Studies in History and Philosophy of Biological and Biomedical Sciences xxx (2012) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Studies in History and Philosophy of Biological and Biomedical Sciences

journal homepage: www.elsevier.com/locate/shpsc

Causation in the sciences: An inferentialist account

Julian Reiss

Erasmus University, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Causation
Inference
Evidence
Objectivity
Biomedical sciences
Social sciences

ABSTRACT

I present an alternative account of causation in the biomedical and social sciences according to which the meaning of causal claims is given by their inferential relations to other claims. Specifically, I will argue that causal claims are (typically) inferentially related to certain evidential claims as well as claims about explanation, prediction, intervention and responsibility. I explain in some detail what it means for a claim to be inferentially related to another and finally derive some implication of the proposed account for the epistemology, semantics and metaphysics of causation.

1. Introduction

That the concept of 'cause' is of extraordinary importance for the biomedical, social and other sciences hardly needs explanation or defence these days. Of course, this wasn't always the case. We all remember reverberations of Bertrand Russell's 'the law of causality is . . . a relic of a bygone age', just like the monarchy (Russell, 1913) and of Karl Pearson's view of causation as 'another fetish amidst the inscrutable arcana of even modern science' (Pearson, 1911, p. vi).

There is another difference between now and then. When Russell and Pearson carried out their attacks on the notion of 'cause' they knew what they were attacking because there was a dominant theory of what causation was, *viz.* some special kind of regularity. A hundred years ago, then, philosophers knew what causation was, they just disagreed on whether the idea was useful and important.

Today we are facing the opposite predicament. Philosophers of causation no longer stand in need of justifying their academic interest. Even philosophers of physics, whose views on causation are traditionally close to Russell's, seem to have discovered the notion as useful and important for a variety of cognitive and practical purposes (Price & Corry, 2007). However, while a virtually universal agreement as to its philosophical significance has emerged, it became increasingly unclear what we are talking about when we are talking about causation.

There are no less than five families of 'standard views' on causation: regularity, counterfactual, probabilistic, process/mechanist and agency/interventionist (see for instance Reiss, 2009 for a discussion). Each of these aspires to provide a full-fledged theory of causation, but each of these is subject to counterexamples that prove quite recalcitrant. One understandable response has been to try to improve the theory as much as possible, and to bite the bullet with respect to the remaining cases. Another response, equally understandable, has been to abandon the project of trying to find an account that gives necessary and sufficient conditions of application and to become a 'causal pluralist' instead. Neither response is very satisfactory. There is obviously something undesirable in ignoring counterexamples or classifying a case as (not) being one of causation against widely held convictions. Pluralism, in turn, raises more questions than it answers. If causation is 'many things', how do we explain that there is just 'one word' (Cartwright, 2004; this point was raised in Russo & Williamson, 2007)? If there are many concepts of cause, how many are there, how do we know, and what is it that makes all of them concepts of *cause*?

This paper aims to advance this debate by providing a more rational basis for addressing issues concerning mainly the semantics of causation, but in its wake also its epistemology and

E-mail address: reiss@fwb.eur.nl

1369-8486/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.shpsc.2012.05.005

metaphysics. That rational basis is a version of inferentialism, a semantic theory that regards the meaning of a sentence as constituted by their inferential connections. I understand inferentialism about causation less as a mature *theory* of causation than as a theoretical framework within which one can address various issues concerning causation in a philosophically fruitful way. In order to give some substance to this claim I will discuss three issues in some detail below: the relation between causal claims and evidence for these claims; conceptual monism versus pluralism; the objectivity of causation.

To say it upfront, I am not trying to argue directly in favour of inferentialism as a theory of the meaning of causal claims, and even less in its favour as a theory of meaning in general. Rather, what I aim to demonstrate is the fruitfulness of inferentialism when applied to causal claims. This may be read as an indirect argument in favour of inferentialism but the focus here is on the substantive issues concerning the semantics, epistemology and metaphysics the framework helps to address.

2. Inferentialism

Hume bequeathed us not one but two theories of causation (Beauchamp, 1974; Beauchamp & Rosenberg, 1981). The better known of the two is the regularity theory according to which an event *X* causes another event *Y* if and only if *Y* regularly follows *X* and the two are contiguous. According to the other theory, *X* causes *Y* if and only if the occurrence of an *X* leads the observer to anticipate *Y*, to expect *Y* to follow. The second account may be called 'necessity' (Beauchamp, 1974; Beauchamp & Rosenberg, 1981), 'associationist' (Spohn, 1993) or 'projectivist' (Beebee, 2007a, 2007b) theory.

The second theory can be given an inferentialist reading. Accordingly, to grasp the semantic content of '*X* causes *Y*' is to infer that *Y* will happen upon learning that *X* has, or that *X* had already happened upon learning that *Y* has. Thus, Hume's second theory (in an inferentialist reading) remains wedded to a regularity view of causation: '*X* causes *Y*' means to infer *Y* to follow *whenever X* has happened; and to infer *X* to have happened *whenever Y* has.

Few people today endorse a regularity account of causation. To see what is wrong with it, take a philosophers' favourite example: 'Smoking causes lung cancer'. We would be ill-advised to expect every smoker to develop lung cancer because in fact only a small proportion of smokers do; we'd be equally ill-advised to blame every individual's contraction of lung cancer on smoking because a substantial proportion of sufferers never smoked. Nevertheless, to grasp the semantic content of 'Smoking causes lung cancer' can be understood as making certain inferences, for instance about the probability of developing the disease conditional on smoking.

