

**Occasional Paper Series
Paper 43**

***The role of cause and effect in education as a
social science***

Stephen Gorard

Cardiff University School of Social Sciences

ISBN 1 872330 614

Research Team

Executive Group:

Professor Stephen Gorard (Director)
Professor Gareth Rees
Professor John Furlong
Dr Laurence Moore
Professor Ken Prandy
Dr Ray Crozier

Research Staff:

Dr Chris Taylor
Patrick White
Helen Taylor (Project Administrator)

Contact Details

ESRC TLRP Research Capacity Building Network
Cardiff University School of Social Sciences
Glamorgan Building
King Edward VII Avenue
Cardiff
CF10 3WT

Tel. 029 2087 5345
Fax. 029 2087 4678
Email. TaylorH1@Cardiff.ac.uk

www.cardiff.ac.uk/socsi/capacity

The role of cause and effect in education as a social science

Stephen Gorard
Cardiff University School of Social Sciences
King Edward VII Avenue
Cardiff, CF10 3WT
02920-875113
email: gorard@cardiff.ac.uk

Abstract

This paper is intended as a stimulus to discussion on the nature of cause:effect models, and their role in educational research. It is clearly not intended to be definitive. Nor is it based on new evidence. It considers in turn models based on no causation at all, models based on causation alongside other explanatory processes, models based solely on causation, and variations of the latter. Among these are purportedly weaker forms of causation, including Granger causation, social determinants, and so-called fuzzy generalisation. Cause:effect models have been based variously on temporal and spatial conjunction, the impact of interventions, and theoretical explanations. The paper concludes that cause:effect provides a powerful, persuasive and near-universal explanation for social and psychological processes, despite our inability to detect it directly. Because of this inability, cause:effect models are more impressive if they have all three of the characteristics just noted. This conclusion, if accepted, has implications for the nature of educational research.

Introduction

Cause:effect models abound in educational research. Other than in purely descriptive work (e.g. 'the achievement gap between boy and girls in 1999 was 17%'), a research report that did not at least imply a causal model might look rather odd. Causes are central to our notion of understanding why things work as they do, but they are just as central to the less sophisticated 'what works' approach (e.g. 'become an effective teacher/school by doing what effective teachers/schools do'). Yet despite this prevalence, social science research methods courses and textbooks tend to overlook the discussion of causal models completely, or else prepare the novice researcher simply with the negative advice that a correlation is not the same as causation. If, over time, the income of the Archbishop of Canterbury tends to rise in line with the street price of cannabis this is not evidence that the Church of England makes money from drug-dealing. In these standard books, everyone is reminded therefore what is not a cause, and what a cause is not. In some methods books there is a section on the potential and limitations of experiments which points to their unique selling point - the claim to be a direct test of cause and effect.

But this is a scarce and recent phenomenon in social science outside psychology. In general, the concept remains untaught and undiscussed. It is the 'skeleton in the cupboard of philosophy' (Salmon 1998).

Descriptive and exploratory research

Some research is, and should be, solely descriptive. It is anyway an essential first step to doing exploratory work for 'before asking why we must be sure about the fact' (De Vaus 2001, p.2). It is, in my opinion, far too common that researchers set out to explain and explore a phenomenon that does not actually exist (Gorard 1998, as cited in Hillage et al. 1998). Recent examples that I have been involved with include attempts to explain: the school-mix effect; the growing gender gap in attainment, and increasing socio-economic segregation in school compositions. The fact that we can create a plausible theory to explain imperfectly understood notions such as these is not evidence that they must exist. Such research should, rather, routinely start from a re-analysis of relevant existing datasets, and base the ensuing exploration on the patterns uncovered in the preliminary work (Gorard 2001).

Even in the exploratory phase the role of 'unfettered' theory remains limited. Whatever methods of data collection and analysis are used, the subsequent theory is an attempt to reconcile empirical findings with pre-existing 'common ground'. This theory is, either explicitly or no, a causal theory (and for the present I interpret the term 'cause' in its widest possible sense), since it suggests how the empirical findings arose. The underlying causal model is perhaps clearest when the research is based on an experimental intervention (although implicit even in the most complex designs). The major claim made for the use of experimental designs (Fisher 1935) is their ability, used correctly, to uncover and test causal mechanisms (in Badger et al. 2000). This is not a moot point, since the difference between experimental and other forms of evidence can be crucial. It can save lives (Roberts 2000). This is not the place to discuss the practical/ethical merits and demerits of experimental approaches in educational research, which are covered elsewhere (e.g. Gorard 2001). Their advocacy by some commentators rests largely on this notion of causation, which makes it timely, for a variety of reasons, to consider here the nature of causal modelling in rather more detail than we usually do.

