosophy of Science
King's College London
Strand, London WC2R 2LS

REFERENCES

BUTTERFIELD, J. [1989]: 'A Space–Time Approach to the Bell Inequality', in J. Cushing and E. McMullin (*eds*), *Philosophical Consequences of Quantum Theory*. University of Notre Dame Press.

CARTWRIGHT, N. [1983]: *How the Laws of Physics Lie*. Clarendon Press.

HESSLOW, G. [1976]: 'Discussion: Two Notes on the Probabilistic Approach to Causality', *Philosophy of Science*, **43**, pp. 290–2.

PAPINEAU, D. [1976]: 'Ideal Types and Empirical Theories', *British Journal for the Philosophy of Science*, **27**, pp. 137–46.

PAPINEAU, D. [1986]: 'Causal Factors, Causal Inference, Causal Explanation', *Aristotelian Society Supplementary Volume LX*, pp. 115–36.

PAPINEAU, D. [1989]: 'Pure, Mixed and Spurious Probabilities and Their Significance for a Reductionist Theory of Causation', in P. Kitcher and W. Salmon (*eds*), *Scientific Explanation: Minnesota Studies in the Philosophy of Science, vol. XIII*, University of Minnesota Press, pp. 307–48.

REDHEAD, M. [1987]: *Incompleteness, Nonlocality, and Realism*. Clarendon Press.

SUPPES, P. [1970]: *A Probabilistic Theory of Causality*. North-Holland.