Peng, C.K., Havlin, S., Stanley, H.E., and A.L. Goldberger (1995): Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series *Proceedings of the NATO Dynamical Disease Conference*, edited by L. Glass, Chaos 5: 82–87.

Samanidou, E., Zschischang, E., Stauffer, D., and T. Lux (2007): Agent-based models of financial markets, *Reports on Progress in Physics* 70: 409–450.

Sornette, D. (2003): *Why Stock Markets Crash: Critical Events in Complex Financial Systems*, Princeton: Princeton University Press.

Sornette, D. (2006): *Critical Phenomena in Natural Sciences – Chaos Fractals, Selforganization and Disorder: Concepts and Tools*, Berlin: Springer.

Sterman, J. (2000): *Business Dynamics: Systems Thinking and Modeling for a Complex World*. New York: Irwin McGraw-Hill.

Voit, J. (2001): *The Statistical Mechanics of Financial Markets*, Berlin: Springer.

Woodward, J. (2003): *Making Things Happen – A Theory of Causal Explanation*. Oxford: Oxford University Press.

Zvelebil, M.J., and J.O. Baum (2007): *Understanding Bioinformatics*. Garland Science.

·42·

Third time's a charm: Causation, science and Wittgensteinian pluralism

Julian Reiss

Abstract

Pluralism about causation seems to be an attractive option as the term seems to defy analysis in terms of necessary and sufficient conditions. This chapter examines a specific form of conceptual pluralism about causation, one that has been termed 'Wittgensteinian'. I will present three such accounts in detail. All three accounts share the rejection of attempting to define 'cause' in terms of necessary and sufficient conditions, and they regard instances of causal relationships to share family resemblance at best. After criticizing and rejecting two earlier accounts, I will develop an alternative that, to the best of my knowledge, does not suffer from the deficiencies of its fellows and is more firmly grounded in some of Wittgenstein's ideas about meaning.

42.1 Introduction

Pluralism about causation is an attractive option. All theories of causation face counterexamples and all attempts to fix them lead to new counterexamples. Though, as always in philosophy, guarantees are hard to come by, there is ample *prima facie* evidence that there is no single essential property or set of essential properties that is shared among all causal relations. In response, a growing number of philosophers have considered pluralistt stances towards causation (Anscombe 1971; Campaner and Galavotti 2007; Cartwright 1999; 2007; De Vreese 2006; Godfrey-Smith 2009; Hall 2004; Hitchcock 2003; Longworth 2006aa, b; Psillos forthcoming; Reiss 2009; Russo and Williamson 2007; Weber 2007).

Pluralism about causation is, however, more of an assortment of ideas than a definite theory.[1] Most fundamentally, one can distinguish pluralism about causation at three different levels:

[1] In a recent survey paper, for instance, Chris Hitchcock distinguishes no less than nine forms of pluralism (Hitchcock 2007).

- *evidential pluralism*: the thesis that there are more than one reliable ways to find out about causal relationships;
- *conceptual pluralism*: the thesis that 'cause' and its cognates has more than one meaning; and
- *metaphysical pluralism*: the thesis that there is no one kind of thing in the world that makes a relationship causal.

This chapter is concerned with a specific form of conceptual pluralism about causation, one Chris Hitchcock terms 'Wittgensteinian' (Hitchcock 2007, pp. 216–7). I will present three such accounts in detail. All three accounts share the rejection of attempting to define 'cause' in terms of necessary and sufficient conditions, and they regard instances of causal relationships to share family resemblance at best. After criticizing and rejecting two already existing accounts, I will develop an alternative that, to the best of my knowledge, does not suffer from the deficiencies of its fellows and is more firmly grounded in some of Wittgenstein's ideas about meaning.

42.2 Wittgensteinian pluralism, takes one and two

Wittgenstein famously claimed that we cannot give a definition of the concept 'game'. He asks us whether all games – board games, card games, ball games, Olympic games – had something in common and observes that although some kinds of games have some characteristics in common there is no one characteristic or set thereof common to all instances games. Hence, we cannot define 'game' in terms of necessary and sufficient conditions (Wittgenstein 1953, §66). Instead, he argues, 'we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail'. Further, 'I can think of no better expression to characterize these similarities than "family resemblance"; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc. etc. overlap and criss-cross in the same way. – And I shall say, "games" form a family' (§66–7).

Although the focus of her paper is an attack on two Humean dogmas – that causes necessitate their effects and that causal relations are not observable, Elizabeth Anscombe presents an account of causation that understands 'cause' as analogous to 'game' (Anscombe 1971 [1992]). She explains (*ibid.* p. 93; emphasis original),

> The word 'cause' itself is highly general. How does someone show that he has the concept *cause*? We may wish to say: only by having such a word in his vocabulary. If so, then the manifest possession of the concept presupposes the mastery of much else in language. I mean: the word 'cause' can be *added* to a language in which are already represented any causal concepts. A small selection: *scrape, push, wet, carry, eat, burn, knock over, keep off, squash, make* (e.g. noises, paper boats), *hurt*. But if we care to imagine languages in which no special causal concepts are represented, then no description of the use of a word in such languages will be able to present it as meaning cause.

If such causatives or 'thick causal verbs' (Cartwright 2004)[2] are understood as constituting the meaning of 'cause', the account faces various problems. To see these, let us define:

Wittgensteinian Pluralism X causes Y if and only if X stands in relation $r \in R$ to Y, where each element of R can be described using a causative in Anscombe's sense.

An immediate problem with this formulation is that causal relations are typically transitive but it is hard to describe the resulting relation using a causative. Consider the following example. *A child upsets a glass of milk. The milk flows on the table, creating a white puddle. Observing the puddle alarms a parent who rushes to fetch a cloth and wipe it off.* It is perfectly meaningful to say that the child (or the child's action) caused the cloth to be milky. But the child didn't *wet* or *stain* or *soak* the cloth. A possible solution would be the following amendment:

*Wittgensteinian Pluralism** X causes Y if and only if X stands in relation $r \in R$ to Y, or such that there is a chain of relations $Xr_1C_1r_2C_2 \ldots C_{n-1}\ r_nY$ with $r_1, r_2, \ldots r_n \in R$, where each element of R can be described using causative in Anscombe's sense.

In this formulation there may remain problems regarding transitivity because it builds transitivity into the concept of cause and not all causal relations are transitive (see for instance McDermott 1995). I will not pursue difficulties relating to the transitivity of causation any further here because they are not specific to the Wittgensteinian account at stake here.

