Single Case Causes: What is Evidence and Why

Nancy Cartwright, Durham University

CHESS Working Paper No. 2015-02
Durham University
March 2015

CENTRE FOR HUMANITIES
CHESS
ENGAGING SCIENCE AND SOCIETY

CHESS working paper (Online) ISSN 2053-2660

Single Case Causes: What is Evidence and Why

Nancy Cartwright

Nancy Cartwright
Department of Philosophy
Durham University
50 Old Elvet
Durham DH1 3HN
E-mail: nancy.cartwright <at> durham.ac.uk

1. The problem setting

How do we establish singular causal claims? It seems we do this all the time, from courtrooms to cloud chambers. Nevertheless, there is a strong lobby in the evidence-based medicine and policy movements that argues that we cannot make reliable causal judgments about single cases in these areas. So we cannot tell whether a policy or treatment 'worked' for any specific individual. The reason often seems rooted in the idea that a singular causal claim, if not equivalent to, at least implies, a counterfactual claim: c caused e → e would not have occurred had c not occurred.[1] Couple this with the idea that we cannot have compelling evidence about what would have happened in any actual case had things been different and you end up with serious doubts about the possibility of warranting singular causal claims.

There are obvious problems with this line of thought. First, we all know familiar counterexamples to the simple unadorned implication from singular causal claim to counterfactual. Second, there are reasons to think we can have compelling evidence about what would have happened if things had been different in specific ways, for instance by warranting a universal law that implies that –e will be always be consequent on –c. Third, on standard accounts of how positive randomized-controlled-trial (RCT) results can legitimate causal conclusions, what gets established is a result about causal relations in single individuals: at least some individuals in the treatment group were caused by the treatment to improve. But the RCT is the gold standard in evidence-based policy and medicine, and among the very folks who object that we cannot establish singular causal claims in various health and social policy domains. I suppose then that the thought is that although we can establish causal claims about single individuals, these are what we might call *anonymous* claims*:* c caused e in somebody (we know not whom). The worry is whether we can have compelling evidence about *individualized* claims: c caused e in this specific individual.

This worry matters when it comes to evaluating policies and treatments. Practitioners can have very strong feelings that a treatment they employed for a particular individual was the cause of improvement in that individual where the treatment is not the kind of intervention

[1] This idea is widespread not just in the evidence-based medicine and policy movements but also among philosophers. Tim Maudlin (2002, 143) for instance claims that, "When we think we know a cause of some event, we typically assent to the corresponding Hume counterfactual", where the 'Hume counterfactual' corresponding to the claim that c caused e is simply the counterfactual 'e would not have occurred had c not occurred'.

that can be tested in an RCT, nor that we have well established theories about that allow us to conclude that it should work for individuals like this one. It seems equally unsatisfactory simply to dismiss the practitioners' views as it is just to go along with them. Nor do we need to, I shall argue here. I shall do so by providing a catalogue of kinds of evidence that are often available that *can* speak in favour of singular causal claims, and do so without routing through a counterfactual on the way.

2. Some categories of evidence for 'c caused e in individual j'

The first distinction that can help provide a useful categorisation for types of evidence for singular causal claims to draw is that between direct and indirect evidence:

- ***Direct:*** Evidence that looks at aspects of the putative causal relationship itself to see if it holds.
- ***Indirect:*** Evidence that looks at features outside the putative causal relationship that bear on the existence of this relationship.

Indirect. The prominent kind of indirect evidence is evidence that helps eliminate alternatives. If e occurred in j, and anything other than c has been ruled out as a cause of e in j's case, then c must have done it. This is what Alexander Bird (2010, 345) has called 'Holmesian inference' because of the famous Holmes remark that when all the other possibilities have been eliminated, what remains must be responsible even if improbable.

It is important to make a simple point at the start. I aim to lay out a catalogue of kinds of evidence that---if true---can speak for or against singular causal claims. How compelling that evidence is will depend on:

- how strong the link, if any, is between the evidence and the conclusion,
- how sure we can be about the strength of this link,
- how warranted we are in taking the evidence claim to be true.