Causation and inference therefore remain closely related. An inferentialist semantics turns that relation into a theory of meaning. It holds that the meaning of a sentence or utterance is constituted by its inferential relations to other sentences. What this amounts to when applied to causal claims is that the meaning of a claim such as 'Smoking causes lung cancer' is constituted by the propositions with which it is inferentially related such as 'Smoking raises the probability of lung cancer' and 'If John hadn't smoked, his chance of developing lung cancer would have been considerably smaller'.

This general inferentialist account has to be made more precise in at least two ways. We should, first, say more about the 'other sentences' with which a causal claim is inferentially connected. What kinds of sentences are these typically? We should, second, say more about the inferential relation. What kind of inferential relation is it?

Before giving answers to these questions, I want to limit the investigation in two ways. First, I will only consider *generic* causal claims of the form '*X* causes *Y*', where both *X* and *Y* are variables (or event-*types*). I shall not be concerned with claims about actual causation. Several philosophers believe that actual causation is the more basic relation (among others, Cartwright, 1989; Lewis, 1973; Spohn, 2006). That may well be true but the inferential treatment of actual causal claims will have to wait until another day. What is clear is that the *relation* between generic and actual causal claims is not at all straightforward (see Hitchcock, 1995). An account of generic causal claims is therefore unlikely to be made redundant immediately by advances on the actual front.

Second, I will focus on causal claims as they are typically found in the social and biomedical sciences. I am not interested *per se* in conceptual analysis but rather in an account of causation that helps us to understand scientific practice in these domains.

2.1. Inferential base and target

The first question concerns the kinds of sentences with which a claim of interest is inferentially related. Here I take my cue from Wittgenstein. Wittgenstein, in his post-Tractatus period, was one of the originators of the inferential conception of meaning. He remarked (Ambrose, 1979, pp. 19-20):

> Some people say that the question, "How can one know such a thing?," is irrelevant to the question, "What is the meaning?" But 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, which is what the question, "What would it be like for it to be true?," asks for.

Focus on the middle sentences first. The meaning of a sentence is given by the sentences from which it follows and that follow from it. Let us call these sentences the 'inferential base' and 'inferential target', respectively. What constitutes typical inferential bases and targets?

The inferential base for a causal claim is, of course, given by sentences constituting or describing the evidence for it (depending on whether one thinks of evidence as consisting of sentences itself or rather as non-linguistic entities such as facts). Descriptions of specific applications of methods of causal inference such as experiments, randomised trials, regressions, applications of structural equations models, Bayes' nets, expert judgements, meta-analyses and so on are all possible members of the inferential base for a specific causal claim. I will say more about this in Section 3.

Causal claims are not normally established for their own sake but rather because they are considered useful for the attainment of certain purposes. These purposes are sometimes of a cognitive nature, relating to the scientific explanation of outcomes of interest or of regularities. To give an example, in the period 2006-2008 the world experienced an unprecedented rise in food prices. In roughly the same period, the use of commodities (including foods) by financial investors has increased greatly. Does the speculators' interest in commodities explain the food price inflation? If we had evidence to the effect that commodity speculation causes food price inflation, we could use the knowledge of this causal claim in conjunction with other items of knowledge (for instance, that there are no other factors that fully explain the food price inflation, that nothing happened that offset the causal effect of speculation and so on) to infer the explanatory sentence 'Speculation explains/helps to explain the food price inflation 2006-2008' (Gilbert, 2010).

In econometric circles, such a claim might be inferred for purely scientific reasons. Other actors will want to hold speculators responsible for high food prices and ensuing starvation and

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

malnutrition. They will want to infer an evaluative claim concerning praise or blame from the causal claim: 'High food prices (and with them in part also starvation and malnourishment) can be blamed on speculators'. The causal claim would have to be conjoined with other sentences for this inference (for instance, speculators are human agents and not compelled in their actions) but the claim about moral responsibility would nevertheless be inferentially connected with the causal claim.

Finally, causal claims allow inferences concerning various practical endeavours such as prediction and control. If it is true that commodity speculation is a cause of episodes of famine and malnourishment, policy makers might consider making speculation with certain kinds of essential commodities illegal in order to help prevent future famines or make them less likely.

Sentences relating to the cognitive, evaluative and practical content of a causal claim form what I call its 'inferential target'. They are the target because these kinds of sentences are what we ultimately seek in our causal investigations; they constitute the 'cash value' of knowing causal claims. Together, the inferential base and target of a causal claim constitute the *inferential system* for that claim.

2.2. Entitlement to infer

The second question concerns the nature of the inferential connection. I propose the following formulation: a sentence is inferentially connected with a causal claim (and therefore part of the claim's inferential system) if and only if the scientific community that asserts the claim is entitled to infer it from the causal claim or to infer the causal claim from it (in conjunction with other sentences the community holds). Three aspects of this characterisation require commenting: the nature of the scientific community who is entitled to make the inferences; the nature of the entitlement; and the nature of the inferences that are being made. For reasons that will become plain shortly, I will go through these back to front.