There are, however, two objections that require closer attention. The first is that this proposal limits causation to cases where there is an active agent, mechanism or process that produces the effect, and not all cases in which 'cause' is used meaningfully involve such an agent, mechanism or process. The second objection is that the account fails to provide a criterion to distinguish genuine causatives from non-causal transitive verbs.

The first objection concerns cases of causation by absences. Absences can figure in causal claims both on the side of the cause as well as on the side

[2] Hitchock (2007) regards Cartwright's theory as a form of Wittgensteinian pluralism. This theory is one of physical causation rather than meaning and therefore not necessarily subject to the criticisms raised here.

of the effect. Cases of the former type are omissions. For instance, Billy's failure to water the plants caused their wilting. Cases of the latter type are preventions. For instance, Suzy's catch caused the ball not to hit the window; it prevented the shattering of the window. In neither case can the abstract 'cause' be substituted by a more concrete causative. Whatever Billy did when he failed to water the plants, he did not *desiccate*, *dehydrate* or *dehumidify* them. Billy did not act, he failed to act. Likewise, Suzy (or Suzy's catch), while stopping the *ball*, did nothing to the *window*.

Proponents of process or mechanistic theories of causation bite the bullet and deny that omissions and preventions are genuine cases of causation. Phil Dowe, for one, uses a (counterfactual) concept of pseudo-causation to describe such cases (Dowe 2000). Peter Machamer thinks that these are not cases of causation, but that can be causally explained (Machamer 2004, 35f.):

> Non-existent activities cannot cause anything. But they can, when other mechanisms are in place, be used to explain why a given mechanism did not work as it normally would, and why some other mechanism became active. Failures and absences can be used to explain why another mechanism, if it had been in operation, would have disrupted the mechanism that actually was operating. Maybe we should draw a distinction and say they are causally relevant rather than causally efficacious. They are not, to use an old phrase, true causes.

But such responses cut no ice when the meaning of causal claims is at stake. Neither ordinary language nor the language of science makes a difference to whether the causal relation involves 'presences', i.e. entities that can act and be acted upon or absences of such entities. Below I will discuss in detail an example from the health sciences that involves causation by absences at the generic level. In some cases it may not even be clear whether or not a relatum is present or absent, and causal language can be used to describe the case perfectly meaningfully (Schaffer 2004).

The second objection was that the Anscombe account lacks a criterion to distinguish causatives from non-causal verbs. How do we demarcate verbs that belong in the category used to describe the relation *R* from those which don't? Certainly not all *verbs* belong in this category. Even though many causal processes are involved in someone walking, we don't describe a causal relation by saying 'Billy is walking'. Nor are all *transitive verbs* causal: 'Billy measures five foot nine' does not describe a causal relation. There are many relations that are non-causal and that can be described using transitive verbs: '*A* entails *B*', '5 and 7 sum up to 12', 'H_2O consists of two hydrogen and one oxygen molecules'; 'The fall in the barometer reading predicts the storm'.

It seems to be the case that once we discover that a certain transitive verb applies to some situation, it is an *additional* discovery that this verb belongs to the set of causal verbs. Moreover, there are numerous verbs that can have causal and non-causal meanings: *determine*, *induce*, *fix*, *lead to*, *depend on*. And

perhaps this phenomenon is more wide-spread than seems at first sight. Many verbs have numerous meanings, only some of which are causal in the way required for Anscombe's account to work. '*To scrape*' means '(1a) *to remove from a surface by usually repeated strokes of an edged instrument*' (causal) or '(1b) *to make (a surface) smooth or clean with strokes of an edged instrument or an abrasive*' (causal) but also '(2a) *to grate harshly over or against*' (non-causal); '*to carry*' means '(1) *to move while supporting*' (causal) but also '(14b) *to provide sustenance for (land carrying 10 head of cattle)*' (non-causal); '*to eat*' means '(3a) *to consume gradually*' (causal) but also '(1) *to take in through the mouth as food*' (non-causal).[3] Thus, for every verb we have to discover that it can be used causally and for some we have to discover in addition that it is used causally on a given occasion.

A potential way out is to say that certain cases of causal verbs are paradigm cases, and whether or not a new verb is causal is determined by its family resemblance with paradigm cases. This, however, is an unpromising route. Take, for the sake of the argument, Anscombe's verbs: *scrape*, *push*, *wet*, *carry*, *eat*, *burn*, *knock over*, *keep off*, *squash*, *make* (*e.g.* noises, paper boats) and *hurt* as paradigms, and *yield* as a yet-to-be-determined case. How *could* we say that 'yield' bears a family resemblance to, say, 'scrape'? Any two things are similar and dissimilar in many, perhaps indefinitely many ways. There simply is no sense in which two things are similar to each other *simpliciter*. Rather, things are similar *with respect to* some feature or another. 'Yield', then, is supposed to be similar to 'scrape' with respect to its causal content, but how do we determine that without having an independent grasp on the concept of cause?

An alternative to Anscombe's theory, also Wittgensteinian in spirit, is to regard causation as a cluster concept. For the concepts of ordinary language, we apply one or the other of the standard tests for causality. To take an example, consider the claim 'Jim used a blanket to *smother* the fire'. First of all, presumably on this occasion we mean by this something like 'Jim used a blanket to suppress the fire by excluding oxygen' (*cf.* definition (2c) from Merriam–Webster). Did Jim's action cause the fire to end? Yes: Had Jim not thrown the blanket over the fire, it would have persisted; Jim's action increased the probability of the fire's death; covering a fire with a blanket is an effective strategy to end it; there is a regularity between covering fires with blankets and their end; there is a mechanism by which the blanket kills the fire; and so forth. Unless the case answers positively to some or all of these tests (I will discuss the details of how many tests have to be satisfied in the next section), we do not have a case of causation. Hence, satisfying the tests is basic for causation, not the application of a verb that's presumed to be causal.[4]

[3] If that is not convincing, '*ingest*' and '*absorb*' can very clearly be used causally and non-causally. All definitions are taken from the Merriam-Webster online dictinary www.merriam-webster.com. Accessed on 27.10.2009.

[4] Stathis Psillos makes a very similar point about the Machamer–Darden–Craver (MDC) notion of 'activity', focusing on the counterfactual test (Psillos 2004, p. 314; emphasis original): 'Activities,

Are we committing a fallacy here, mistaking test for identity or truth conditions? I don't think so. If '*X* causes *Y*' is true if and only if '*XRY*' is true, where *R* is a relation (or an activity or capacity) described by a thick causal verb, then we need some principled way of telling which verbs do describe relationships that are causal. And this cannot be done, or so I've been trying to argue, unless we have an independent concept of cause. The tests I've mentioned are meant to help us in determining which transitive verbs are causal, not to define causation.