All three of these are hostages to ignorance… which is always the case when we try to draw conclusions from our evidence. In any particular case we may not be all that sure about the other factors that need to be in place to forge a strong link between our evidence claim and our conclusion, we may worry whether what we see as a link really is one and we may not be all that sure about the evidence claim itself. The elimination of alternatives is a special case

where the link is known to be strong: If we have eliminated alternatives then the conclusion follows without the need of any further assumptions. But, as always, we still face the problem of how sure we can be of the evidence claim---have we really succeeded in eliminating all alternatives? No matter what kind of evidence claim we are dealing with, it is a rare case when we are sure our evidence claims are true and we are sure how strong our links are or even if they are links at all. That's why, when it comes to evidence, the more the better.

Direct. I have identified at least four different kinds of direct evidence possible for the individualised singular causal claim that c caused e in j:

1. The character of the effect: Does e occur at the time, in the manner and of the size to be expected had c caused it? (For those who are familiar with his famous paper on symptoms of causality, Bradford Hill (1965) endorses this type of evidence.)
2. Symptoms of causation: Not symptoms that c occurred but symptoms that c caused the outcome, side effects that could be expected had c operated to produce e. This kind of inference is becoming more and more familiar as people become more and more skilled at drawing inferences from 'big data'. As Suzy Moat puts it "People leave this large amount of data behind as a by-product of simply carrying on with their lives". Clever users of big data can reconstruct a great deal about our individual lives from the patterns they find there.[2]
3. Presence of requisite support factors (moderator variables): Was everything in place that needed to be in order for c to produce e?
4. Presence of expectable intermediate steps (mediator variables): Were the right kinds of intermediate stages present?

3. A diagrammatic example

Let me illustrate with one of those diagrammatic examples we philosophers like, this one constructed from my simple-minded account of how an emetic works.

Imagine that yesterday I inadvertently consumed a very harmful poison. Luckily I realised I had done so and thereafter swallowed a strong emetic. I vomited violently and have subse-

[2] At a *Spaces of Evidence* conference, Goldsmiths, University of London, 26 Sept 2014. See (Moat et al. 2014).

quently not suffered any serious symptoms of poisoning. I praise the emetic: It saved me! What evidence could I have for that?

- Elimination of alternatives: There are very low survival rates with this poison. So it is not likely my survival was spontaneous. And there's nothing special about me that would otherwise explain my survival having consumed the poison. I don't have an exceptional body mass, I hadn't been getting slowly acclimatised to this poison by earlier smaller doses, I did not take an antidote, etc.
- Presence of required support factors (other factors without which the cause could not be expected to produce this effect): The emetic was swallowed before too much poison was absorbed from the stomach.
- Presence of necessary intermediate step: I vomited.
- Presence of symptoms of the putative causes acting to produce the effect: There was much poison in the vomit, which is a clear side effect of the emetic's being responsible for my survival.
- Characteristics of the effect: The amount of poison in the vomit was measured and compared with the amount I had consumed. I suffered just the effects of remaining amount of poison; and the timing of the effect and size were just right.

4. What makes these evidence?

A catalogue is important. It drives home the point that there is a good variety of different kinds of evidence that we may be able to establish in particular cases that can speak for individualised singular causal claims. I arrived at these by looking across disciplines and practices, at what I see us actually doing in courtrooms, in physics experiments, in deciding who stole the cookies from the cookie jar. They all, I believe, make common sense. Can we say anything more in their favour? Yes.