The paper proceeds by considering three positions in relation to causal models - that they exist, that they do not exist, and that they exist alongside non-causal phenomena. It shows that there is no logical or empirical reason to reject any of these positions, but that educational researchers, by the nature of their remit, are committed to the first. The paper continues by outlining some of the desirable characteristics of a 'good' causal model, and concludes with a brief discussion of the relevance for the future of publicly-funded educational research.

Are outcomes caused?

It is not possible to detect a cause empirically or prove that one exists philosophically. We can never directly sense a cause. We merely induce their existence from our experience of the association of two or more events, and this is nothing more than a habit of mind - immutable though it appears (Hume

1962). A very similar process is observed in both classical and operant conditioning, where the association of two things leads the conditioned subject to behave in the presence of one thing as though it implied the presence of the other. Yet unlike conditioning in dogs or pigeons, the process of induction has been presented as the chief criterion of demarcation between what is considered 'science' and what is not. This is why, despite important developments by Kuhn (1970) and then Lakatos (1978), it has tended to remain the 'skeleton in the cupboard of philosophy' (Russell, in Ayer 1972). Our notion of cause is little more than a superstition. Alternative criteria for the definition of scientific, as opposed to non-scientific, endeavours have been suggested (e.g. by Popper, in Magee 1973). The problem with these is that, despite the claims of their advocates, they do not remove the problem of shaky philosophical foundation of induction.

For example, when observation leads us to question a belief because it brings two beliefs into contradiction we tend to stick with the most familiar of the two concepts (Goodman 1973), which suggests that Popper's notion of falsification does not actually eliminate inductive logic. To use Popper's own example. No number of observations of white swans can prove that all swans are white, for even if all white swans could be accounted there may be other swans that are not white. This, in essence, is the problem of induction and therefore of causation. However often two things appears together (whiteness and 'swanness'), they cannot prove a link. On the other hand, as Popper observes, only one observation of non-white (black perhaps) swan is needed to falsify the proposition that all swans are white. He suggests therefore that scientists proceed not by trying to prove their propositions, but by falsifying them. This is one basis for the purported difference between science and non-science. The problem with this is that Popper, and his advocates, are ignoring the crucial distinction between formal and real-world (henceforth 'Aristotelian') logic. In formal logic, a contradiction such as 'A entails B' and 'Here is an A which is not B' cannot be explored further. The contradiction shows that there is a flaw somewhere in the prior logical chain, perhaps in one or more of assumptions, since both statements cannot be true. Logic does not help us find the flaw (any more than mathematics can help us find a cause, see below). Since A and B are ideal terms we do not attempt to tinker with them and overcome the contradiction. Contradiction is not the same as falsification.

However, in the real world, where A and B become swans and white and so refer to actual objects, at least one of the terms could be misapplied to the real world object in question. Therefore, we can at least consider the possibility that only one of the propositions is falsified by the contradiction. This is what Popper does without making this step explicit. He then states that it is clear which proposition is wrong, so clear that the alternative is usually dismissed as merely 'playing with words' (Thouless 1974). But this clarity is, like induction, actually only a habit of mind as well. In the example, Popper proposes that we change the definition of swan to include the possibility that some swans are black, and does not even bother to argue against the alternative. The other way out of the contradiction is equally logical (even if it appears implausible because of our habit of mind). We could change the definition of black to exclude the possibility of being applied to swans. Thus the thing that looks like a swan is actually not because it is black. The choice is between changing our definition of swan or of black. In this example we tend to prefer changing the definition of the least familiar term, and swan is a much less general term than black. In fact the same is true in every example of 'falsification' that I can think of. What seems like a logical argument for falsification could actually be an appeal to the same non-logical phenomenon of

familiarity that underlies induction, and therefore causation. Put like this the notion of causation sounds, and indeed is, difficult to justify. What are the alternatives?

What if outcomes were not caused?

Another possibility to be recognised and examined is that the concept of causation, on which the apparent pre-eminence of experimental methods rests, is an illusion. It is possible to imagine and describe social life, and events more generally, without reference to causes. Since this is so, and we cannot see, smell, hear, measure or register causes directly it may be unwise to assume that they exist. In fact, an argument could be advanced that this is the most parsimonious and therefore the most scientific explanation of our observations.

A perfectly plausible alternative is one based purely on random events. A large table of pseudo-random numbers can contain arithmetic sequences, and passages of repetition, without us denying their essential randomness. The sequence '0 1 2 3 4 5 6 7 8 9' is as likely to be generated randomly as any other sequence of ten digits, such as '3 2 7 5 8 8 4 5 1 9'. Both are equally 'random' in the sense that we mean when describing such tables. In the same way perhaps the apparent regularities and repetitions that we observe more generally would be expected in a large (possibly infinitely large) universe. On this, admittedly rather extreme view, all scientific propositions are like the superstitions of a gambler who believes that stroking a rabbit foot improves their odds, or of a pigeon in a Skinner box repeating pointless actions in face of an accidental reinforcement schedule.

