Brief Note

Introducing the Language of Causal Analysis

Michael Joffe
Department of Epidemiology and Biostatistics, Imperial College London, Norfolk Place, London W2 1PG, United Kingdom.
m.joffe@imperial.ac.uk

The paper "A simple method for causal analysis of return on IT investment" [1] is welcome. The authors have taken an important step forward by introducing an explicit assessment of causation into their analysis. Traditionally, researchers have been timid in their approach to causality, tending to avoid causal language. This is misplaced. It is true that in the statistical/empirical sciences such as epidemiology, associations are all that we can ***observe***, and that these are not necessarily causal. But it is impossible to make sense of them without invoking causal relations – for example, a long-established rule of thumb in multivariate analysis has been to adjust for covariates ***that are not on the causal pathway*** of interest. Besides, in most cases, the importance of an association – scientifically and also practically – is mainly derived from what it tells us about a causal process. And yet, the practice has typically been for explicit causal thinking to occur not in the design of the study or the set-up of the analysis, but only afterwards, in assessing the findings; or at least, papers are written in this way, the causal inference being left until the Discussion section of a paper, where it is "smuggled in" [2], rather than being part of the Methods section. It is more appropriate to develop and use causal language in a rigorous fashion.

The paper employs the strategy of looking at the three main variables of interest, noting that six time orders are ***logically*** possible but not all are plausible in the real world, and then discarding those that are implausible. First, the options where use of IT precedes IT investment are discarded, whereas those where IT use follows investment are retained. In the second stage, the authors also exclude the possibility in which IT use occurs after IT investment, but where this is mediated through revenue growth, i.e., IT investment → revenue growth → IT use.

The paper uses substance-specific knowledge of causal relationships to make an *a priori* judgement as to whether a particular set of causal relationships should be a candidate hypothesis for investigation and estimation. Put like this, in the abstract, it sounds suspect to someone who is used to the mind-set that inferring relationships from data is the only legitimate procedure, and that any kind of *a priori* imposition of causal structure is unwarranted. Clearly this strategy is valid only to the extent that the *a priori* judgement is sound; it needs to be explicitly justified in terms of previous evidence and substantive knowledge, and/or (as here) by plausibility in the light of how the world works.

The second-stage argument rests on the validity of the following propositions:(i) there is no plausible way in which IT investment would lead to revenue growth, apart from through IT use, and (ii) there is no reason to suppose that revenue growth would of itself lead to IT use in the absence of IT investment. The first-stage argument relies on the idea that IT use follows investment and not *vice versa*, but this is open to the objection that there could in addition be a pathway in the opposite direction between these two variables, because the more IT is used, the more likely the organisation is to invest in more of it. This two-way causation is an example of reinforcing (or "positive") feedback, and one has to be careful to consider such possible complications. In this particular case, it is likely that the "reverse" pathway would be weaker than the one that was considered plausible, but ideally one would like to see evidence on this.

If the authors' judgements are accepted, two possibilities remain, both having IT investment followed by IT use. In one alternative, IT use increases productivity and this leads to revenue growth. In the other, IT investment and use result from revenue growth, because it is the increased revenue that makes the investment possible. These can conveniently be represented graphically:

(1) IT investment ⟶ IT use ⟶ revenue growth
(3) revenue growth ⟶ IT investment ⟶ IT use

In principle both could occur, which would again be an example of reinforcing feedback; however, the analysis presented finds evidence only for the first pathway. This does not rule out the possibility that revenue growth may lead to IT investment and thence to IT use – for example, the time lag could be longer than the duration of the study.

As with any single study, definitive conclusions cannot be drawn from this paper on its own. More sophisticated analysis is possible for future studies, both statistically and conceptually, as the authors themselves state. Statistically, it might be possible to use the successive years in the dataset to refine the analysis, using time series methods to document the order in which changes occur in one variable and then another, so providing further time-order evidence for causal direction. This would be more convincing with a dataset that covers a longer period.

Conceptually, one needs to consider whether these diagrams may be incomplete, in that additional causal pathways may need to be included. We have already touched on the possibility of reverse causation, in this case raising the issue of feedback, but in other circumstances the reversed direction may be an alternative to the favoured hypothesis, rather than an addition to it. For example, if blood cholesterol levels are observed to be lower among people who later develop colorectal cancer, does this mean that low cholesterol predisposes to (causes) this disease? Or, that the cancer reduces the blood concentration in its early, preclinical, stages? It is unlikely to be both – and important to know, because if it were the former it would affect the judgement whether having a low cholesterol level is wholly a good thing or not, which in turn would affect whether treatment would be recommended.