Francis Longworth has developed this proposal in detail. He regards causation as a cluster concept, by which he means the following (Longworth 2006a, p. 112f.):

Cluster concept. There are a number of features that are relevant to, or 'count towards' an individual's being an instance of the concept. X is a cluster concept if and only if the following conditions are jointly satisfied:

(1) The presence of the entire set of features (the 'cluster set') is sufficient for the concept to be applied.

(2) No feature is necessary.

(3) At least one feature from the cluster set must be instantiated.

Longworth suggests that (perhaps, among others) the following features are members of the cluster set (Longworth forthcoming; this is a paraphrase):

- *Counterfactual dependence* ('*E* counterfactually depends on *C*');
- *Lawlike regularity* ('There is a law such that "whenever *C*, then *E*"');
- *Manipulability* ('Changing *C* is an effective strategy to change *E*');
- *Probability raising* ('$P(E|C\&K) > (C|K)$, where *K* is a set of background factors');
- *Mechanism* ('There is a local physical process from *C* to *E*');
- *Responsibility* ('*C* is [morally] responsible for *E*').

Counterexamples to univocal theories of causation show that none of these features is necessary for causation. For example, cases of redundant causation

such as bonding, repelling, breaking, dissolving etc., are supposed to embody causal connections. But, one may argue that causal connections are distinguished, at least in part, from non-causal ones by means of counterfactuals. If "*x* broke *y*" is meant to capture the claim that "*x caused y* to break," then "*x* broke *y*" must issue in a counterfactual of the form "if *x* hadn't struck *y*, then *y* would have broken." So talk about activities is, in a sense, disguised talk about counterfactuals'. Notice that Psillos doesn't say '*x* broke *y*' *means* '*x* caused *y* to break', leaving open the possibility of extra content.

Though the authors seem to disagree, I believe that the MDC notion of 'activity' is very close to Cartwright's notion of thick causal verbs in that thick causal verbs describe activities. Hitchcock makes a similar observation (2007, p. 300), pointing out that a difference lies in the fact that MDC use activities as building blocks for their more fundamental notion of a mechanism.

demonstrate the non-necessity of *counterfactual dependence*, in indeterministic cases that of *lawlike regularity* and so forth. However, some subsets of the cluster set are sufficient, e.g. counterfactual dependence and responsibility; production and responsibility; and dependence holding fixed some *G* and responsibility.

Longworth argues that his cluster theory is superior to other accounts in that it explains the truth of five theses regarding the concept of causation (2006a, p. 100; the discussion of how the cluster theory meets these desiderata occurs on pp. 119ff.):

1. *Counterexamples*: There are many extant univocal theories of causation and all of them have counterexamples.
2. *Disagreement*: There are some cases about which individuals disagree in their intuitive causal judgements.
3. *Vagueness*: There are borderline cases of causation.
4. *Error*: Individuals' intuitions are sometimes clearly mistaken.
5. *Degrees of Typicality*: Some cases of causation appear to be 'better' or more typical examples of the concept than others.

Univocal theories must fail because they inflate a single feature of causation into a necessary and sufficient condition; hence, there are counterexamples. Disagreements and vagueness obtain because it is not always clear what precise subset of criteria is sufficient for the application of the concept. Individuals' intuitions are sometimes mistaken because they take the fact that the envisaged scenario has one feature from the cluster set as sufficient to apply the concept while closely analogous cases (which have that and only that feature) are judged differently. Degrees of typicality, naturally, stem from the fact that scenarios have smaller and larger numbers of features from the cluster set.

42.2.1 Understood as account of our ordinary concept of causation

Longworth's account is successful. I know of no case of causation that has *none* of the mentioned features. Whether or not a case that has some but not other features is judged as causation depends on the subsets of the cluster set we take to be sufficient. Longworth does not give a final answer to that question but this flexibility is an advantage of the account. Language is in flux and the subsets of features that are taken to be sufficient for causation and how important the satisfaction of each criterion is each may change over time.

According to this theory, then, 'cause' is ambiguous, vague, gives rise to disagreements in individuals' judgements as well as occasional error, and it comes in degrees. But what seems advantageous from the point of view of our ordinary concept of causation may turn out to be unfavourable for science and policy. For science and policy we require concepts that have a definite

meaning and clear conditions of application. Disagreements, so they arise, should be resolvable with reference to an external standard, not individuals' intuitions.

Perhaps it is not a problem for our ordinary concept of causation that some people believe that the father's inattention was a cause of the child's drowning while others think that it was only a quasi-cause because the there was no physical process of the appropriate kind; or that, for some, the fact that a murderer's parents met at a ball in Vienna is a cause of her criminal deed while for others this thought appears ridiculous. For science and policy having clear answers to such questions matters greatly. In determining whether the father should be held liable for his child's accident, we don't only have to know whether certain normative considerations apply but also whether he was *causally* responsible for the accident. And it won't do to answer the question whether he was causally responsible with 'according to some intuitions yes, according to others, no'. Nor will it do to answer 'in some sense, yes; in another, no'.

The account that I develop in the three sections that follow might answer the question 'does X cause Y?' with 'in some sense, yes; in another, no', depending on the case. But unlike other forms of conceptual pluralism, this one has a methodology built into it how disagreements can be resolved. One could say that it makes cause *unspecific* rather than *ambiguous*. 'Cause' here is an unspecific term that is specified by what I will call an 'inferential analysis': an analysis of what set of propositions the claim in which 'cause' occurs is inferentially connected with. So let us now look at what causation has to do with inference.

42.3 Causation and inference

To develop my own Wittgensteinian account of causation I need to digress for a moment. My account builds on the idea that causation and inference are intimately related. This is most easily seen in Hume's theory of causation because within that theory causation and inference are the two sides of the same medal.

In Hume's theory, for any two independent, spatially contiguous and temporally ordered events A and B, if one knows that A causes B, one is entitled to infer B upon observing A. And if one is entitled to infer B upon observing A, one knows that A causes B. The problem is only that one cannot know that A causes B because one cannot see it. Concomitantly, one is never entitled to infer B upon observing A because the future might not resemble the past. The problems of causation and induction thus collapse into one.