2.2.1. The nature of the inferences

There are, broadly speaking, two models of the nature of inference. According to one model, all valid inferences are licensed by the logical form of the propositions that are inferentially connected with each other. This model, accordingly, holds that all inferences are *formal* inferences. Deductive inferences are formal inferences. For instance when we infer, by universal instantiation, from 'All men are mortal' and 'Michael Jackson is a man' that 'Michael Jackson is mortal', then the validity of the inference does not depend on what concepts are used in place of 'men', 'mortal' or 'Michael Jackson'. We could substitute, say, 'mortal' by 'five feet ten inches tall' and validly infer that Michael Jackson was five feet ten inches tall. Such substitutions will often affect the soundness of a reasoning episode, as in this case, but not their validity.

As adequate as the formal model of inference is for deductive reasoning, as problematic does it become when inductive inferences are concerned. This is because of the fact that there are no universally justifiable forms or schemas of inductive inference (*cf.* Norton, 2003). Consider just two examples, both of which are very powerful in some instances, and hopeless in others. The first is inductive generalisation: the inference from 'Some A are B' to 'All A are B'. This form of inference is justifiable depending on whether the A's belong to a class of things that are very homogeneous such as natural kinds as well as on the nature of the B's.

When the A's are protons, for example, we are licensed to infer from 'Some protons have a mass of $1.672621777{*}10^{-27}$ kg' to 'All protons have a mass of $1.672621777{*}10^{-27}$ kg'. When the A's are bars of chocolate, by contrast, we should not infer from 'Some chocolate bars have a mass of $1{*}10^{-1}$ kg' that 'All chocolate bars have a mass of $1{*}10^{-1}$ kg'. Note that the nature of the B's is important too. Whereas we are not licensed to infer from 'Michael Jackson, Patrick Swayze, Matt Damon, Eddie Murphy are 5′10″ tall' that 'All men are 5′10″ tall', it is not unreasonable to infer from the mortality of myriads of men in the past that all men are mortal.

Similar considerations apply to analogous inferences. They have the general form: 'A is x'; 'B is like A'; therefore, 'B is x'. If A is 'a cobra', x is 'lays eggs' and B is 'a python', then we are licensed to make the inference on the grounds that both A and B are alike in that they are both types of snake. The inference is not licensed if x is 'is poisonous' instead. Not very many snakes are poisonous, and therefore being a snake on its own doesn't provide sufficient grounds for the analogical inference in this case.

Norton, 2003 shows that all schemas of inductive inference that have been proposed, including Bayesianism, the error-statistical approach and inference to the best explanation either suffer from counterexamples such as the above or else are vacuous. He proposes a 'material account of induction' in which all induction is local and grounded, not in the form of the propositions involved, but rather in matters of fact that hold only in specific domains.

Norton's is a specific instance of material accounts of inference that include Wilfried Sellars' (Sellars, 1953), Robert Brandom's (Brandom, 1994, 2000) and, very recently, Ingo Brigandt's (Brigandt, 2010). According to material accounts of inference, inferences are licensed (that is, supported or warranted) by the content of the concepts that figure in the propositions which are inferentially connected; thus, whether an inference is licensed or not depends on the content of the concepts the inferences are made about.

There are three main differences between the material and the formal accounts that worth pointing out here. First, because material inferences derive their warrant from the content of the concepts that are involved, inferential warrant is not normally preserved upon substitution of concepts. According to the material model, then, all inferences are like inductive inferences. Deductive inferences are a special case in that their warrant depends only on the logical terms that figure in the inferentially connected propositions. Deduction is a special case of induction, not the other way around.

Second, the formal model of inference regards most ordinary inferences as enthymematic. Thus, to infer 'The iron is being heated up to 1000 °C, therefore it will glow' is a formally invalid inference whose validity can be restored by adding a premiss such as 'Whenever iron is being heated up to 1000 °C, it will glow'. The material account, by contrast, regards the original inference as warranted without the addition. It is warranted because it is part of the content of the concept 'iron' that it will glow at 1000 °C.

This does not mean that all inferences are trivially warranted just because someone makes them. Rather, the content of a concept is a fact about what is known about the concept at the time the inference is made. Therefore, a language user may be mistaken in making this or that inference about propositions containing certain concepts because the content of the concepts do not license the inference. That is, he may not be *entitled* to make the inference. I will say more about the nature of the entitlement in the next subsection.

Third, the kind of inference will, according to the material account, also depend on the nature of the concepts and propositions involved. In some cases (say, when chemical elements are involved and we are inferring the melting point of all samples from observations of the melting point of a small sample), the inference will be categorical (albeit subject to constant air pressure and other conditions). In other cases, the inference will be probabilistic (say, when an inference is made from a statistical proposition to a single case). Most inferences will be subject to *ceteris paribus* clauses, but whether or not this is the case will also depend on the concepts

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

involved. Formal accounts of induction, by contrast, model all inductive inferences as instances of a unique type.

Before turning to the nature of the entitlement on which the inferentialist theory is based, let me point out that adopting a material model of inference does not beg the question concerning any of the more specific questions regarding causation I will address later on. In particular, the material inference account does not by itself answer the question whether 'cause' operates more like a term of logic—as most philosophers would presume—or more like an empirical concept such as gene. All the material account does is make that question an *a posteriori* one. In order to see what kind of concept 'cause' is we have to attend to the material inferences scientific communities are entitled to make when they use the word. It may well be the case that once we analyse these inferences, 'cause' works analogously to 'implies' or some other term of logic. That is, it may well turn out that 'cause' entitles language users to inferences independently of the other concepts that figure in a causal claim and of other background beliefs. But the reverse may also be the case: that 'cause' depends in its content on the causal relata, and possibly other terms and background beliefs. Either way, according to the material account, this would be an empirical discovery. Similarly with the kind of inference licensed. *A priori*, the term 'cause' might license inferring a deterministic claim ('doing *X will* result in *Y*'), a probabilistic claim ('doing *X* will result in *Y* with a probability *p*' or 'doing *X* will *probably* result in *Y*' or 'doing *X might* result in *Y*'), or a *ceteris paribus* claim ('doing *X* in the right conditions will/will probably/might result in *Y*'). Which one of these inferences is in fact licensed by the term 'cause' (allowing for the possibility that different inferences are licensed on different occasions) is an empirical discovery. None of this should be very controversial: I take this to be the central (and still correct) insight of Hume's contribution to the philosophy of causation.