Many studies of causality nowadays use a 'causal structural equations framework'. Although functional form should vary to suit the situation, for simplicity the equations are often written as linear forms where the exact form does not matter for the lesson drawn, in which case they look like this:

Sample form for a causal structural equation set:

$x_2\ c{=}\ a_{21}x_1 + a_2$

$x_3\ c{=}\ a_{31}x_1 + a_{32}x_2 + a_3$

$x_4\ c{=}\ a_{41}x_1 + a_{42}x_2 + a_{43}x_3 + a_4$

....

where c= signifies that the two sides of the equation are equal and that all the factors on the right-hand side are causes of those on the left. I write the equations this way because this is the way you will frequently see them, with *a*s and *x*s, which suggests that *a*s are parameters and not ordinary variables representing causes. In general that will not be the case. Few factors we pick out and represent by variables can produce a contribution to an effect by themselves. They need helping factors, or auxiliaries: oxygen in the room if striking a match is to produce a flame. The *a*s represent the auxiliaries, some of which may be factors already represented by an *x*. The difference between *a*s and *x*s is that a causal structural equations system lays out equations representing the causal principles for a set of quantities represented by the *x*s but does not tell you what brings about the quantities represented by *a*s.

To adopt a causal structural equations framework thus is to embrace, even if only implicitly, J.L. Mackie's view that causes are INUS conditions for their effects, or where variables with more than two values (yes the feature represented obtains, no it doesn't) are concerned, that causes are INUS conditions for contributions to the effect, where in the linear form of structural equations we are supposing, each separate term that adds linearly represents a contribution.[3] An INUS condition for some event e is an *I*nsufficient but *N*ecessary part of an *U*nnecessary but *S*ufficient condition for the occurrence of e or, for multi-valued variables, for a contribution to e. (Mackie 1965). Each one of the terms on the right-hand side of a structural equation $x_j\ c{=}\ a_{ji}x_i$ represents a set of conditions which together are sufficient to produce a contribution to the effect represented by x_j, and both components of a term $a_{ji}x_i$ ---a_{ji} and x_i--- represent INUS conditions for the production of this contribution to x_j. Structural equations provide the means to represent in quantitative terms the contribution of each cluster of INUS

[3] Note that holding the view that all causes are INUS conditions does not commit one to the view that all INUS conditions are causes, a view that faces well-known counter-examples (see e.g. Cartrwright, 1989 or Baumgartner 2008, 339). Note also that, because adopting a causal structural equations framework is to embrace the view that causes are INUS conditions, accounts based on such frameworks are subject to any sound objections raised against this component of Mackie's view (see e.g. Baumgartner 2008, 342-346).

conditions represented by a term $a_{ji}x_i$ to the effect x_j. They also allow for the representation of the way the different terms each of which is sufficient for the production of some contribution to x_j interact to determine the value of x_j.

Structural equations are used in discussions of both singular causation and generic causation. In the latter case they represent general causal principles or what are sometimes called 'causal laws'. For me the difference is negligible since I take generic causal claims to be claims about individuals in a particular population. But our topic here is singular causation. In using a structural equations framework in this context, I shall suppose that outcomes in a given individual are fixed by a set of causal principles true of that individual and that these principles can be expressed in a set of equations like those above, where the effect is on the left and the causes for that individual of that effect are on the right. It is assumed that,[4]

Causal Principles:

1. The causal relation is irreflexive (nothing causes itself),
2. the causal relation is asymmetric (if a causes b, b does not cause a),
3. causes occur temporally before their effects,
4. there is a principle that holds for every effect at every time,
5. the principles satisfy a kind of transitivity requirement: for any cause X that appears in a principle, substituting a full set of causes of X for X in that principle yields a true principle,[5] and
6. any true relations that hold among quantities appearing in the causal principles, are guaranteed to hold given the truth of the relations described in the causal principles.[6]

Because of CP1., CP2. and CP3., a set of equations meant to represent causal principles will take the familiar block triangle form of my sample set. It is important to note that 1-6 constrain the notion of a causal principle but will not generally pick out a unique set of equations.

[4] Note that these are necessary conditions, not sufficient.

[5] It should be noted that this does not imply the more contentious claim that singular causal claims are transitive. See below for further discussion.