But they do so only because Hume held a regularity view of causation, and that view is well known to be false. Without the regularity view, the relation

between causation and inference is less tight. Few of us hold that an effect *must* follow its cause – an effect might fail to follow its cause for instance because an intervening factor prevents it from doing so or because the cause is indeterministic. Therefore, an observer of the cause is not entitled to infer the effect (but rather something weaker such as 'the probability of the effect is high' or '*ceteris paribus*, the effect will obtain'). Likewise, few of us hold that if an agent is indeed in the position to infer a later event from an earlier that the earlier event *must* be the cause of the later – for instance because the relation may be due to a common cause such that earlier and later event are epiphenomena. Knowing that A is regularly followed by B then does not entitle a language user to infer that A causes B (but rather something weaker such as the disjunctive proposition "'A causes B" or "A and B share a common cause" or "there is some non-causal reason for the association between A and B"'). More tenuously than in Hume, causation and inference are nevertheless related.

An inferentialist theory of the meaning of causal claims explains simply and elegantly why this should be so. Inferentialist theories of meaning hold, roughly, that the meaning of an expression is given by its inferential connections to other expressions. According to some interpreters, Wittgenstein held such a theory in the period between the *Tractatus* and developing the theory of meaning as use in the *Philosophical Investigations*. For instance, in his *Remarks on the Foundation of Mathematics* he says (quoted from Peregrin 2006, p. 2):

> The rules of logical inference cannot be either wrong or right. They determine the meaning of the signs... We can conceive the rules of inference – I want to say – as giving the signs their meaning, because they are rules for the use of these signs.

Building on this idea I propose the following for causal claims. The meaning of a causal claim is constituted by the system of propositions with which it is inferentially connected; that is, the system comprised of those propositions that entitle a language user to infer the causal claim as well as those she is entitled to infer from it.

Let us call such a system an 'inferential system for causal claim CC' or short 'inferential system-CC'. An inferential system-CC can roughly be divided into inferential base, inferential target and the causal claim CC itself. The inferential base (for CC) comprises all those propositions *from which* a language user is entitled to infer CC. The inferential target (of CC) comprises all those propositions that a language user is entitled to infer *from* CC.

Scientists seldom establish causal claims for their own sake but rather because they take them to be conducive to the more ultimate goals of science such as scientific explanation, policy and prediction (to give some examples). If a causal claim together with the relevant background knowledge entitles a user to infer a scientific explanation, a policy claim or a prediction, then these latter propositions constitute what I call the inferential target of the causal

claim. In concrete terms, consider a claim such as 'aflatoxin is hepatocarcinogenic' ('exposure to aflatoxin causes liver cancer'). An epidemiologist might be interested in explaining the population-level correlation between aflatoxin exposure and liver cancer and thus whether it is due to the carcinogenicity of the substance; a policy maker in inferring 'controlling aflatoxin is an effective strategy to reduce mortality'; finally, a person exposed to aflatoxin in knowing whether consumption of aflatoxin will lead to (an increased chance of) liver cancer *in him* and thus in prediction. Below, I will illustrate the kinds of propositions that must be part of the inferential base in order for a language user to be entitled to these inferences in the context of this case.

Here I will say no more about inferential systems-CC in general save two brief remarks. First, the inferences that form the connections between the propositions contained in it are material rather than formal inferences. Formal models of inference (such as *modus ponens*), as the name suggests, are valid in virtue of their form and independently of the propositions that they take as arguments. Material inferences, by contrast, are valid due to the content of the propositions. To illustrate, consider John Norton's example of contrasting the two inferences 'Some samples of the element bismuth melt at 271°C, therefore all sample of the element bismuth melt at 271°C' and 'Some samples of wax melt at 91°C, therefore all samples of wax melt at 91°C' (Norton 2003, p. 649). It is subject and domain specific (or as Norton calls it, 'material') background knowledge that entitles a language user to the former but not the latter inference. In this case, that background knowledge includes the empirical generalization that chemical elements tend to share physical properties and the fact that bismuth is an element whereas wax is a generic name for a variety of substances. Importantly, proponents of theories of material inference hold that it is not the case that there must be implicit premisses that turn the material argument into a formally valid one once made explicit. Rather, the inferences are licensed by the material facts concerning the subject matter of the propositions involved (Norton 2003; Brigandt forthcoming).

Second, I use the rather clumsy formulation 'inferences a language user is entitled to' in an attempt to strike a balance between a descriptive and prescriptive perspective on meaning. It is clearly the case that ordinary folk as much as sophisticated scientists sometimes make mistakes when inferring a causal claim from evidence or some other claim in the inferential target from a causal claim. It would therefore be incorrect to take those inferences language users actually make as the basis for meaning. On the other hand, there aren't many hard-and-fast rules that philosophers can use to prescribe scientists and ordinary folk what inferences they should and shouldn't make. The best guide to what's doable and what isn't is scientific practice and therefore I won't make highly general claims about what a language user is entitled to. Instead, in the next section I will show how tightly inferential base and target are connected on the basis of a brief analysis of two brief case studies.

42.4 An inferentialist analysis of two causal claims

In this section I consider the kinds of material inferences a user is entitled to make when she knows, first, that 'aflatoxin causes liver cancer' and second, that 'lack of sunlight causes multiple sclerosis'. In particular I will ask under what conditions knowing the causal claim entitles the user to infer (a) a more specific causal claim; (b) a claim about explanation; (c) a claim about policy; (d) a claim about prediction; and (e) a mechanistic claim.

42.4.1 Is aflatoxin carcinogenic *in humans*?

The carcinogenicity of aflatoxin is more like Norton's wax example than his bismuth example in that there is a great deal of variability of the toxicity of substances among different species and populations in general. Aflatoxin turns out to be carcinogenic in human populations but the inference could only be made on the basis of population-specific evidence.[5] Thus, in general, when the causal claim concerns the toxicity of a substance, language users are entitled to inferences about a given population only when the inferential base contains evidence claims about just that population.

42.4.2 Does the carcinogenicity of aflatoxin explain the (human) population-level correlation between the substance and incidence of liver cancer?

It turns out that the inferential base for the human population specific causal claim contains mostly evidence regarding the mechanism of its operation. That is, it contains a claim such as 'There exists a pathway through which aflatoxin produces cancerous growths in liver cells'. For at least two reasons this claim does not entitle to infer the explanatory claim. First, the existence of one or several mechanisms through which aflatoxin causes and therefore increases the chance of liver cancer is compatible with the existence of further mechanisms through which aflatoxin prevents the disease. In this particular case, it is implausible that there should exist a pathway such that exposure to aflatoxin is actually beneficial (e.g. Steel 2008, p. 116). But this is an additional claim the inferential base must contain, which in no way follows from the claim about the carcinogenicity of aflatoxin.