2.2.2. The nature of the entitlement

When should we regard a community of language users to be entitled to making certain inferences? There are two traditional answers to epistemic justification, both of which are equally unsatisfactory: transcendentalism and descriptivism. The transcendentalist believes that there are external global standards of justification that can be grasped by philosophers and others and used to evaluate knowledge claims independently of context and purpose of the investigation. Transcendental epistemic standards concerning causation have proved very elusive. That is, no global standards of evaluation have been found that do not involve large numbers of counterexamples. Take the evidence-based medicine movement's view that randomised clinical trials (RCTs) are the 'gold standard' of evidence for causal claims in medicine. If one took that to be an external standard for justifying causal claims, many claims that are in fact accepted by the medical community would come out as unjustified: be it the implication of smoking in lung cancer, the efficacy of acetylsalicylic acid to relieve headaches or the efficacy of parachutes to prevent deaths by crushing into the ground. Similarly, it is often pointed out that if the medical profession were to accept only RCT-based results (prospectively speaking), it would harm its ability to make sound public health decisions (*e.g.*, Worrall, 2002 on the Extra-Corporeal Membraneous Oxygenation or ECMO case; Lie & Miller, 2011 on circumcision and HIV). In other words, taking an external standard for the justification of causal inferences would make the account descriptively inadequate and epistemically problematic.

Descriptivism of course does not suffer from the problem of descriptive inadequacy. It holds, roughly speaking, that whatever standards a scientific community holds conveys epistemic justification. The problem with that view would be that a scientific community could not be mistaken in adopting this or that standard. Suppose the medical community did adopt the evidence-based movement's view that RCTs are the 'gold standard', where the latter is to be interpreted as at least necessary for justified causal inferences. On a descriptivist view, there would be no way to criticise the medical community for doing so and appraisals both inside the profession such as Jan Vandenbroucke's (*e.g.*, 2004) or from the outside such as Nancy Cartwright's (*e.g.*, 2007) or John Worrall's (*e.g.*, 2002) would be futile.

There is a third way, however, and that is contextualism (see for instance Kincaid, 2004; Reiss, 2008). The contextualist rejects the idea that there are global standards to which every justifiable knowledge claim has to conform. Standards are always local and contextual. They are local in that they apply first and foremost to specific episodes of scientific reasoning, which obtain in particular fields, periods and sometimes regions, and not to science as a whole. They are contextual in that they are relative to specific purposes of the inquiry, the questions asked, and the background knowledge that can be presumed.

Contextualism avoids the main drawbacks of transcendentalism and descriptivism. By allowing standards to vary between domains, periods and specific queries, it is likely to be more descriptively adequate than transcendentalism. At the same time it does not lose all normative force as descriptivism does. Criticisms are possible as long as they are contextually motivated. One can, for example, point out that a scientific practice is unable to address certain kinds of questions (which would have critical bite of course only to the extent that it is these questions that were salient at the time) or that certain things were already known at the point of the inquiry but ignored by the scientists, to the detriment of the aims of the inquiry or the well-being of affected people. Arguably, contextualism has a greater normative force than transcendentalism since the criteria the transcendentalist offers are often themselves hard to justify except in a circular, regressive or dogmatic manner. The contextualist dispenses with ultimate justification of course but the normative claims he does make have a much thicker network of mutually enforcing vindications than the transcendentalist could offer, simply because a justification, if it can be given at all, has to respect local empirical facts, background knowledge and the goals and purposes of the query.

According to the contextualist, then, whether a scientific community is entitled to make an inference depends on what the background beliefs and the purposes and goals of the investigation are, and on what is known about the evidence and the causal relation in question. Contextualism too does not beg the question vis-a-vis someone who holds that there are certain features that all causal relations have in common. If that were the case, contextualist inquiries will have that result: justifying causal claims on the basis of evidence, or claims that are in the inferential target on the basis of a causal claim is independent of the nature of the causal claim, background beliefs and purposes of the query. However, it would be a result of an investigation, not an assumption.

2.2.3. The scientific community

An aspect of the inferentialist theory of meaning as it has been introduced so far which may be considered troublesome is that meaning is a thoroughly social affair. This is certainly troublesome for the Humean who seeks the meaning of 'cause' inside his own head. To Hume, for a concept to be meaningful, it had to be associated with an idea, which itself was a remembered sense impression. 'Cause' was meaningful only to the extent that it was seen or felt. Humean associationism has long been given up as a general theory of meaning. It is about time that associationism concerning causation be given up too.

In the framework proposed here meaning is social because entitlement is, and entitlement is social because there is no guarantee that an individual holds all the beliefs that are jointly necessary to

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

warrant an inference. This, I believe, is true quite generally, but it is easiest to see when one considers a feature characteristic of at least some modern science. Causal claims in the social and, more ostensibly, the biomedical sciences are often established by what has come to be known as 'big science': large numbers of researchers who bring different kinds of expertise to a common research project. Drug trials, for instance, are now frequently conducted in multiple centres, sometimes located in more than one country. The same is now true of the randomised field evaluations that have risen in prominence over the last years.