[6] So, consider, e.g., a cause c with two effects, e_1 and e_2. Supposing determinism, e_1 obtains iff e_2 obtains. That is not among the causal principles. But it obtains on account of the causal principles.

To my mind this is not amiss because causal notions are not reducible to non-causal notions; any attempt to characterise one causal concept will involve mention of other equally causal concepts. For instance, 1-6 can pick out a unique set of principles if 6 were to be replaced by 6': any true relations that hold among quantities appearing in the causal principles, hold on account of (or 'because of' or 'are generated by') the relations described in the causal principles. But the 'on account of' relation in 6' is hardly more transparent than the concept of causal principle I am trying to explicate.

Then to connect equations and principles we make the canonical assumption that the equations are right when they represent the right principles. So,

> *Causal correctness for structural equations*:
> A causal structural equation set is correct for individual iff each equation represents a true causal principle that holds for that individual.

The point of the equations is that the factors represented could take different values---I might or might not take the poison, might or might not take the emetic, might or might not survive. They represent all the different arrangements of these values that are possible for the individual. So they imply counterfactuals. There are two different attitudes to the connection between the equations and the counterfactuals.

Judea Pearl (2000), and I, and various others take the equations as basic: they represent the principles that 'nature' or the underlying physical or psychological or social structure sets. I also take the notion of causation represented in 'c=' as primitive. I explicate this primitive notion of causation by describing constraints on the equations, like the ones I laid out. Also, as I construe causation, principles like these usually depend on some underlying system that gives rise to them and that we can use to explicate them. In nice cases we can even derive the principles from facts about the underlying system, its parts and what they do in interaction, as when we derive Kepler's laws from the structure of the planetary system and Newton's theory. I call these underlying systems, *nomological machines*; philosophers of biology subsequently called them *mechanisms*. All of that helps give meaning to the equations and the general principles governing individuals that these equations represent. The counterfactuals for any individual then fall out from the causal principles that hold for the individual.

Others, like Christopher Hitchcock (see e.g. Hitchcock 2007), who has championed the use of structural equations for singular causation, take nature (or the underlying structure) to set directly a very great number of counterfactuals as true; sometimes they are even nondenumerable. In this case the primitive undefined concept is 'What would happen if...' The equations then are mere summaries of these counterfactuals. It is clear that Hitchcock also takes it to be a fact about nature that the set of counterfactuals set by nature can be summarised in sets of equations that have the requisite form (i.e., the relations between variables is irreflexive, asymmetric, time ordered, etc.).

Which you take to be the case does not matter to my point in this paper. I mention the two different attitudes because I want to underline that it does not matter. Whichever way we look at it, I want to show that we can ground the kinds of evidence I have described in such a set of equations. If you follow Hitchcock, you then will see the evidence as grounded in counterfactuals. On my view, the equations record principles governing the individual which then imply counterfactuals. But on neither account is there any *particular* counterfactual that implies, or is implied by, the singular causal claim. Nor will the kinds of evidence I catalogue speak immediately to any counterfactual. If it is evidence for some counterfactual, it is indirect evidence. It is evidence that speaks immediately to features of the equations, which then in turn imply counterfactuals. We thus avoid having to tackle the problem that troubles the singular-cause sceptics I described at the start, of how one can have direct evidence of what would have happened if things had been different. What constitutes the evidence as evidence of causation is that it tells us about features of the equations, not that it tells us directly about any counterfactuals.

Structural equations generally involve multiple-valued variables, often continuous ones. But a lot of the literature on singular causation, and in particular my very simple-minded example, uses two-valued variables: c occurs or it does not, e occurs or it does not; I take the emetic or I don't, I vomit or I don't; my health is near its starting state or it is not. In this case it is simpler to use the notation of Boolean logic (if we suppose that causation is deterministic). The resulting expressions are familiar from J.L. Mackie's aforementioned discussions of causes as INUS conditions. They look like this:

$$e \; c \equiv V_i a_i c_i.$$

Here c≡ means that the factors represented by variables on the right-hand side are causes of those on the left and that the left- and right-hand sides of the c≡ sign will be equivalent in truth value for all allowed combinations of values of the quantities represented on the right-hand side. The variables represent occurrences at different times, so the same feature (height, weight, exhibiting certain symptoms, etc.) occurring at different times will be represented by a different variable.