However, this view, while intellectually coherent, means the end of scientific endeavour and, by definition, is not one that can be logically espoused by anyone engaged in research on teaching and learning. Similarly, an economist believing that market indicators were actually following a 'random walk' could not earn a living as a predictor of these indicators, except as a charlatan.

Nevertheless, causes are seen by some respected commentators as pre-scientific. Pearson (in Goldthorpe 2001) as early as 1892 was calling the idea of causes a 'mere fetish', which was holding up the advance of correlational techniques in statistics. Russell (in McKim and Turner 1997) argued in 1968 that physics no longer seeks causes as they simply do not exist. Causality is perhaps a relic of a bygone age, like the theory that infections were caused by demons invading the body perhaps. The best we can apparently hope for is the identification of 'relatively invariant functional relationships among measurable properties'. So Russell, like Pearson, would argue that scientific laws are idealised correlations. Mathematical statements or systems of equations can describe systems but they cannot express either intention or causality. If we drop a ball in a round bowl it will come to rest in the centre. We may predict this, and say that this was 'caused' by gravity, but we can see neither the cause nor the gravity, and the cause could not be expressed mathematically. This becomes clearer if we drop two balls in the bowl. We can model the final resting places of both balls mathematically but we cannot use this to decide which ball is 'causing' the other to be displaced from the centre of the bowl. The events are mutually determined and this system of mutual determination is what the equations express (Garrison 1993).

In economics as well as physics some commentators have moved away from causal explanations. Wages and interest rates might be inversely related over time, but rather than deciding that one causes the other it might be more realistic to describe them as mutually determining. Mathematics (including statistics) is like formal logic (see above). It can be used to show that systems are, or are not, in equilibrium, and to predict the actual change in the value of one variable(s) if another variable(s) is changed. However, this prediction works both ways. If $y=f(x)$ then there will be a complementary function such that $x=f'(y)$. Which variable is the dependent one (on the left-hand, predicted side) is purely arbitrary. Nothing in mathematics can overcome this. Nevertheless, non-causal mutuality (or concomitance) could be a perfectly reasonable and reasonably useful interpretation of many such sets of events.

What if cause and non-cause co-exist?

Another position worthy of consideration in relation to the existence of causes is that they exist alongside non-caused events. One version of this stance was taken by those advancing the teleological argument for the existence of a god. Their argument was that everything has a cause, so it is possible to follow the causal chain back to the first cause which was, for the want of a better term, god. Ignoring the simple counter-argument that the existence of a first cause actually refutes the first premise (i.e. everything has a cause), it is clear that such advocates are allowing both causes and non-caused phenomena to exist in the same universe. In practice, such an approach is now followed by economists who present evidence for rational choices as a causing agent. These choices, such as those involved in human capital theory, do not appear to work for individuals but only at aggregated levels. One interpretation therefore is that individuals operate using idiosyncratic processes that only appear to be rational when grouped. More overtly, this position was taken in the twentieth century by physicists and others believing that events at some levels are random (uncertain) while at higher levels of analysis they are patterned. In social science this belief appears in models, both quantitative and qualitative, in which the predictable components of behaviour are seen as causal in nature, and the unpredicted (and unpredictable) parts are seen as random error terms or individual whimsy (Pötter and Blossfeld 2001).

My contention here is that all of these positions, while logically possible, are currently as invalid for the practising social scientist as the model of entirely random events. The number of potential explanations for any finite set of observations is actually infinite (created by simply adding more and more redundant clauses to a proposition for example). We overcome this practical problem, and foster cumulation, by concentrating only on the simplest explanations available. These are the most parsimonious, seeking to explain the observations we make without using additional propositions for which there is not already evidence. They are also the easiest to test, and to falsify in the Popperian model. We have no direct evidence for either causes or random events (Arjas 2001), so to use either one of them in an explanation involves making an assumption. To explain a set of observations using both involves making two assumptions, and is therefore unparsimonious. We have enough trouble establishing whether causes exist or not. To allow them to exist alongside unrelated phenomena makes most social scientific propositions completely untestable (for the falsification of a purported cause can always be gainsaid by

the 'whimsy' element). Perhaps this is why the social science of education shows so little progress over time.

Uncertainty could also be merely unpredictability, and it would be arrogant to assume that if we cannot yet predict a set of events then there is no more predicting to be done. Chaos theory is clearly causal but it allows for unpredictability due to complications in computation from the initial states (Gleick 1988). This unpredictability could stem from our inability to predict causatory events, or from our misunderstanding of the basic randomness of events (see above). Both explanations are plausible, but currently untestable. Using both processes together is unnecessary, and trying to combine them into one description often leads to logical difficulties anyway. For example if sub-atomic events are really random, but have an effect on larger processes which are themselves causes, then following the causal chain argument the larger 'causes' are themselves randomly determined and therefore random. And if 'random' events can have a cause then they are not random.