Second, the population-level association is likely to be confounded. In the given case it is infection with the hepatitis-B virus (HBV) that may be responsible for the association. Populations subject to high exposure to aflatoxin are

[5] Steel (2008) argues that the example is a case of successful extrapolation from a claim about animal models (in particular Fischer rats) to humans. I am doubtful whether he is right (Reiss forthcoming). But even if we go along with Steel, the reasoning he presents depends in large part on evidence regarding the *human* metabolism. The important point is that causal claims about toxicity are almost always population specific.

also populations where HBV prevalence is high, and HBV is a known cause of liver cancer. Moreover, HBV is known to *interact* with aflatoxin but in ways that are not fully appreciated (Wild and Ruggero 2009). That is, the carcinogenicity of aflatoxin itself depends on whether or not the compound is co-present with other causes of liver cancer, and it may be the case that even though aflatoxin causes liver cancer in some humans, in populations also affected by HBV aflatoxin is causally irrelevant for cancer (or is even a preventative) so that the association is entirely due to the carcinogenicity of HBV. It is thus no surprise that in one and the same article we can read the following statements: 'Aflatoxins, which are the metabolites of some *Aspergillus* species, are among the most potent hepatocarcinogens known'; 'Several ecological studies have shown a correlation between liver cancer incidence and aflatoxin consumption at the population level, but findings are not entirely consistent'; and 'Case-control studies with dietary questionnaires or biomarkers of recent exposure to aflatoxin have also provided inconsistent results' (Henry *et al.* 1999, p. 2453).

Thus, it may or may not be that the association between exposure to aflatoxin and liver cancer incidence can be explained by the causal claim. Hence the inference cannot be made on the basis of the causal claim alone. In addition, knowledge about other pathways through which the compound affects liver cancer as well as about confounders and modes of interaction is required.

42.4.3 Is control of aflatoxin an effective strategy to reduce mortality of the affected populations?

The usual approach to controlling aflatoxin exposure is to set standards for a maximum level of contamination of finished food products. According to the best available estimates lowering the standard does indeed achieve a small reduction of liver cancer incidence (*ibid.*). However, for two reasons setting stricter contamination standards is not considered a good strategy to reduce mortality. First, higher food standards will lead countries to limit the import of affected products, which may mean that the least contaminated foods and feeds are exported, leaving the more highly contaminated products in the most affected countries. Second, it may lead to food shortages in those countries (*ibid.*). Thus, controlling aflatoxin is not an effective strategy to reduce mortality in the affected populations because the intervention, while decreasing mortality along one path – through aflatoxin consumption and liver cancer – increases mortality along another, viz. food deprivation.

Such an intervention would certainly be 'ham-fisted', to use Elliott Sober's term (Sober 2009). A ham-fisted intervention is one that affects the target variable through pathways that do not go through the cause variable of interest. But there is no guarantee that there exist interventions that are not ham-fisted. Nor is there a guarantee that an intervention that affects, if at all, the effect (mortality) only through the cause (exposure to aflatoxin) leaves the causal relation intact. Especially in the social sciences interventions might be structure altering and therefore unable to be exploited for policy purposes. Again, therefore, a claim about policy can only be inferred when a number of additional pieces of knowledge is contained in the inferential base.

42.4.4 Does exposure to aflatoxin predict liver cancer in the individual case?

Just as there is much variability between species, there is often much variability within a single species. Therefore, whether the causal claim is relevant for an individual depends on whether or not the individual belongs to the precise population for which the causal claim has been established. In the aflatoxin case, the toxicity of the substance depends on details of the metabolism that are widely shared among humans, hence establishing carcinogenicity for some humans is likely to be relevant for all humans (and this, once more, is an additional proposition that has to be part of the inferential base if a prediction is to be made). However, even if that is the case, three possible circumstances may drive a wedge in between the truth of the causal claim and successfully using the claim for prediction. First, even if aflatoxin is toxic in most humans, some may have a rare genetic make-up that makes them immune to aflatoxin (that this is not an idle possibility is demonstrated by the fact that some species such as mice are immune). Second, even if a given individual is susceptible to aflatoxin, intervening factors may prevent the causal relation from realising. People might swallow antidotes or die before aflatoxin has made its way through the metabolism. Third, even if the individual is susceptible and nothing intervenes, the cause may fail to produce its effect because the mechanism operates indeterministically. None of these possibilities can be excluded without additional evidence.

Let us now examine a case in which a causal claim has been established by means of epidemiological – that is, probabilistic – data. It has long been known that there is a characteristic pattern in the global distribution of multiple sclerosis (MS): high latitude is associated with a high risk for MS (Kurtzke 1977). But it is difficult to disentangle genetic factors and various environmental factors such as nutrition and culture. Strong evidence that sunlight exposure is the relevant factor came from a quasi natural experiment in Australia. Australia presents a very favourable case for causal analysis because it displays enormous latitudinal spread and climatic variation at the same time as genetic and cultural homogeneity (van der Mai *et al.* 2001, p. 169; references suppressed):

> In Australia, a more than sixfold increase in age-standardized MS prevalence has been demonstrated from tropical Queensland to Tasmania. Within Europe and the United States, there is also an at least two- to threefold gradient of increasing MS prevalence

with increasing latitude. These geographical differences were initially interpreted to represent environmental influences which varied by latitude, such as climatic factors, dietary characteristics and infectious agents. More recent analyses indicate that geographical MS variation, at least in North America, may result from a complex interplay of genes and environment. The marked Australian latitudinal gradient found in the national prevalence survey of 1981 is unlikely to be explained by genetic factors only, because the gradient is evident even among UK and Irish immigrants to Australia, a population subgroup that is predominantly Caucasian. These findings together with the large latitudinal spread across the continent, stretching from 10° to 44° South in latitude, and a uniform health care system provide a good opportunity to examine the relationship between latitude-related factors and MS.

[...]

The aim of this study was to conduct an ecological analysis of the extent to which UVR [ultraviolet radiation] levels might explain the regional variation of MS in Australia. We contrasted the relationship between UVR and MS prevalence with that of UVR and melanoma incidence, because the latter association has previously been demonstrated to be causal.