Now we can address the question of truth conditions for singular causal claims. I shall suppose that singular causal claims are *equation-based:*

> *Equation-based assumption*:
> Where a claim of the form 'c caused e in individual j' has a determinate truth value, a causal structural framework applies. That requires that there is a set of equations that hold for j that have the characteristics implied by CP1.--CP6.

We can read 'hold for' here in either of the two senses I described. We can suppose, as I do, that a set of equations is correct if it represents correctly a set of causal principles that obtain for j. Or we can see a set of equations, as Hitchcock does, as a convenient summary of a vast network of certain kinds of counterfactuals about j and j's situation.

The equation-based assumption is clear in the work of those who use a causal structural equations framework to treat singular causal claims. Hitchcock, as I have noted, uses these equations to analyse singular causal claims even though he grounds the truth of these claims in a network of counterfactuals and eschews the principles that I take the equations to represent. My suspicion is that detailed investigation would reveal that many other counterfactual accounts, as well as many manipulation accounts, will also be committed to causal structural equations. After all, these kinds of accounts suppose that there are facts of the matter about what would have happened and what would have caused what if things had been different; and it is difficult to imagine that a summary of these facts would fail to have the characteristics implied by CP1.--CP6. So the equation-based assumption will have wide appeal, and with it then, so too should my catalogue of evidence types. I am less sure about alternatives, like process theories of singular causation or G.E.M. Anscombe's view that there are a host of causal relations that we identify with thick causal descriptions like 'pushing' or 'lapping up', where what matters is that one happening produces the other or the second 'comes out

of' the first. In cases where this is all there is to singular causation I am not sure that my evidence types genuinely are evidence.

Given the equation-based assumption, truth-conditions for singular causal claims fall straight out of the structural equations framework:

> 'c caused e in individual j' iff c occurs and e occurs and there is a causal equation that holds of j and there is some set of factors a that actually occur, s.t. under that equation 'c&a' conjointly are sufficient to cause e.[7]

It should be noted that this is a 'top-down' or 'general to specific' explication of singular causation: the central facts that make a singular claim true are facts about causal principles. It differs in this respect from other ways of clarifying causal notions that I have discussed that take singular causation as primitive and use that to explicate concepts of causal laws (see e.g. Cartwright 1989, Chapter 3).

The question that usually arises whenever an account of the truth conditions for singular causal claims is put forth is: Is this account faithful to our intuitions, especially as regards cases of symmetric overdetermination, cases of preemption or the causal status of omissions? Ned Hall and L.A. Paul's *Causation: A User's Guide* (2013), for instance, proceeds entirely under the assumption that fit with our intuitions in a handful of contentious cases---including cases of symmetric overdetermination, preemption and omissions---is the yardstick by which to evaluate accounts of singular causation.[8]

Let me make two remarks here. First, insofar as the truth conditions specified above do not require that the occurrence of e counterfactually depend in any way on that of c, the account I offer is impervious to cases of redundant causation (e.g. overdetermination or preemption) that create trouble for counterfactual approaches because they are cases in which e does not

[7] This will have to be suitably modified when we are not using just dichotomous variables. Essentially it will go over to something like this: 'X = x causes a contribution of size ax to outcome Y iff there is a set of factors a that take net value a s.t. under a true equation the net contribution of X to Y given X = x and A = a is ax.'