A more complex solution is to construct a model that involves both causation and other competing explanations of a non-determinist nature, such as intentionality through personal choice. Gambetta (1987) describes educational decisions, for example, as a product of what is available to the individual, what the individual wants and, indirectly, the social conditions which shape the individuals' intentions. However, explanations such as this are unparsimonious on two counts. Firstly, within the causal model, if a cause can be either direct or indirect, an infinite number of possible intermediate steps can always be created between the observed direct 'cause' and the hypothesised indirect one (Blalock 1964). If either explanation fits the data, the simplest solution is the best and the simplest solution cannot be both solutions at the same time. Similarly, the problem with causation is not that there are events that it cannot explain, but that it is impossible to measure. Therefore, there is no value in mixing it up with a model of intention which is also perfectly capable of explaining decisions by itself but which is also not open to observation by social scientists. Given that there is no way of deciding between them empirically, either causation or intention can be adopted (it makes little practical difference which at this stage). There is no empirical justification for working with both at the same time (any more than there is for working with causation and randomness). Rather, in a causal explanation, an intention or an individual choice can be an outcome (of social or family background for example) and a cause. The argument is actually about the nature of the cause (or effect), not about whether it is a cause.

Notions of causality

We have perhaps, as shown above, excessive confidence in the notion of causation in social science, but if we question it what is left? At one extreme if the events we observe in our fieldwork are random, with the apparent patterns appearing by chance in an infinite universe or through self-delusion like the figures observed by ancients in the night sky, then there is no social science (and no need for research). We would have to return our grant and RAE-derived funding to the tax-payer!

Even allowing for the existence of genuinely random events alongside cause and effect produces problems. If the two sets of events do not interact, then our explanation is unparsimonious (why not

three types of mutually exclusive events, or thirty?). If the types do interact then randomness 'trumps' causation. A random event cannot be caused in any meaningful sense, and an event caused randomly is random (and we are back to the first extreme). One conclusion is therefore that these unobservable causes are not necessary in philosophical terms, but that they are fundamental to social science, and to learning and understanding more generally. It follows from that, if accepted, that most debate is about what the nature of causation is rather than whether it is. When psychologists argue the nature/nurture controversy, or sociologists debate the relative importance of structure and agency, for example, they are simply arguing about what the relevant causes are.

Excluding a middle-way on this issue leaves two general approaches. Events are not caused (random, unpredictable and inexplicable) and the apparent regularities are due to chance. Events are caused (determined, potentially predictable and explicable) and the apparent exceptions are due to lack of knowledge. The first of these has been covered above. The second allows several different interpretations, and it is important to recognise and examine some of these as well so that we can be clearer when discussing/implying causality in our research which of these interpretations is in operation. The remainder of this section outlines a variety of characteristics for causal models, and suggests the current level of agreement about them.

Association and causation

One way of viewing causation is as a stable association between two elements. Where one is present the other is also, and when one is absent the other is also. It is the constant conjunction that suggests that all possible futures will be like all pasts (Hume 1962). This view of causation has two main problems: we know that it opens us to superstition, and it does not allow for intermittent association (see above). Skinner's accidental reinforcement schedule is a powerful reminder of the dangers of allowing causal models to be based only on association. Skinner's intermittent reinforcement schedule shows us how difficult it might be to shake such causal models once they have been accepted.

We can be easily fooled by association (hence the common caveats about correlations in standard textbooks), especially where these associations involve large numbers and are backed by expertise or apparent authority (Brighton 2000). A case in point appeared in one of the first Programme seminars for the ESRC Teaching and Learning Programme, where the leaders of two projects made the same argument. They accepted that correlation was not the same as causation, but suggested that multi-level multi-variable linear regression was able to detect causes. But linear regression however complex is still based on correlation, having all of the same limitations with the added disadvantages of being harder to understand, and not producing individual-level predictions and residuals. A similar point was made recently by Johnson (2001) about the false distinction in the US between 'causal-comparative' studies using analysis of variance techniques, and 'correlational' studies. Even though comparative models involve comparison between two or more groups (and like correlational techniques are becoming increasingly complex), they do not provide positive evidence of causation in non-experimental designs. It is, perhaps, simply their complexity and the apparent authority of statisticians who understand them that makes others prepared to accept the falsehood. Techniques such as analysis of variance were anyway developed for use in experimental trials, in which there was little or no measurement error

(agricultural applications). It is not immediately clear that they have a value in non-experimental situations, especially in social science where the measurement error is high (Field and Wilkinson 2001).

