

The perception of causation

Eric Blair

28 November 2003

So here's the structure of just about every academic paper: author asks a question in the way of 'hey, I wonder if A causes B?' and then goes out and gathers some data, and then checks to see whether the data verifies that A and B are related. The rest of today's pontifications will pile on caveats to this simple structure.

Statistics only disprove The first is that statistical methods on the observation side is aimed only at falsifying claims, not verifying them. The correct wording of a typical conclusion would be: 'The study failed to reject the possibility that A and B are related.' The media will clean this sentence up to say, 'The study found a link between A and B.' More forceful, less passive, less true.

Causation Next, there are two motivations for asking about A and B: observation and persuasion. Our academics like to imply that they're just innocently inquiring about the world around them, but people get married to their ideas quickly—especially in the social sciences, where there's often a gut political belief that the academic would like to back up. Further, the goal of all persuasive essays is not to say that A and B are linked, but that A causes B. We readers don't entirely mind our academics writing causal-persuasion-oriented papers instead of observational papers, because causation is interesting and observation alone isn't.

If I had to point to one thing that distinguishes humans from stuff, it's the perception of causality. It is impossible for humans to think without using a causal framework, and it's impossible for machines to think causally. You can't do a statistical derivation to prove objectively that one thing causes another, since correlation is never causation. Further, if you have two things which we've decided cause another, you can't do anything to show that thing one was more of a cause than thing two, even though we humans often care deeply about that very question.

An academic paper really, really is an artistic expression—even more than painting or self-immolating performance art—because the paper's intent is to make the reader perceive the causal functioning of the world (that inherently human thing) as the author perceives it. The funny thing is that the persuasive support for the communication of perception lies in statistics, which can provide

evidence for or against causation, but not causation itself. You get papers that are 100% machinery, but whose intent (if they haven't forgotten it) is to persuade the reader.

Linear models The primary tool for the testing of an A-is-related-to-B hypothesis is the linear regression, which is based on lots of assumptions that nobody ever checks. Notably, most of these regressions are about nonlinear systems. I mean, really guys: it's a *linear* regression because all of the methods embodied in the process are from *linear* algebra, and yet people will run a frigging regression on anything, without paying a second of heed to whether the thing they're modeling is really even vaguely linear. And throwing in an 'oh, it's log-linear, so just take the log of everything and you're hunky-dory' just doesn't cut it unless you have more justification for log-linearity than 'oh, it looks all loggy.'

In sum, our dependence on linear regression makes the world a worse place, because it is often not applicable for observational purposes, and at the same time says nothing about causation. Going away from the linear regression form, you've got two statistical options. The first is to make things a lot more complex. E.g., do everything as a maximum likelihood estimation of some weird function that you dreamed up. This allows you to get quantitative values where there were none before, which is observationally useful, but is even less persuasive than the simpler regression.

The other option is to get much simpler than the linear regression: just describe what you see. Gather your variables, calculate their means and variances, and tell us whether the mean is different from zero. A good description of the situation with a few connecting paragraphs is often a much more persuasive argument than a table of forced regressions. I wish there were more papers like that.

Modelling of nonlinear systems That said, let's get back to the subject of me. I draw up models, many computer-based, which are intended to describe a situation. A good model will do well on both of the above fronts: on the observational level, it comes up with data which is comparable to the data we cull from reality, and on the causal level, it should embody the causal processes that seem similar to that of the real world. I.e., a model which always gets the right answers but does so in an intuitively displeasing fashion will not make it as a subjectively good model.

So the problem is that the folks at [name of center at institution], of which I am including myself, model situations where the causal process is based on the interaction of lots and lots of people. Instead of, 'Oh, variable A rises, which causes variable B to fall,' we have, 'People have a set of heterogeneous characteristics with distribution A, and if they interact in a specific fashion which seems basically descriptive of reality, then B results. We can predict how changes in A will affect B.' [That includes both my all-computer work on migration and my computerless paper-and-pencil analysis of conviviality.]

It's observationally great, but dissatisfying with respect to the perception of causation. You don't feel smarter after reading the paper. You have no causal factoid to put in your back pocket and pull out at cocktail parties. But if you're looking for predictions about a nonlinear system, then you're much better off with one of these computer models than an invalid but easily comprehensible linear-regression based causal story.

The MPU It's a common complaint that modern academics are wankers because they often break down their papers into the minimal publishable unit (MPU). Back in the day, an author would expound for a hundred pages on a topic, and may produce one paper which summarizes his/her/its findings about the world and its workings. Nowadays, the lament goes, the cynical academic would break that down into a thousand short papers, thus multiplying his/her/its citation count by a thousand. Papers like Milgrom & Robert's 1980 magnum opus on auction theory just aren't getting published anymore.

I think this is true, but the blame here is in the wrong place—yes, papers are more likely to be exactly one MPU long, but that is because the readers demand it to be so. Referees don't have time to read a hundred technical, meandering pages; they often don't have all that much interest in the topic that's been thrown on their desk. They don't feel smarter after reading the darn thing.

A paper which does not embody a simplistic causal story will fail. This is a horrible thing for the study of systems where there are no simplistic but true causal stories to be had.

Our hubris People like their gods to have faces, and like them to generally be kind of like people, but more powerful. Stories about deities have a causal framework much like that of stories about people; where they don't fit the normal human form, thousands of pages get written on why the omnipotent one is only behaving in an unhuman way on the surface. Similarly, academics expect the workings of the world to function in a basically comprehensible manner, and get annoyed when the world doesn't.

So my lament is that the models of [name of center at institution] throw such academics for a loop: my laptop can keep track of more numbers than any human, and it's pretty darn easy for me to have it do things which are conceptually simple, but are beyond easy A-causes-B explanation. They are observationally sound—and often better than simpler stories, but are not causally persuasive unless people are willing to accept causal stories which are complex beyond the workings of their own brains.