[8] Glymour et al. (2010) defend the view that this is the wrong way to proceed if one's concern is to arrive at an account of singular causation that is faithful to our intuitions in general.

counterfactually depend on c and yet, our intuitions seem to tell us, c causes e.[9] And since (i) claims of the form 'o caused e in j'---where o is an omission---are typically assumed to have determinate truth values and (ii) omissions can be INUS conditions, whether the account given above allows for omissions to be causes will depend on whether there is any equation of the form 'e c= ao' representing a true causal principle holding for j in the relevant set of causal structural equations. In other words, issues regarding the causal status of omissions arise upstream from the account of singular causation given above. They arise, so to speak, at the stage at which the true causal principles that hold for j are laid out.

The second remark is a reminder that causation is what I, after Otto Neurath, call a *Ballung* concept (See e.g. Cartwright et al. 1996, Part 3 or Cartwright and Runhardt, 2015): A concept made up of a congestion of ideas, with boundaries that are not clearly and unequivocally delineated. Different aspects of the concept will be stressed and different boundaries will be drawn in different ways depending on the use to which the concept is put and on the context in which it is put to use. My concern here is not to provide an account of singular causation that is faithful to our intuitions in a handful of test cases. Rather I aim to offer a characterisation that serves a different set of related ends: First, it makes more precise our *Ballung* concept of causation in a way that can be useful in a variety of scientific and policy settings; second, we can catalogue a variety of kinds of evidence about whether the singular causal relation as thus characterised obtains; and third, we can make it clear why the facts we take to be evidence for singular causal claims as characterised are in fact evidence for such claims.
I do not claim that this is the only nor the best way of making the concept of singular causation more precise. To the contrary. I would argue that there is no one right way to do it, and that is because our ordinary concept is a *Ballung* concept. It is not precise by its very nature. But sometimes, as in science and often in policy, we need a precise concept. Different ways

[9] A very helpful anonymous referee is concerned about overdetermination---two factors both sufficient for the same effect according to a correct equation---in asymmetric cases where intuitions want to count one as a cause but not the other. But the INUS/equation-based account is deliberately designed to do just that: two determining causes *are* both causes; it is a central feature of the Mackie INUS account that there is more than one way to skin a cat! The referee worries that this undermines the job of causes to serve as effective strategies: we might spend effort to ensure the presence of one of the over determiners when another is already there. I agree that that's true and we would like to know the full structural causal equation to avoid that. But that does not show that the overdeterminer is not an effective strategy, but rather that it is not a cost effective strategy.

of making it more precise will allow it to do different jobs in different contexts. Perhaps we can develop a concept more useful for many of the contexts I have in mind by refining what I propose here or even taking a very different tack. What matters is that the concept we develop be able to do the job we require of it AND that we stick with the sense characterised throughout. It is no good gathering positive evidence using a method that is good for evidencing singular claims as made precise in one way and then drawing inferences that are licensed by some other sense. That is to do science by pun.

With the concept as characterised by the truth conditions above, it is apparent why the five types of evidence in our catalogue are indeed evidence for the claim 'c caused e in individual j'. I can illustrate with our emetic example.

In the emetic example, the principles that fix what happens in the specific situation I was in yesterday when I swallowed the poison will look like this:

1. t(1): v_1 c≡ φ
2. t(2): v_2'c≡ m v ψ
3. t(2): v_2 c≡ χ
4. t(3): v_3 c≡ π
5. t(4):h_4 c≡ (s&v_2') v Ω
6. t(4): h_4 c≡ (s&m) v (s& ψ) v Ω
7. t(4): h_4' c≡ a&-m
8. t(5): h_5'' c≡ ε
9. t(5): p_5 c≡ s&m

Here V represents mild vomiting, V' represent violent vomiting, M represents taking the emetic, H represents the degree of health I actually experienced at t(4), H' and H'' represent different degrees of health (where H ≠ H' ≠ H''), S represents absorbing a small amount of poison, P represents the actual amount of poison measured in vomit, and A represents taking an antidote for the poison. The numerical subscripts attached to each variable indicate the times at which the events represented by these variables are stipulated to occur. Finally, Greek letters represent causes not adumbrated. Principles 1, 3 and 4 for instance relate mild

vomitings at times t(1), t(2) and t(3), respectively, to unspecified causes other than taking the emetic. Notice that 6. follows by substitution of 2. into 5 using the rule Peter Menzies (2012) calls "composition" in his use of structural causal equations.[10] Writing it down does not commit me to assuming transitivity of singular causal facts in general; it just is the case in this case that 6. is true.