Despite all of these caveats, purported causal models based only on association appear throughout the research literature, sometimes dominating entire fields of endeavour (see Gorard et al. 1998). Where economists talk about causation they often mean something much weaker, like Granger-causation or temporal relationships, which takes the post hoc ergo propter hoc fallacy of logic and converts it via a little flourish and an empirical test of 'causality' to a seemingly respectable principle. Granger-causation in economics assumes that we are working with a universe of information. If a variable is eliminated from this universal model, and this produces no change in a second variable then the first variable cannot be the cause of the second (Hendry and Mizon 1999). Otherwise it can be said to 'Granger-cause' it. The practical problem with this empirical approach to causation is that a Granger-cause and a cause are not the same thing but they sound confusingly similar, and anyway no one actually works in the 'universe of information'. Economists use regression models very far from universal in nature, sometimes even bivariate, and still claim Granger-causation which becomes, in essence, a fancy term for a correlation. A similar approach is often used in partialling variance in school effectiveness work. Here the argument is for robust dependence. A variable is not a cause if its influence (regression coefficient) is eliminated by the addition of new variables to the system. But this is clearly nonsense (Goldthorpe 2001). A causal path analysis may show that education leads to a higher income but this is very far from showing that education causes income. Robust dependence is not enough. Only a prediction from theory, or a test via intervention, can take us any further than a purely descriptive mathematical relationship.

Given the difficulty of identifying causes, perhaps the best that can be hoped for is to identify only weak causes or 'determinants'. These could be the producers of the observed effects, or they could be simply the indicators, or sign posts, of a future outcome. In fact, social scientists outside structural equation modelling use many forms of determination which are not strictly causal, and I include here historical and structural analyses (Pötter and Blossfeld 2001). We should also accept a causal model which is probabilistic rather than determinist in nature (Goldthorpe 2001), although we would be unable to decide whether this worked because the world is actually non-determinist, or because it is too complicated to explain fully and so we allow for error. 'The teacher may give what appears to be the same lesson in exactly the same way in a second classroom, but the outcome of the second lesson may be quite different because some un-noted variables of the setting, or the class, or the individuals within the class, are sufficiently different to affect the outcomes' (Bassey 2001, p.7). However, rather than see this as something peculiar to education or social science we should recognise that this is the common form of causal modelling. Simple deterministic causation is rare in reality, where even physical 'laws' are actually generalisations from many differing observations. Water tends to boil at 100 degrees centigrade, but it depends on the atmosphere and the purity of water. I doubt if I, or any readers, have actually witnessed water boiling at precisely 100 degrees. The best we will have done is observed a tendency for what we call water to boil at near that point, and to have created a list of exceptions and conditions. Even such a simple law appears probabilistic, rather more than like the constant conjunction of strict determinism.

Time and causation

'It can be said to be axiomatic to any notion of causality that it only acts forward, that is, a cause must precede its effects in time' (Arjas 2001, p.60). In research, as in life, an easy assumption is sometimes made about the direction of causation that does not really stand up to scrutiny. This assumption is that one event can only be considered a cause of another if it occurs first, therefore if two variables are related then their temporal sequence defines which is the cause. For example, an analysis by Dolton et al. (1994) of data from the longitudinal Youth Cohort Study explained the labour market position of each participant in terms of their position in previous sweeps. Similarly, Gershuny and Marsh (1994) explained each participants' current employment position, as described by employment status, sector and occupational level, in terms of two main determinants, both preceding the current employment position - their initial characteristics and the accumulation of previous employment. They therefore adopted a 'recursive determination model', which was based on the causal chain approach advocated by Blau and Duncan (1967). Technical literature generally suggests methods for analysis of event histories, such as Proportional Hazard modelling, which try and explain occurrence variables in terms of prior events. In fact, a majority of studies use these unidirectional recursive models, such that the predictors or independent variables are themselves unaffected by the outcomes or dependent variables (Berry 1984). The approach was summed up in one study thus, 'what we do now becomes what we are, and what we are in part determines what we do next' (Gershuny and Marsh 1994, p. 69). In their analysis of the determinants of unemployment, variables were entered into the model in the order that they occurred historically, from parents' occupational class through initial education to the work details. The 'effect' of the earlier episodes was assumed to be present throughout the analysis but was found to diminish over time. In this way, the past is seen as affecting the present while both can affect the future, but the future cannot affect the present and the present cannot affect the past.

However, in many respects the assumption of unidirectional causation is unrealistic (Berry 1984). Causality is merely assumed to be time-determined (Hume 1962). The relationships between data which are seemingly in a temporal sequence are often reciprocal (Hagenaars 1990). Many educational decisions may be made in the expectation of a fairly remote future outcome, for example a decision to stay on at school at age 16 in order to become a barrister. Even as early a stage as picking courses of study for GCSE can be dependent on final career intentions (Roker 1991). Rational choices can allow people to jump towards attractive options rather than being 'pushed from behind' (Gambetta 1987). In addition, the direction of the arrow of causation is not at all clear even in well-established links between variables. For example, does a higher family income lead to a better education for the children, or can stress laid on future educational plans also lead to a need for a higher level of income? Does greater investment in training lead to company growth, or are richer companies more likely to spend money on training? Analyses using only a single equation model, such that A is predicted by B and C separately assume no relationship between the predictors. Regression models are now becoming more complex, with multi-equation rather than single equation models, which allow the predictors to influence each other, so that although A may be caused by B, both A and B may be caused by C. However, in theory at least, if two variables can be reciprocally related or even if their error terms are related, then a non-recursive model must be used, but this is seldom seen in social science outside econometrics (Berry 1984). One reason for this is that as it is not possible to deduce the direction of causation from a simple association, some non-recursive models cannot be meaningfully described.