In addition, one needs to consider not just other arrows between the existing variables, but additional variables that may be relevant to the observed association – the well-recognised but not always well-handled issue of confounding. For example, if a particular occupational group is found to have an unusually high risk of lung cancer, is this merely because smoking is particularly common among them or because they are exposed to a carcinogenic agent? Again the implications are important. And, it is possible that both could be occurring: for example, less well-off people have relatively high rates of lung cancer, as well as higher smoking rates, but only a proportion of their higher risk is attributable to the effect of smoking. This can be expressed by saying that there is a causal pathway from low socioeconomic status to lung cancer that is mediated by smoking, and one or more other pathways that are mediated by other factors. It is possible to estimate the proportions of the increased risk that are and are not mediated by smoking [3, 4]. In the current context, there could for example be one or more unmeasured common causes of IT use and of revenue growth.

A natural way to express such concepts is by using diagrams, with variables that have arrows drawn between them to represent causation, the arrow head indicating its direction. Their use has grown in recent years, for example in epidemiology where they are described as DAGs (directed acyclic graphs, meaning that the arrows have directions, and that there are no causal loops in the system) [5–9]. Confounding can readily be depicted as one or more additional variables, with arrows both to the putative causal variable and to the outcome variable (see Figure 1). Such diagrams are intuitive to draw and to interpret – which may introduce a danger of misuse, as there are quite strict rules for their proper use. These have been elaborated during the past twenty years or so, in the branch of statistics known as graphical models [10, 11]. It has been shown that graphs are equivalent to equations, and when properly used they have certain advantages. This is for two reasons: they separate the causal connection from the functional form, so that one can state that X causes Y without specifying that the association is, for example, linear. Secondly, in some complicated situations they are actually superior in handling issues like confounding, because the rules for their proper use perform better than the rules of thumb mentioned earlier [5, 7].

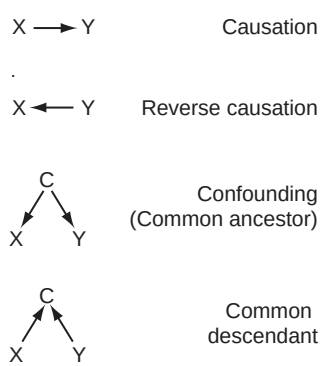

Figure 1. The four types of causal process that can bring about a statistical association.

A key insight in this literature is that causal and associational languages are different [2, 12]. Causal concepts include confounding, because knowledge of causal direction is needed to interpret it, as well as influence, effect, explanation, attribution and intervention. Associational concepts are those that apply to the relationships observed in the data, studied by statistical methods, and include probability, correlation, regression, likelihood, and "controlling for". The latter group can help provide evidence for the former group, but this step requires also a different type of input, one that deals with plausibility in terms of mechanism – ***how*** the change operates, e.g., that IT use could lead to revenue growth by making the firm more efficient – using substantive knowledge. This is because "a causal relationship is one that has a ***mechanism*** that by its operation ***makes a difference***" [13], and the associational language only deals with the difference-making aspect. The two approaches are complementary: in the discovery of causal relationships, either mechanism or difference-making can come first. However, the evidence is not very secure when only one is present, and only becomes robust when these two types of evidence correspond well with one another.

Representing causal relationships as diagrams helps to clarify thinking. In particular, consider the situation discussed by Bradford Hill in his classic paper [14]: an association has been observed that is robust, i.e., statistical analysis has shown that it is highly unlikely to have arisen purely by coincidence. There are only four types of causal process that can bring about such a situation (see Figure 1), three of which we have already considered: "direct" causation, reverse causation and confounding. The latter is shown as having a common parent – in causal diagram language, a "parent" is a variable that has a direct causal influence on another variable; obviously multiple parents can exist (because of multiple causation). More generally, if a diagram is more complicated, the influence can be mediated through one or more intermediate variables, in which case the term "ancestor" is used. The fourth type of causal relationship occurs when the association is brought about instead by a common descendant, typically a selection effect ("Berkson bias"); discussion of this is beyond the scope of the current paper [5].

The central problem in inferring causation from associational evidence is that it is impossible to be certain what pathways are possible for any particular causal system – knowledge is always liable to be incomplete. This is true whether one uses the traditional methods of statistics and epidemiology or whether one adopts the newer diagrammatic approach. The latter enables all the pathways that are known or suspected to be depicted, and therefore has the advantage of explicitness, and it may be easier to spot missing pathways and/or variables. Nevertheless, it is unhelpful to exaggerate the degree of uncertainty, as in RA Fisher's contention that there are "innumerable" ways to explain the association between cigarette smoking and lung cancer, other than direct causation. In this instance, time order rules out reverse causation, and the strength of the observed association requires the existence of one or more confounding variables with an ***extremely*** strong impact both on smoking habit and on lung cancer incidence. It seems highly unlikely that such a factor could be present without quickly becoming obvious.