In 'big science' no individual researcher has all the evidence required to justify knowledge claims. Rather, numerous scientists are required, each one with his or her specific expertise. By and large, then, it is not individual researchers who are justified in making an inference but rather communities of researchers. A direct consequence is that the content of concepts is dispersed across individuals and therefore no one researcher is in full possession of the concept. The meaning of scientific concepts too is one that is dispersed across many researchers.

3. Epistemology, semantics, metaphysics

The proof of any pudding is in the eating, as we all know, and the pudding of causation is no exception. In this section I will show how the inferentialist framework I have outlined above can be used to address some questions that have vexed philosophers of causation in recent years. I will take a selection of questions from the epistemology, the semantics and the metaphysics of causation as examples, formulate them within the inferentialist framework and sketch how they might answered within the framework. To say it upfront, I will *not* argue in favour of a specific answer to each problem. This is for the simple reason that because of the nature of the inferentialist account, these specific answers depend on empirical facts, and therefore to provide a specific answer requires detailed case-based work, which is beyond the scope of this paper. I nevertheless hope that this paper will advance the debate by describing how to address important issues concerning causation within an inferentialist framework and thus providing the reader with useful philosophical *tools*. My own answers to these questions will have to wait for another occasion.

3.1. Causation and evidence

Let us call the regularity, probabilistic, counterfactual, agency/manipulability and process theories the 'standard accounts' of causation (see for instance Beebee, Hitchcock, et al., 2009). All standard accounts are 'verificationist' theories (or developments of verificationist theories) in that they take conditions under which a causal claim can be *tested* for its truth or falsity to provide the *meaning* for the claim. Under the right conditions ('*ceteris paribus*'), many causal claims entail claims about regularities, probability raising, counterfactual dependence, manipulability and transmission of conserved quantities. It is only natural, then, to look for regularities *etc.* in order to verify causal claims. But to identify a verification criterion for a claim with its meaning is to make the same mistake as the logical positivists and other verificationists made and it runs into the same kinds of trouble. One well-known source of trouble for the verificationist is that there are often multiple ways in which a claim can be verified, a fact which is turned into a mystery under a verificationist conception of meaning. Verificationist theories of causation suffer from exactly this problem.

Causal claims can be supported by a great variety of evidence: controlled and randomised trials, (prospective/retrospective) observational studies, case reports and process tracing, thought experiments and analyses of natural experiments, implications derived from theory. Verificationist theories of causation cannot explain this fact. If the meaning of 'causes' is exhausted by 'is regularly associated with', why is it the case that many causal claims are supported by nothing but a singular execution of a well-designed controlled experiment (*e.g.*, Cartwright, 1989)? Moreover, why is it the case that experiments are almost never repeated but often replicated, where 'repetition' refers to the repeating of an experiment under the exact same conditions and 'replication' the execution of a related experiment under similar but not identical conditions? If invariance under intervention is what causation means, why is it the case that causal claims, for instance in epidemiology and the social sciences, are supported by observational evidence where no interventions have been performed? If probability raising is what causation means, then why do historians support causal claims by performing a thought experiment about singular events where no probabilities are involved? And so on.

Verificationist theories of causation turn all these commonplaces into unsolved problems. Let me illustrate the issue by Woodward's interventionist account. Clearly, Woodward's is an account of the meaning of causal claims (Woodward, 2003, p. 38):

> My aim is to give an account of the content or meaning of various locutions, such as X causes Y, X is a direct cause of Y, and so on, in terms of the response of Y to a hypothetical idealized experimental manipulation or intervention on X.

He also notes in various passages in the book that experimental manipulation is not necessary for verifying causal claims. For instance (*ibid.*, pp. 35-6; footnote suppressed):

> There are cases in which one can learn more about causal relationships (assuming one is willing to make certain other assumptions) from observation or from a combination of observation and experiment than from practically possible experiments alone, and there are many cases in which, for moral or practical reasons, one must rely on nonexperimental evidence to reach causal conclusions. A plausible manipulability theory will not deny that reliable causal inference on the basis of nonexperimental evidence is possible, but rather, suggests a specific way of thinking about such inferences: we should think of them as an attempt to determine (on the basis of other kinds of evidence) what the results of a suitably designed hypothetical experiment or manipulation would be without actually carrying out this experiment. For example, for moral and political reasons, among others, one cannot carry out experiments in which some children are randomly assigned to public and others to private school and their subsequent academic careers observed. Nonetheless, it is illuminating to think of attempts to infer from nonexperimental data the effects of private schooling on achievement as attempts to predict what the results of such an experiment would be without actually doing it.

Woodward here merely claims that 'reliable causal inference on the basis of nonexperimental evidence is possible', tells us how to understand a claim inferred from non-experimental evidence but he owes us an explanation for why this should be so. Inferentialism offers a very simple explanation: multiple kinds of evidence are regularly in the inferential base for causal claims. That scientific communities (and indeed, other language users) are entitled to infer a causal claim from a variety of kinds of evidence I take to be too obvious to argue for it explicitly.