Teleology is not a respected phenomenon, and intentions are usually seen as being causes from the future. An opinion poll of researchers would probably discover very little support for the notion of 'backwards' causation. However, a variety of situations and puzzles have been devised which expose how common the teleological explanation of events actually is.

A simple example, from a whole range, of our habit of using backwards causation is as follows. You are appearing in a TV quiz, and are presented with three closed boxes. One box contains a prize and the other two are empty. You are allowed a free guess. If you pick the box with the prize, you win it. You select one of the boxes (box A for example). The compere, who knows the contents of each box, then deliberately opens an empty box (box C for example) and shows it to you. The compere then gives you a chance to change your mind. Do you now have any reason to pick another box (box B in this example) or to stick with your original choice? Put another way, what have you learnt from the opening of box C?

Many readers will argue that they have no reason to change their mind, but that they now have an improved chance of winning whether they stick or pick box B. People tend to claim that whereas they had started with odds of 1 in 3, they now face odds of 1 in 2. But even being tempted by this 'analysis' displays a belief in backwards causation. Nothing that the compere has done in opening the box can change the position of the prize or, therefore, the odds of winning. When the game started you had odds of success of 1 in 3 (with box A). The prize was twice as likely to be in one of the other two boxes, even though one of the other two boxes must be empty. The fact that you now know which of the other two boxes is empty changes little. The prize is still one third likely to be in box A, and two thirds likely to be in one of the other two (which is simplified now to box B). Picking box B is twice as likely to be successful as picking box A. To consider otherwise implies that opening box C can have an effect on the actual position of the prize (cf. the problem of Schrödinger's cat).

Since the entire notion of causation has no solid evidence-base, but is chiefly a habit of mind according to Hume and others, the fact that reverse time causation is also a habit of mind gives it a very similar philosophical and scientific status as the more usual causal models. In fact, in a full determinist model of events it makes as much sense for time to run backwards as forwards (it would, presumably, not be possible to tell the difference anyway).

Plausibility

In evaluating whether a possible theory makes sense, De Vaus (2001) suggests in addition to explaining the co-variation and time sequence, and being plausible, that the proposed dependent variable must be capable of change. While the sex of the student could affect the outcome of an assessment, the reverse could not be true. Sex would be unchanged by the assessment. In fact, we can go further than saying the dependent variable must be capable of change to it can be changed by the independent variable. If there is a relationship between the level of poverty among sixteen-year-olds and their examination results, then the only causal model that makes sense in the short-term is one where poverty leads to examination results.

A possible characteristic of a good causal model is an explanatory process or theory that takes these restrictions on plausibility into account. If causation is a generative process then something must be added to the statistical association between an intervention and an outcome for the model to be convincing. The cause must be tied to some process that generates the effect. The standard example is the clear relationship between smoking and lung cancer. The statistical conjunction and the observations from laboratory trials were elucidated by the isolation of carcinogens in the smoke, the pathological evidence from diseased lungs and so on. From this complex interplay of studies and datasets emerges an explanatory theory - the kind of theory that generates further testable propositions. This is the key role for theory-building in educational research.

Interventions

Another way of viewing causation is via the effect of an intervention. If causes are not susceptible to direct observation, but what they 'cause' is effects, then at least those effects must be observable. We therefore follow the principle of 'no causation without manipulation', and attempt to mimic the classic Fisher experiment. This is the approach used by Pavlov in so far as classical conditioning involved a causal model of learning and extinction. Koch used a very similar approach of intervening and removal to show causation in infections (Cox and Wermuth 2001). Unfortunately in a social science where the subject of study is people we cannot usually expose the same people both to the treatment and not, as might be possible by using two near identical cases in Physics for example. We therefore use statistical approaches (including random allocation to groups) to overcome this limitation. And this, of course, may be why probabilistic models of causation emerge. They may reflect, not the reality of the study, but the limitation of the experimental design.

These same statistical procedures are now more widely used where an intervention is not even attempted but is replaced by further statistical controls such as weighting. There remains fundamental disagreement over the validity of these approaches (McKim and Turner 1997). Prediction, based on correlation alone, does not depend on a causal relationship, nor does it necessarily exhibit causation. This is true however impressive the prediction is - we may accurately predict the severity of a fire from the number of fire engines attending without attributing the cause of the fire to the engines (De Vaus 2001). Day always precedes night and so could explain 100% of the variance in a regression model, but that would not make it the cause. What we need to do is be creative in the consideration of alternative explanations (and then test these if practical). In fact, it is common to encounter the 'fallacy of affirming the consequent' in social science (and I plead occasionally guilty to this one). The fallacy argues that if A is true then B will follow. Then if B appears it means that A is true. While seductive there is no logic to this argument unless it starts more strongly with 'only if'. Otherwise exactly the same argument can be made with Z (or anything else) substituted for A.