Here I have recorded only a small sample of the causal principles I take to be true in this case. I have added some principles that were not called into play to illustrate my point about grounding the 5 types of evidence. So for this case consider,

1) Effect characteristics. Here we can look at,
    a) Timing: The degree of health actually experienced occurred at t(4) as it should were the emetic a cause. h'' occurs at the wrong time. Even if it is equal to h, it is too late for the emetic to be the cause. (E.g. Had the emetic caused that degree of health at t(4), by t(5) natural earing off of the effects would make the actual health at t(5) higher.)
    b) Size of effect: h is the degree of health actually experienced at t(4) and it is what 6. says follows from s&m.
2) Presence of required support factors.
    a) In finding out whether s holds or not we are finding out about the factors in equation 6. without which m cannot produce h.
3) Presence of intermediate stages.
    a) The violent vomiting at t(2) is, as we can see from equations 2. and 5. a necessary condition for m to cause h at t(4).
4) Symptoms of causality
    a) We see from 3. and 6. that the emetic is the cause of the degree of health actually experienced at t(4) if and only if it is the amount of poison in the vomit at t(5) is p, the amount actually measured.

[10] Though Menzies does not explicitly define it, 'composition' is the rule that permits the substitution of the causes $x_1, \ldots, x_n$ of some effect y to y itself in the right-hand side of any structural equation in which y appears.

5) Elimination of alternatives.

   a) No special body mass, etc. These are the factors that should be filled in for ψ and Ω in 6. Showing they did not occur leaves s&m as the only possible cause.

What I have done here can be generalised to any set of structural equations. What I have shown is that,

> *The five types of evidence that one factor produced another in the single case are evidence for that because they inform us about actual values of variables in the causal equations that govern the case.*

Moreover, grounding their claims to be evidence in these equations gives us a lead into how strongly they speak for or against the causal connection. For instance, elimination of alternatives is a clincher. If you can really do it, the connection follows for sure. Similarly with the presence of the requisite auxiliary factors. If they are there then the cause will produce the effect. In special cases like the one imagined here, where a particular side effect occurs just in case the cause in question caused the outcome and it wouldn't occur otherwise, then presence of symptoms of causation will also be a clincher. Presence of intermediate effects is not and for two reasons. They can be present yet not caused by the cause in question even though they could have been; and something can always intervene between the last one considered and the final outcome. When the evidence claims, even if true, cannot clinch the causal connection, what should we do? I have no special advice here. That is our usual position: the evidence underdetermines the conclusion. For better or worse, we can employ our usual solutions.

Noting that the intermediate steps can be present but caused by a different cause suggests a 6^{th} type of evidence we can also adduce:

6) Production of intermediate stages by their predecessors. This kind of evidence will be stronger than the mere presence of the requisite intermediate steps. To warrant the claim that that there is a causal connection between one stage and the next one can appeal to the types evidence already listed.

5. How do we know the equations?

That's a good question. Recall earlier I said that we are hostage to ignorance on three fronts:

- how strong the link, if any, is between the evidence and the conclusion,
- how sure we can be about the strength of this link,
- how warranted we are in taking the evidence claim to be true.

Grounding the types of evidence in structural equations shows us that there is a link and gives an indication of how strong it is. So how sure we can be of this link depends on how sure we are of the equations. And perhaps that's the killer. Maybe we know nothing at all about how this kind of effect could be produced by this cause or any other. Then we won't be able to appeal to the kinds of evidence I have proposed. But we are often in a better position and sometimes in a far better one. We have a theory about the principle, perhaps a general theory, perhaps one very local to this case. Sometimes the theory is already well-supported and can get independent support in this case (and some of this support can come from knowledge about the underlying system that gives rise to the principles we postulate). Sometimes this will be easier because we don't need the theory to be very precise. My emetic example is a case in point.