As the authors state, an important concept in the construction and analysis of causal diagrams is conditional independence. Returning to the example of IT investment and use, and revenue: we can express the simple diagram (1) above as "revenue growth is independent of IT investment, given IT use". This means that IT investment has no effect on revenue growth except for that mediated through IT use: a particular level of IT use is what affects revenue growth, and the fact that this resulted from IT investment does not affect this link in the chain. This is sometimes expressed in terms of what is known: that if you know the value of IT use, then discovering the value of IT investment tells you nothing more – but the property of conditional independence depends on the relationships in the real world, not on our knowledge of it. (The same conditional independence relationship would exist if the direction of causation were reversed, i.e., if one of the possibilities that were rejected at the first stage was true: revenue growth → IT use → IT investment, or even if IT use caused both revenue growth and IT investment. This is one of the properties established in the mathematics of graphical models.) The standard notation is:

$$P(r \mid u, i) = P(r \mid u)$$

where r, u and i respectively represent revenue growth, IT use and IT investment. This equation is read as follows: "the probability of revenue growth given the levels of IT use and IT investment is the same as the probability of revenue growth given the level of IT use alone". It is useful to be familiar with this notation before trying to deepen one's understanding of the use of causal diagrams.

A final point relates to surrogate measures. The authors' measure of IT use was the number of patient records, but they also suggest other possibilities such as the number of calls to the system or of clinicians using the system. They state that the choice of which to employ introduces some level of subjectivity. However, it is possible to use a number of measures to provide estimates for a single underlying concept, in which case they are termed latent variables. This approach has long been widespread in the social sciences, and combines neatly with graphical representation. Traditionally, the method of analysis has involved factor analysis, principal components analysis and/or structural equation models, but others are possible [15].

REFERENCES

[1] Alemi F, Zargoush M, Oakes JL, Edrees H. A simple method for causal analysis of return on IT investment. *Journal of Healthcare Engineering,* 2011, 2(1), 43-53.

[2] Pearl J. Causal inference in the health sciences: a conceptual introduction. *Health services and outcomes research methodology* 2002; **2**: 189–220.

[3] Petersen ML, Sinisi SE, van der Laan MJ. Estimation of direct causal effects. *Epidemiology* 2006; **17(3)**: 276–84.

[4] Goetgeluk S, Vansteelandt S, Goetghebeur E. Estimation of controlled direct effects. *J R Statist Soc B* 2009; **70(5)**: 1049–66.

[5] Glymour MM, Greenland S. Causal diagrams. In: Rothman KJ, Greenland S, Lash TL (eds.) *Modern epidemiology*. Wolters Kluwer/Lippincott Williams & Wilkins, Philadelphia, 2008.

[6] Robins JM. Data, design, and background knowledge in etiologic inference. *Epidemiology* 2001; **11(3)**: 313–20.

[7] Howards PP, Schisterman EF, Heagerty PJ. Potential confounding by exposure history and prior outcomes – an example from perinatal epidemiology. *Epidemiology* 2007; **18(5)**: 544–51.

[8] Hogan JW. Bringing causal models into the mainstream. *Epidemiology* 2009; **20(3)**: 431–32.

[9] Joffe M, Mindell J. Complex causal process diagrams for analyzing the health impacts of policy interventions. *Am J Public Health* 2006; **96**: 473–79.

[10] Pearl J. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. San Francisco, Morgan Kaufman Publishers, Inc., 1988.

[11] Lauritzen SL. *Graphical models*. Oxford, Oxford University Press, 2006.

[12] Pearl J. *Causality: models, reasoning and inference*. New York, Cambridge University Press, 2000.

[13] Joffe M. Causality and evidence discovery in epidemiology. In: Weber M (ed.) *Proceedings of ESF workshops* [provisional title], Springer, *accepted for publication*.

[14] Hill AB. The environment and disease: association or causation? *Proc Royal Soc Med* 1965; **58**: 295–300.

[15] Bishop CM. Latent variable models. In: Jordan MI (ed.), *Learning in Graphical Models*, MIT Press, 1999, pp 371–403. Available at http://research.microsoft.com/en-us/um/people/cmbishop/downloads/bishop-latent-erice-99.pdf [accessed 31 May 2010].