Causation and the law

It is interesting also to consider the legal position of causation. Evidence for causation has been presented in many legal cases - a common theme in occupational medicine for example (Rom 1992). Bradford-Hill's criteria for identifying causation are widely applied (Bradford-Hill 1996). These are a temporal relationship, specificity, biological [i.e. mechanical] plausibility, and coherence. When these have been put to the test in law, the US Supreme Court has ruled (e.g. in *Daubert versus Merrill Dow*, 1993) that experts seeking to prove causation (of the toxic effect of a chemical) have to establish a greater than 50% probability based on common ground, generally agreed techniques and evidence from peer-reviewed publications. The expert does not therefore need to rely on scientific proof or certainty, or even intervention trials necessarily. Helpful, but not all essential, characteristics of the purported causal relationship are:

- confirmation of association in different studies, researchers, populations, and methods;
- frequency of association compared to frequency of either alone;
- exposure to the factor (for the individual or the population) before the onset of the disease;
- predictive relationship between the factor and frequency (biological gradient);
- isolation of factor and use as intervention to create disease;
- coherence with previous knowledge, and plausibility;
- workable previously agreed analogy;
- reduction in disease after removal of factor.

Conclusion

Causes are particularly relevant in a climate of evidence-informed policy-making and practice for at least two reasons. Causes are really only susceptible to testing by intervening and measuring, the technique of randomised controlled trials and related designs which form the basis for the Campbell/Cochrane syntheses. In addition, in order to determine what works in any given situation the intervention must be proposed first (for there are an infinite number of potential interventions). While this creative phase of a study can be, and has been, inspired or serendipitous, the closest we have to a technique for generating such ideas is to try and understand why things work. This is the role of theory - not banner-waving grand theory, but attempts to provide explanations for observed phenomena in ways that are fruitful and actually testable. Thus, I would argue that a useful cause would have the characteristics of all of the models proposed above. It would involve conjunction (relatively stable association), a measurable effect from the intervention, and at least a tentative theoretical explanation. However, the more standard notions (Pötter and Blossfeld 2001) of cause and effect having spatial and temporal contiguity, constant conjunction, and temporal succession have all been brought into doubt in this discussion. Causation is a difficult concept for use in social science for a variety of reasons. It is difficult to define, requiring some form of agency or productive force to be meaningful, although this force is never observed and merely induced from patterns of ordered events (Blalock 1964). Some models may allow causes to operate at a distance, or even require causes to come after effects, and some may wish causes to be only probabilistically associated with effects. All of these are possible, and none of them disturb the conclusions just drawn.

Having resolved this, in practical terms cause/effect is still difficult to isolate. Given the design bias, and sampling and measurement errors in all our work we may end up with estimates, catalysts, determinants, or even 'fuzzy generalisations' rather than simple, almost mechanical, cause and effect models. While perhaps disappointing to some, this is actually inevitable. Our role as researchers is to minimise the bias and the sampling and measurement errors. Statistics, as popularly conceived, can only help with the least important of these - the sampling error (and while statistical procedures describe ideal situations, social scientists conduct their studies in, and are influenced by, real-life social settings, Gephart 1988). Overcoming the rest of the error, the bulk of it in any design, is to do with rigour (and whatever that is will have to be the subject of a further paper). Rigour would transcend any specific approach or method. It is certainly not the prerogative of experiments (whose importance lies in the intervention only). The current paucity of experiments in social science is therefore not an excuse to evade the need for rigour. The same situation is faced in many fields such as archaeology, palaeontology and astronomy, and for more solidly practical reasons perhaps. Even cutting-edge sciences such as molecular genetics use relatively few genuine experimental designs (although the routine benchwork creates controls as a matter of course). The same situation applies in a range of scientific and quasi-scientific settings (Collins and Pinch 1993). 'Physics envy' among social scientists is misplaced, and there remain many useful strategies of a non-experimental nature that enable us to increase our confidence in perceived causal relationships (such as selection modelling, or longitudinal studies combined with triangulation of methods, see Johnson 2001). Further ideas about the design of research to test causal models will also have to be the subject of a further paper.

Notes

The work involved in preparing this paper was undertaken as part of a project funded by the ESRC (grant number L139251106), and a draft was presented at the first internal Teaching and Learning Programme Conference, London 24th May 2001, and Sheffield 7th June 2001.

I would like to thank Harry Collins, Rob Evans, Laurence Moore and James Tooley for their helpful suggestions.