42.4.5 Is there a mechanism from (lack of) sunlight to multiple sclerosis?

Let us suppose then that it is true that lack of sunlight causes MS.[6] The first thing to note is what has been established is a probabilistic causal claim. That is, in a certain population (caucasians, say), lack of sunlight increases the probability of MS, holding fixed other causes of MS. Many of the limitations described above hold here too. For instance, the claim is population relative and without population-specific evidence no inferences can be made about a hitherto unexamined population. Above I also argued that a mechanistic causal claim does not license an inference regarding the corresponding population-level probabilistic claim. Here let me ask the reverse question: does a population-level probabilistic causal claim entail anything about mechanisms? My answer is once more no but the reasoning requires some elaboration.

When some time passes between the occurrence of a cause and the onset of an effect, it is plausible to assume that there exist some intermediaries that transport the causal message from cause to effect. In the type of biomedical cases I have been talking about, there lie long stretches of time between cause and effect, often many years. There is some evidence, for instance, that

[6] If it is indeed the case, as I believe it is, that this causal hypothesis is widely accepted in the biomedical community, the vitamin-D/MS link provides an interesting case study against the so-called Russo–Williamson thesis according to which both mechanistic as well as probabilistic evidence is required to establish a causal claim (Russo and Williamson 2007). Whereas parts of the vitamin-D metabolism are understood fairly well, the etiology of MS is still completely unknown (e.g. Ramagopalan and Giovannoni 2009).

sunlight exposure during age 6–15 is an important risk factor associated with MS (van der Mei 2003). The onset of the disease typically occurs much later, between the ages 20 and 40 (van Amerongen *et al.* 2004).

Sunlight is required for the skin to metabolise vitamin-D3. UV-B radiation photolyses provitamin D3 to previtamin D3, which, in turn, is converted by a thermal process to vitamin-D3. Vitamin-D3 is biologically inactive but when converted into 1, 25 – $(OH)_2D$, the hormonally active form of vitamin-D, involved in an abundance of biological functions including calcium homeostasis, cell differentiation and maturation and, most relevantly, immune responses. How precisely 1, 25 – $(OH)_2D$ affects MS is unknown but studies with mice have shown that the hormone successfully prevents the onset of experimental autoimmune encephalomyelitis (EAE), which is recognized as a useful animal model for MS (van Etten *et al.* 2003). Moreover, there is some evidence that vitamin D interacts with the major genetic locus which determines susceptibility to MS (Ramagopalan *et al.* 2009).

None of this shows, however, that there is a mechanism from sunlight exposure to onset of MS. It is the *lack* of sunlight that causes vitamin-D *deficiency*. As vitamin D is an important *preventer* of MS, it is the *absence* of vitamin D that causes MS. Now, one might call this a (sketch for a) mechanism. But it is important to see the differences between the causal relations involved in this example and those involved in other cases such as the aflatoxin case that was described above. Exposure to aflatoxin causes cancer through a series of intermediate stages, all of which contain markers that have a clear (and, in fact, unique) association with the toxin. At least in principle, therefore, the causal effect of aflatoxin on liver cells could be learned by both forward as well as backward chaining. Forward chaining uses the early stages of a mechanism to make inferences about the types of entities and activities that are likely to be found downstream and backward chaining reasons conversely from the entities and activities in later stages about entities and activities appearing earlier (Darden 2002, p. 362). Forward chaining thus would start with the consumption of aflatoxin, examine the various stages of its metabolism and eventually establish an effect of an aflatoxin metabolite on liver cells. Backward chaining proceeds by examining these cells, asking what could possibly have caused the characteristic mutation and then backtracking further. As the mechanism is fully present in each individual in which aflatoxin has caused liver cancer, it could (again, in principle) be discovered on the basis of a single individual.

The role of sunlight is not analogous to a chemical compound making its way through the human metabolism. Sunlight is a factor that enables the skin to synthesise vitamin D, which, after several transformations, plays an active role in regulating immune responses among other things. There would be no use in attempting forward or backward chaining in an individual suffering from MS. Even if that individual were deficient in vitamin D, there would

be no sense in which 'lack of sunshine' could be regarded as 'the' cause of the deficiency, analogously to the sense in which exposure to aflatoxin is 'the' cause of the presence of its various metabolites in the blood stream. We might say that lack of sunlight was among the causes of the vitamin-D deficiency because of the truth of the counterfactual 'had the individual been more exposed to sunlight, her vitamin-D levels would have been higher'. But alternative antecedents (e.g. 'had the individual eaten more oily fish' or 'had the individual taken dietary supplements') also make the counterfactual true and with it the associated causal claims. Such counterfactual claims we judge in turn on the basis of population-level epidemiological – i.e. probabilistic – data.

Aflatoxin is an entity that damages liver cells by way of various activities the compound and its metabolites engage in. Nothing analogous is true in the sunlight/MS case. Using the well-known Machamer–Darden–Craver definition of a mechanism according to which 'Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions' (Machamer *et al.* 2000, p. 3), it is straightforward to conclude that there is a mechanism in the former but not in the latter case.[7]

Another way of describing the difference is the following. If it is true that at the population level aflatoxin causes liver cancer, then there must be some individuals whose liver cancer was brought about by aflatoxin. But it is not the case that if at the population level lack of sunlight causes MS, there must be some individuals whose MS was brought about by lack of sunlight. When a mechanism is present, a causal generalization entails something about singular causal relations. When no mechanism is present, there is no such entailment either.

42.5 Re-enter Wittgenstein

Even the more patient among the readers might have wondered by now what these musings about inference have to do with Wittgenstein, pluralism and Wittgensteinian pluralism. Let us look at Wittgenstein first.

Wittgenstein is famous for having remarked that 'the sense of a proposition is the method of its verification' in a conversation with the Vienna Circle (McGuinness 1985, p. 352). But apparently he himself expressed outrage when the 'verification principle' was attributed to him (Anscombe 1995, p. 405) and at least according to some interpretations (e.g. Medina 2001; Peregrin 2006) held an *inferentialist* theory of meaning in the period between the *Tractatus* and developing the theory of meaning as use in the *Philosophical Investigations*. For instance, in his *Remarks on the Foundation of Mathematics* we can read (quoted from Peregrin 2006, p. 2):

> The rules of logical inference cannot be either wrong or right. They determine the meaning of the signs... We can conceive the rules of inference – I want to say – as giving the signs their meaning, because they are rules for the use of these signs.