According to the inferentialist account of causation proposed here, the meaning of a causal claim is not exhausted by its verification conditions but the two are nevertheless closely related. How they are related is described by José Medina in an article on the

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

theory of meaning Wittgenstein held between the Tractatus and the Philosophical Investigations in which he explains the passage from Wittgenstein I quoted earlier (Medina, 2001, p. 308; emphasis original):

> 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?" us 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." [Ambrose, 1979, pp. 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 in part determines what these inferential connections are. It is a consequence and not an assumption of this view that propositions describing different kinds of evidence frequently happen to be in the inferential basis of a causal claim.

One specific question one might want to ask is whether different kinds of evidence—such as 'probabilistic evidence' about the probabilistic relations among causal variables of interest and 'mechanistic evidence' about the causal process that connects a cause with its effect—are (normally) required for establishing a causal claim. Federica Russo and Jon Williamson 2007 argued that this is the case in the health sciences and suggested that it might be true also elsewhere such as in the social sciences (see also Russo & Williamson, 2011). This claim has later been called the 'Russo-Williamson Thesis' and discussed quite widely (e.g., Gillies, 2011; Illari, 2011, Claveau this volume).

Within the inferentialist framework proposed here the Russo-Williamson Thesis would be formulated thusly: In the health sciences (and possibly elsewhere), probabilistic and mechanistic evidence must (normally) be in the inferential base for a causal claim. In order to establish it inductively, we would have to take a representative sample of causal claims accepted within the respective sciences and show that the relevant communities of researchers are (normally) entitled to inferring the causal claim only when the inferential base contains sentences describing both probabilistic and mechanistic evidence. In order to refute the thesis, we have to find a range of cases where a scientific community is entitled to make the causal inference on the basis of one type of evidence alone.

Biomedical researchers in fact do accept a wealth of causal claims despite lacking understanding of the underlying mechanism responsible for a causal relation (or in the jargon of clinicians, despite the lack of 'biological rationale'). For instance, lithium has been accepted to be effective in the treatment of mania since the late 1960s (the FDA approved of its use in 1970), and yet we still read in Wikipedia: 'The precise mechanism of action of Li+ as a mood-stabilizing agent is currently unknown' (entry 'Lithium pharmacology', accessed on 17. November 2011).

It is likely that at least sometimes communities of biomedical researchers are entitled to infer the causal claim despite the lack of biological rationale. Generally speaking, a community of researchers can be said to be entitled to infer a causal claim from the evidence whenever all alternative explanations of the evidence can be ruled out. What that means depends on the case. A measured correlation between two variables can for instance be spurious, by which I mean that it can be explained by selection bias, measurement error or other biases in the inference from sample to population. Or it may be genuine but explainable by reverse causation or a common cause. And so on. It is unlikely that the ruling out of alternatives must always proceed on the basis of evidence about the mechanism of action. But that's a matter of empirically proof. *A fortiori*, what is normally the case is a matter of empirical proof, and is therefore beyond the scope of this paper.

3.2. The concept of 'cause'

How many concepts of 'cause' are there? Whether causal pluralism is true or not, and if true, what kind of pluralism we should endorse is another issue that has been intensely debated in the recent literature (*e.g.*, Campaner & Galavotti, 2007; Cartwright, 2004; Godfrey-Smith, 2009; Hall, 2004; Longworth, 2006; Reiss, 2011; Williamson, 2006a, 2006b). Here I should be concerned mainly with what Chris Hitchcock, 2007 calls 'extramural pluralism' about causation, *viz*. the idea that different accounts of causation (such as the probabilistic, the mechanist, the counterfactual *etc*. accounts) define different concepts of causation; that '*X* causes *Y*' when understood probabilistically refers to something else than '*X* causes *Y*' when understood mechanistically or probabilistically and so on. Hall, 2004, for instance, is a pluralist in that he believes that there are two fundamental ideas of cause, *viz*. 'production' and 'dependence'. Woodward, 2003, by contrast, is a monist in that he thinks all causal relations can be understood as relations that are invariant under hypothetical interventions.

On the face of it, the inferentialist framework presented here makes 'cause' an abundantly pluralist notion because meaning of 'cause' varies dramatically with the sentence in which it is used and the historical context in which it is uttered. 'Cause' means something different when it appears in 'Smoking causes lung cancer' from when it appears in 'Lack of sunlight causes multiple sclerosis (MS)', and when each claim is uttered in a different historical context simply because there are different sentences in the inferential base and target on each occasion. Let me make two remarks about this plurality of meanings.

On the one hand, this multiplicity and un-fixedness of the meaning of 'cause' may seem unusual given most philosophers of causation have tried to reduce it to one or a small number of conditions of application, but it is by no means implausible. Cigarette smoke, for instance, contains some 60 carcinogens which affect lung tissue in a variety of ways including a suppression of the lung's efficiency to remove particulate matter by bringing about alterations in the ciliary activity of the bronchial mucosa and DNA damage. By contrast, sunlight exerts a protective effect on people by allowing skin exposed to UV-B radiation to synthesise vitamin-D, which is involved in a variety of biological functions including immune responses. How precisely it affects the onset of MS is unknown but studies with mice have shown that a derivative of the vitamin successfully prevents the onset of experimental autoimmune encephalomyelitis, which is recognised as useful animal model for MS. Lack of sunlight therefore causes by the absence of a factor that is involved in a protective mechanism for the disease. Smoking and lack of sunlight cause in quite different ways. Hence, there is some initial plausibility that the meaning of the term differs when embedded in different sentences.