References

- Arjas, E. (2001) Causal analysis and statistics: a social sciences perspective, *European Sociological Review*, 17, 1, 59-64
- Ayer, A. (1972) *Russell*, London: Fontana
- Badger, D., Nursten, J., Williams, P. and Woodward, M. (2000) Should all literature reviews be systematic?, *Evaluation and Research in Education*, 14, 3&4, 220-230
- Bassey, M. (2001) A solution to the problem of generalisation in educational research: fuzzy prediction, *Oxford Review of Education*, 27, 1, 5-22
- Berry, W. (1984) *Nonrecursive Causal Models*, London: Sage
- Blalock, H. (1964) *Causal Inferences in Nonexperimental Research*, Chapel Hill: University of North Carolina Press
- Blau, P. and Duncan, O. (1967) *The American Occupational Structure*, London: John Wiley
- Bradford-Hill, A. (1966) The environment and disease: association or causation?, *Proceedings of the Royal Society of Medicine*, 58, 285
- Brighton, M. (2000) Making our measurements count, *Evaluation and Research in Education*, 14, 3&4, 124-135
- Collins, H. and Pinch, T. (1993) *The Golem: what you should know about science*, Cambridge: Cambridge University Press
- Cox, D. and Wermuth, N. (2001) Some statistical aspects of causality, *European Sociological Review*, 17, 1, 65-74
- Daubert versus Merrill Dow (1993) 509 US 579, 113 S. Ct 2786
- De Vaus, D. (2001) *Research design in social research*, London: Sage
- Dolton, P., Makepeace, G. and Treble, J. (1994) Measuring the effects of training in the Youth Cohort Study, in McNabb, R. and Whitfield, K. (Eds.) *The Market for Training*, Aldershot: Avebury, p. 195
- Field, A. and Wilkinson, L. (2001) Getting your numbers wrong, *The Psychologist*, 14, 6, 316
- Fisher, R. (1935) *The design of experiments*, Edinburgh: Oliver and Boyd
- Gambetta, D. (1987) *Were they pushed or did they jump? Individual decision mechanisms in education*, London: Cambridge University Press
- Garrison, R. (1993) Mises and his methods, pp.102-117 in Herbener, J. (Ed.) *The meaning of Ludwig von Mises: contributions in economics, sociology, epistemology, and political philosophy*, Boston: Kluwer Academic Publishers
- Gershuny, J. and Marsh, C. (1994) Unemployment in Work Histories, p.66, in Gallie, D., Marsh, C. and Vogler, C. (Eds.) *Social Change and the Experience of Unemployment*, Oxford: Oxford University Press
- Gleick, J. (1988) *Chaos*, London: Heinemann
- Goldthorpe, J. (2001) Causation, statistics, and sociology, *European Sociological Review*, 17, 1, 1-20
- Goodman, N. (1973) *Fact, fiction and forecast*, New York: Bobs-Merrill
- Gorard, S. (1998) *The middle way*, BERA Internet Conference, www.scre.ac.uk/bera/debate/index.html
- Gorard, S. (2001) *Quantitative methods in educational research: the role of numbers made easy*,

- London: Continuum
- Gorard, S., Rees, G. and Jephcote, M. (1998) The role of contour lines in school improvement, *Research Intelligence*, 66, 30-31
- Hagenaars, J. (1990) *Categorical Longitudinal Data: Log-linear, panel, trend and cohort analysis*, London: Sage
- Hendry, D. and Mizon, G. (1999) *The pervasiveness of Granger causality in econometrics*, Nuffield College Oxford, (mimeo)
- Hillage, J., Pearson, R., Anderson, A. and Tamkin, P. (1998) *Excellence in research on schools*, Sudbury, DfEE Publications
- Hume, D. (1962) *On Human Nature and the Understanding*, New York: Collier
- Johnson, B. (2001) Towards a new classification of nonexperimental quantitative research, *Educational Researcher*, 30, 2, 3-14
- Kuhn, T. (1970) *The structure of scientific revolutions*, Chicago: University of Chicago Press
- Lakatos, I. (1978) *The methodology of scientific research programmes*, Cambridge: Cambridge University Press
- Magee, B. (1972) *Popper*, London: Fontana
- McKim, V. and Turner, S. (1997) *Causality in crisis? Statistical methods and the search for causal knowledge in the social sciences*, Indiana: University of Notre Dame Press
- Pötter, U. and Blossfeld, H. (2001) Causal inference from series of events, *European Sociological Review*, 17, 1, 21-32
- Roberts, I. (2000) Randomised trials or the test of time?: The story of human albumin administration, *Evaluation and Research in Education*, 14, 3&4, 231-236
- Roker, D. (1991) *Gaining the Edge: The education, training and employment of young people in private school*, London: City University
- Rom, W. (1992) *Causation*, in *Textbook of Environmental and Occupational Medicine*
- Salmon, W. (1998) *Causality and explanation*, New York: Oxford University Press
- Thouless, R. (1974) *Straight and crooked thinking*, London: Pan