According to this theory, then, the meaning of an expression is given by the role it plays in our inferential practises. On this view, then, there is a perfectly natural and simple explanation why causation and inference are so intimately related: the meaning of a causal claim is given by its inferential role.

How do we know with what other expressions a given expression is inferentially connected? This is where in Wittgenstein's theory of verification comes in. José Medina explains its role as follows (Medina 2001, p. 308; emphasis is Medina's):

> That the verificationism of the Satzsystem view is at the service of an inferentialist semantics becomes explicit when Wittgenstein remarks that the import of asking of a proposition 'What is its verification?' is that 'an answer gives the meaning by showing the relation of the proposition to other propositions. That is, it shows *what it follows from and what follows from it.* It gives the grammar of the proposition.' [Wittgenstein 1979: 19–20] So, for Wittgenstein, verificationism seems to be a heuristic tool that enables us to analyze the content of propositions in terms of their *inferential use.*

Thus, whereas the meaning of an expression is given by its inferential connections with other expressions in a system of propositions, its method of verification determines what these inferential connections are. This latter point is precisely what I've argued in the preceding section: the method of verifying a causal claim – of evidentially supporting it – determines with what other claims it is inferentially related.

Moreover, it is easy to see how this theory of meaning leads to a form of pluralism about causation. If its inferential connections to other propositions constitute the meaning of a causal claim and the kinds of propositions from which a causal claim can be inferred and those that can be inferred from a causal claim differ from claim to claim, the case for pluralism has been made. Very roughly, we can define identity conditions for causal claims as follows. Suppose the term 'cause' is used on two different occasions and it is not known whether it has the same meaning on both occasions. Two such claims would have the form 'X α-causes Y' and 'Z β-causes W'. We can then say that 'α-causes' has the same meaning as 'β-causes' (on these occasions) to the extent that 'X α-causes Y' is inferentially connected to the same kinds

[7] This is not to deny that there is something *similar* to a mechanism at the type level. It is certainly true that the variable 'exposure to sunlight' is causally relevant to the variable 'vitamin-D level', which in turn is relevant to the variable '1,25-(OH)2D', which, finally, is relevant to the risk of MS. One way to put my point is to say that that if we want to call that a mechanism we can infer at best a mechanism of this type but not a mechanism of the type that mediates the influence of aflatoxin on liver cancer.

of propositions regarding the relation between X and Y as 'Z β-causes W' is inferentially connected to propositions regarding the relation between Z and W. If, to give a fictional example, both 'X α-causes Y' and 'Z β-causes W' have been established by RCTs and both license claims about effective strategies (such as 'promoting X is an effective means to raise the chance of Y' and likewise for Z and Y), then 'α-causes' means the same as 'β-causes' (on these occasions).

There is no guarantee that the kinds of propositions found in inferential base and target are the same for different instances of 'cause'.[8] Different methods of supporting a causal claim license different kinds of inference: this is just what the previous section aimed to establish. Therefore, the meaning of 'cause' in 'Aflatoxin causes liver cancer' and 'Lack of sunlight causes MS' differs – as these claims differ both with respect to the kinds of propositions in their inferential base as well as those in their inferential target.

42.6 Conclusions

The advantages of the account proposed here over its two Wittgensteinian competitors are easy to see. Unlike Anscombe's account inferentialism has no difficulty with cases of causation by absence, as was shown in the discussion of the causal claim about lack of sunlight and MS. The issue whether or not a given transitive verb is a genuine causative simply doesn't arise.[9] Unlike Longworth's account, inferentialism doesn't make causal claims ambiguous or vague or both. There is a definite set of propositions with which any causal claim is inferentially related. True, we might not always have a very clear idea of what these sets are. But this is a question of epistemology, not of semantics.

Finally, inferentialism has an answer to Jon Williamson's challenge: 'If one can't say much about the number and kinds of notions of cause then one can't say much about causality at all' (Williamson 2006, p. 72). It is certainly the case that the type of pluralism entailed by an inferentialist theory of meaning is of the indeterminate variety in that number and kinds of notion of cause are not

[8] Though if the Russo–Williamson thesis were true, researchers in the health sciences did indeed always require both difference-making evidence and evidence about mechanistic connections in order to establish causal claims, and in addition the kinds of propositions one is entitled to infer from causal claims were also the same, then conceptual monism about causation in the health sciences, which they favour, would be supported. I do not think that that thesis is true, and I think that my second case can serve as a counterexample (footnote 5) but it is interesting to note that the thesis (plus one further assumption) entails conceptual monism under an inferential conception of meaning.

[9] An issue that does arise is the parallel one of *justifying* the inferences among base, causal claim and target. But this is one we ought to leave to science. As I claimed above, the best guide to what works and what doesn't is scientific practise, and there is no reason why this area should be exempt from the general principle.

fixed once and for all times. But, as the inferentialist analyses of section four have shown, there is a great deal one can say about causality.

Acknowledgements

Though developed and motivated independently, the view on causation I present here resembles some work on scientific representation by my former colleagues in Madrid, Mauricio Suárez (Complutense University) and Jesús Zamora Bonilla (UNED). I received financial support from two projects of the Spanish ministry of education, FFI2008-01580 and CONSOLIDER INGENIO CSD2009-0056.

References

Anscombe, Elizabeth (1971) [1992]. *Causality and Determination*, Cambridge: Cambridge University Press.

Anscombe, Elizabeth (1995). Ludwig Wittgenstein (Cambridge Philosophers II), *Philosophy* **70**: 395–407.

Brigandt, Ingo (2010). Scientific reasoning is material inference: Combining confirmation, discover, and explanation, *International Studies in the Philosophy of Science* 24(1): 31–43.

Campaner, Raffaella and Maria Carla Galavotti (2007). Plurality in causality, in Peter Machamer and Gereon Wolters (eds), *Thinking About Causes: From Greek Philosophy to Modern Physics*, Pittsburgh (PA): University of Pittsburgh Press: pp. 178–99.

Cartwright, Nancy (1999). *The Dappled World*, Cambridge: Cambridge University Press.

Cartwright, Nancy (2003). Causation: One word, many things, *Philosophy of Science* 71: 805–19.

Cartwright, Nancy (2007). *Hunting Causes and Using Them*, Cambridge: Cambridge University Press.

Darden, Lindley (2002). Strategies for discovering mechanisms: Schema instantiation, modular subassembly, forward and backward chaining, Supplement to *Philosophy of Science* 69(3), vol. 2: S354–S365.