That the historical context in which a causal claim is uttered matters is also not implausible. When the mechanism of action for a(n already accepted) causal relation can been established, a host of new research avenues will open up, and the claim can be put to different uses such as, for instance, giving kinds of explanations that previously could not be given. It is not implausible that after such a discovery the meaning of 'cause' changes in the same way as the meaning of, say, 'water' changed after it was discovered that it is H_2O or the meaning of 'electron' changed after its charge could be measured for the first time.

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

On the other hand, we may ask whether there are commonalities of meaning among the different occasions when the term 'cause' is used; that is, if there are 'family resemblances' among the different uses. And of course there are. All causal claims are to some extent similar because they all have similar sentences in their inferential base and target: sentences describing certain types of evidence in the former, and sentences expressing certain kinds of explanatory and evaluative claims and claims about prediction and policy interventions in the latter.

In Reiss, 2011 I argued that there are important differences among different causal claims because not every causal claim warrants inferring the same kinds of claims in the inferential target, and what is in the inferential target is determined in part by what is in the inferential base. In particular, I argued that communities of scientists are not always entitled to infer the claims 'There is a mechanism from *X* to *Y*' and 'An intervention in *X* will change *Y*' when they hold the causal claim '*X* causes *Y*'. Rather, whether there is entitlement to these inferences from the causal claim depends on what is in its inferential base. Thus, for instance, when only claims about a substance's mechanism of action are in the inferential base, the inference to a claim about interventions will often fail because the mechanism might be fragile or there are different mechanisms in place which mutually cancel (as is well known in the literature).

Independently of whether my earlier arguments are successful, what is incontrovertible is that the family resemblances we consider to hold among causal claims depend on the purpose of the investigation. For some purposes we may want to put more emphasis on the commonalities, for others on the differences, and different purposes will determine precisely what commonalities and differences we emphasise. Whether 'cause' is one thing or many will therefore receive a different answer depending on what the context is in which we ask the question.

3.3. The objectivity of causation

The inferentialist account of causation can also illuminate issues concerning the objectivity of causation. Whether causation is objective ('in the objects') or subjective ('in the mind of the beholder') has been an intensely debated issue since Hume wrote on the subject. Hume's own position is subject to interpretation among other things because, as mentioned above, we find at least two theories of causation in his writings. According to one, the regularity theory, causation is objective. Regularities hold between objective events occurring in the world. According to the second, the necessity/associationist/projectivist theory, for *X* to cause *Y* means for the observer of an occurrence of *X* to have the expectation that *B* will occur. Causation is subjective in that there are no causal relations without observing minds.

It is possible, however, to resolve that tension. According to one view (Beauchamp & Rosenberg, 1981), the subjective theory is conceptually more basic but it can be objectified by grounding expectations in objective regularities. Causation is therefore fundamentally subjective in that a mental feature (an expectation) makes a relationship causal, but to the extent that that mental feature is itself a product of objective regularities (which, according to Hume, it was to a large extent), causation imports a degree of objectivity.

The inferentialist account proposed here is similar to this interpretation of Hume's account. Causation is fundamentally subjective in that what warrants the appearance of the term 'cause' or a cognate in a claim is the claim's being inferentially related in certain ways to certain other kinds of claim, and to be inferentially related is to play a role in—human—reasoning practices. But, and this is what makes causation partly objective, these reasoning practices may not arbitrary but rather shaped by objective facts. These facts are not however facts about regularities. There is a fascinating history to be told about the reception of the regularity account in biomedical science (see for instance Codell-Carter, 2003) but it is no longer accepted, and probably for good reason.

To look for possible sources of objectivity, first consider the inferential base. We have seen above that (sentences describing) a variety of different kinds of evidence can be and frequently are in the inferential base for a causal claim. There is very little of substance one could say in general about what entitles a community of researchers to infer a causal claim from the evidence. One attempt would be: a community of researchers is entitled to infer a causal claim *CC* from evidence *E* if and only if it can rule out all alternative explanations (to *CC*) of *E*. What 'ruling out an alternative explanation' means depends on the nature of the evidence and the causal claim in question.

The problem from the point of view of trying to locate objectivity in the evidence-causal claim relation is that whether or not something is an alternative explanation of the evidence depends on what is known by the community and can therefore vary between communities. Suppose for instance that a correlation is in the inferential base for a causal claim '*X* causes *Y*'. Let us ignore the possibility of the correlation being spurious or due to sampling error or other non-causal explanations for the sake of the argument. Let us also ignore reverse causation. That leaves our causal claim and one alternative, *viz.* the explanation of the correlation by the existence of a common cause. I submit that the latter cannot be ruled out on the basis of objective evidence alone. This is simply due to logic of negative existential statements: there cannot be conclusive evidence that '*Z* does not exist', whatever *Z* is, and conclusive the evidence would have to be in order to be objective. This is not to say that one cannot rule out that alternative explanation. But that ruling out will be relative to what is known by the community.

One might counter that this may be true for observational evidence but surely experiments can settle the question objectively. But this objection is easily rebuttable. If the experiment is a randomised trial, the chance of confounding goes to zero only as the sample size approaches infinity. For finite sample, there is always a possibility of confounding (which, of course, may be ruled out, but only on the basis of what is known at the time). 'Totally controlled experiments' (the term is Nancy Cartwright's, see Cartwright, 1989) are rare in the social and biomedical sciences but where they happen, control can never be 'total'. Rather, when one is lucky, every factor that is known to potentially influence the result can be controlled, but in that way the result is again relative to what is known by the community conducting the controlled experiment.