Finally, getting a lot of positive evidence of these various types will help confirm both the theory and the singular claim at once by the hypothetico-deductive method. My theory says if I swallow the emetic and vomit violently, that will cause me to get better. I do swallow the emetic and I do vomit violently and I do get better. The theory also tells me this should happen if, but only, if I have absorbed just a small amount of poison by the time of vomiting. Those are things the theory tells me to expect and so when they occur, they speak in favour of the theory. How strongly? That is a good question. But this is a standard problem for the hypothetico-deductive method, not one peculiar to cases of singular causation. If our worry is that I got better after this but perhaps not on account of this, the sequencing, the timing, the size of effect, the presence of auxiliary factors, the absence of auxiliaries for other causes hypothesised by the theory all contribute to the case for both the theory and the singular claim. But clearly no amount of positive consequences of our hypotheses will clinch the hypotheses for us. Again, that is just the usual situation that we have learned to live with, sometimes well and sometimes not.

6. Conclusion

G.E.M. Anscombe argued that we can often warrant singular causal claims just by looking at the causal relation in the single case: We can see the cat lapping up the milk. RCTs, currently touted as gold standard for warranting causal claims in evidence-based medicine and social policy, and increasingly across the social sciences, don't look at the causal relation in any single case. Their warrant for claims about causal connections in individual cases is entirely indirect.

I offer a strategy in the middle. I have outlined a concept of single-case causation and catalogued a number of kinds of evidence we can have directly about the single case itself to warrant this kind of causal claim about it. I have also pointed out that our warrant for taking any of these *as* evidence in a particular case will depend on how secure we can be about the related hypotheses about how the case is governed. Our hypotheses teach us how the cause should be expected to produce its effect in this case; our evidence speaks to whether matters have indeed proceeded in the way to be expected. So the strategy I offer is not useful in cases where we know nothing about how the putative cause should produce its effect and are in no position to propose plausible hypotheses with enough detail to tell us enough about what to expect. But then, how often are we in that unfortunate position?

Acknowledgements

I would like to thank Alex Marcellesi for research assistance and thank the participants in my winter 2013 graduate seminar on ‘Evidence and Singular Causes’ at UC San Diego as well as an anonymous referee for help with the ideas and details of this paper.

References

Baumgartner, Michael. (2008). “Regularity Theories Reassessed”, *Philosophia*, 36: 327-354.

Bird, Alexander. (2010). “Eliminative Abduction: Examples from Medicine”, *Studies in History and Philosophy of Science*, 41: 345-352.

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

Cartwright, Nancy, et al. (1996). *Otto Neurath: Philosophy between Science and Politics*. Cambridge: Cambridge University Press.

Cartwright, Nancy and Runhardt, Rosa. (2015). “Measurement”, in N. Cartwright and E. Montuschi (eds.), *New Topics in the Philosophy of Social Science*. New York: Oxford University Press.

Glymour, Clark, et al. (2010). “Actual Causation: A Stone Soup Essay”, *Synthese*, 175: 169-192.

Hill, Bradford. (1965). “The Environment and Disease: Association or Causation?”, *Proceedings of the Royal Society of Medicine*, 58: 295-300.

Hitchcock, Christopher. (2007). “Prevention, Preemption, and the Principle of Sufficient Reason”, *Philosophical Review*, 116: 495-532.

Mackie, J.L. (1965). “Causes and Conditions”, *American Philosophical Quarterly*, 2: 245-264.

Maudlin, Tim. (2002). *The Metaphysics Within Physics*. New York: Oxford University Press.

Moat, Helen Susannah, et al. (2014). “Using big data to predict collective behavior in the real world”, *Behavioral and Brain Sciences*, 37: 92-93.

Pearl, Judea. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.