De Vreese, Leen (2006). Causal pluralism and scientific knowledge: An underexposed problem, *Philosophica* 77 (1): 125–150.

Dowe, Phil (2000). *Physical Causation*, New York: Cambridge University Press.

Godfrey-Smith, Peter (2009). Causal pluralism, in Helen Beebee, Peter Menzies and Christopher Hitchcock (eds), *Oxford Handbook of Causation*, Oxford: Oxford University Press: pp. 326–37.

Hall, Ned (2004). Two concepts of causation, in John Collins, Ned Hall and Laurie Paul (eds), *Causation and Counterfactuals*, Cambridge (MA): MIT Press: pp. 225–76.

Henry, Sara Hale, F. Xavier Bosch, Terry C. Troxell and P. Michael Bolger (1999). Reducing liver cancer – Global control of aflatoxin, *Science* **286**(5449): 2453–4.

Hitchcock, Christopher (2003). Of Humean bondage, *British Journal for the Philosophy of Science* **54**(1): 1–25.

Hitchcock, Christopher (2007). How to be a causal pluralist, in Peter Machamer and Gereon Wolters (eds), *Thinking About Causes: From Greek Philosophy to Modern Physics*, Pittsburgh (PA): University of Pittsburgh Press: pp. 200–21.

Kurtzke, John (1977). Geography in multiple sclerosis, *Journal of Neurology* **215**(1): 1–26.

Longworth, Francis (2006a). *Causation, Counterfactual Dependence and Pluralism*. PhD Thesis. University of Pittsburgh.

Longworth, Francis, (2006b). Causation, Pluralism and Responsibility, *Philosophica* **77**(1): 45–68.

Machamer, Peter (2004). Activities and causation: The metaphysics and epistemology of mechanisms, *International Studies in the Philosophy of Science* **18**(1): 27–39.

Machamer, Peter, Lindley Darden and Carl Craver (2000). Thinking about mechanisms, *Philosophy of Science* **67**(1): 1–25.

Mackie, John (1974). *The Cement of the Universe: A Study of Causation*, Oxford: Oxford University Press.

McDermott, Michael (1995). Redundant causation, *British Journal for the Philosophy of Science* **46**: 523–44.

McGuinness, Brian (1985). Wittgenstein and the Vienna Circle, *Synthese* **64**: 351–8.

Medina, José (2001). Verificationism and inferentialism in Wittgenstein's philosophy, *Philosophical Investigations* **24**(4): 304–313.

Norton, John (2003). A material theory of induction, *Philosophy of Science* **70**(4): 647–670.

Peregrin, Jaroslav (2006). Meaning as inferential role, *Erkenntnis* **64**: 1–35.

Psillos, Stathis (2004). Glimpse of the secret connection: Harmonizing mechanisms with counterfactuals, *Perspectives on Science* **12**(3): 288–319.

Psillos, Stathis, forthcoming, Causal pluralism, unpublished manuscript, University of Athens.

Ramagopalan, Sreeram and Gavin Giovannoni (2009). Can we predict MS?, *The Lancet Neurology* **12**(8): 1077–79.

Ramagopalan, Sreeram, Narelle Maugeri, Lahiru Handunnetthi, Matthew Lincoln, Sarah-Michelle Orton, David Dyment, Gabriele DeLuca, Blanca Herrera, Michael Chao, Dessa Sadovnick, George Ebers and Julian Knight (2009). Expression of the multiple sclerosis-associated MHC class II allele HLA-DRB1*1501 is regulated by vitamin D, *PLoS Genetics* **5**(2): e1000369. doi:10.1371/journal.pgen.1000369.

Reiss, Julian (2009). Causation in the social science: Evidence, inference, and purpose, *Philosophy of the Social Sciences* **39**(1): 20–40.

Reiss, Julian, forthcoming, Review of Daniel Steel's Across the Boundaries: Extrapolation in Biology and Social Science, *Economics and Philosophy*.

Russo, Federica and Jon Williamson (2007). Interpreting causality in the health sciences, *International Studies in the Philosophy of Science* **21**(2): 157–70.

Schaffer, Jonathan (2004). Causes need not be physically connected to their effects: The case for negative causation, in Christopher Hitchcock (ed.), *Contemporary Debates in Philosophy of Science*, Oxford: Blackwell: pp. 197–216.

Sober, Elliott (2009). Reichenbach's cubical universe and the problem of the external world, *Synthese* (online first).

Steel, Daniel (2008). *Across the Boundaries: Extrapolation in Biology and Social Science*, Oxford: Oxford University Press.

van Amerongen, B.M., C.D. Dijkstra, P. Lips, C.H. Polman (2004). Multiple sclerosis and vitamin D: an update, *European Journal of Clinical Nutrition* **58**(8): 1095–109.

van Etten, E., D.D. Branisteanu, L. Overbergh, R. Bouillon, A. Verstuyf and C. Mathieu (2003). Combination of a 1,25-dihydroxyvitamin D(3) analog and a bisphosphonate prevents experimental autoimmune encephalomyelitis and preserves bone, *Bone* **32**: 397–404.

van der Mei, I., A.-L. Ponsonby, L. Blizzard, and T. Dwyer (2001). Regional variation in multiple sclerosis prevalence in Australia and its association with ambient ultraviolet radiation, *Neuroepidemiology* **20**: 168–74.

van der Mei, I., A.-L. Ponsonby, T. Dwyer, L. Blizzard, R. Simmons, B.V. Taylor, H. Butzkueven, and T. Kilpatrick (2003). Past exposure to sun, skin phenotype, and risk of multiple sclerosis: case-control study, *British Medical Journal* **327**: 1–6.

Weber, Erik (2007). Conceptual tools for causal analysis in the social sciences, in Federica Russo and Jon Williamson (eds), *Causality and Probability in the Sciences*, London: College Publications, pp. 192–207.

Wild, Christopher and Ruggero Montesano (2009). A model of interaction: Aflatoxins and hepatitis viruses in liver cancer aetiology and prevention, *Cancer Letters* **286**: 22–8.

Williamson, Jon (2006). Causal pluralism versus epistemic causality, *Philosophica* **77**(1): 69–96.

Wittgenstein, Ludwig (1953). *Philosophical Investigations*, Trans. Elizabeth Anscombe. Oxford: Blackwell.

Wittgenstein, Ludwig (1979). *Wittgenstein's Lectures, 1932–35*, edited by Alice Ambrose, Oxford: Blackwell.