Objectivity therefore cannot be located in the relation between inferential base and causal claim. How about the relation between claim and inferential target? Let us recall that the inferential target contains sentences that constitute the 'cash value' of a causal claim, that is, sentences describing the kinds of relations (or states of affairs) the knowledge of which is the ultimate purpose of establishing a causal claim, for instance explanatory sentences and sentences concerning predictions and policy interventions. Can such sentences be known objectively?

Explanation is a notoriously interest-relative affair. Too varied are the conditions under which an explanation is regarded as successful in order to be able to argue that explanatory success is objective. An explanation can, of course, be objectively successful given the standards prevailing in a community at the point the explanation is made but that just shows that explanatory success is relative to the standards of the community.

This leaves predictions and sentences expressing possible policy or therapeutic interventions. I will once more leave the development of a positive account of how success at predictions

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

and interventions may lend causal claims objectivity for another occasion. Let me just make a few general remarks here. The way I see it causal claims are objective to the extent that the predictions and interventions made on their basis are successful, and to the extent that success at predicting and intervening can be established objectively. Both these factors are domain specific. Many economic methodologists baulk at the idea of regarding predictive success as a theory confirmation because 'prediction is impossible' in economics (*e.g.*, McCloskey, 1998). In previous work I have pointed out that all highly generic arguments concerning predictive success in economics and other social sciences such as those given by McCloskey are unconvincing (Reiss, 2007) but it is certainly much harder to anticipate events successfully in these sciences than in others. Accordingly, the objectivity of causal claims comes in degrees, and some sciences may establish causal claims of higher, others of lesser objectivity.

The last point to make is that the kind of objectivity lent to causal claims by the success of predictions and interventions that are inferentially connected with them is fairly limited. In particular, it is not the 'corresponding to the objects' sense of objectivity. As is well known, predictions and interventions can be successful without being based on the true causal structure. The way the account of causation proposed here is set up, many of the false causal structures will have been ruled out in the process of establishing the causal claim from the evidence. But given that a community can only rule out *known* possible alternative explanations of the evidence there remains always the possibility that the true explanation is overlooked, the inference made and the prediction or intervention accidentally successful.

4. Conclusion

In this paper I have given an account of the semantics of causal claims, which locates the meaning of causal claims in their inferential connections to claims about evidence for causal claims and cognitively or practically useful claims about explanation, prediction and possible interventions. I have also given an account of what it means for two claims to be inferentially connected in terms of a scientific community's entitlement to make an inference. The cash value of the account is that it allows to address a range of vexing questions concerning causation in original and insightful ways. In this paper I looked at issues in three areas: causation and evidence, the concept of 'cause' and the objectivity of causation and sketched how to employ the inferentialist account to address these. Though I left more specific details of the answers to future work, the inferentialist account has already demonstrated its worth by: a) showing how inferentialism fares better than any of the standard accounts of causation in explaining the relation between causal claims and evidence in their favour; b) showing ways to determine how many concepts of 'cause' there are; and c) showing wherein the objective content of causal claim lies.

The main strength of the proposed account is, in my view, that it makes all these issues partly empirical, *a posteriori* issues. The most significant is the following. To most philosophers the linguistic function of 'cause' is akin to that of logical term such as the material conditional or biconditional: the appearance of 'cause' in a sentence is associated with certain inferences independently of what is being related causally. According to the proposed account, all inferences are fundamentally material and therefore depend on the other terms in the statement. However, formal inferences are the limiting case of material inferences. Thus, should it turn out to be the case that scientific communities are entitled to certain kinds of inferences on account of the appearance of the term 'cause' alone, causal inferences would be formal (or, more accurately, formalised material inferences). For instance, a view of causation similar to Woodward's in which language users are entitled to infer that if one were to intervene on X, Y would change whenever X causes Y (and independently of what X and Y are) is a possible consequence of the proposed account but not built into it.

The same is true of the relationship between causation and evidence and the objectivity of causation. It is a possible consequence of the account proposed here that both mechanistic and probabilistic evidence must be in the inferential base, not a built-in assumption. And once more, causation is fundamentally subjective because of its tight connection to human reasoning practices but an objectification of causal relations is a possible consequence rather than an assumption of the account.

Another question that I unfortunately had to leave open is the relationship between this and other, related accounts of causation. I know of at least three candidate accounts that bear more than superficial relationship to the theory presented here. There is, first, Helen Beebee's 'projectivist' reading of Hume's account of causation (Beebee, 2007a, 2007b). Second, there is Wolfgang Spohn's account of causation which is based on his theory of ranking functions (Spohn, 2006, 2012). Third, there is Jon Williamson's and Federica Russo's epistemic theory of causation (*e.g.*, Russo & Williamson, 2007, 2011; Williamson, 2005, 2006a, 2006b). This matter too will have to await future work.

Acknowledgements

I would like to thank Nancy Cartwright, José Díez, Roman Frigg, Javier González de Prado, Carl Hoefer, Max Kistler, Caterina Marchionni, Phyllis McKay Illari, Samir Okasha, Federica Russo, David Teira, Marcel Weber, the EIPE Reading Group 'Section C' and audiences at the Institut d'Histoire et de Philosophie des Sciences et des Techniques Paris, Erasmus University Rotterdam, UNED (Madrid), Kent University (Brussels Campus) and Geneva University for valuable comments.

References

Ambrose, A. (Ed.). (1979). *Wittgenstein's lectures, Cambridge 1932–1935. From te notes of Alice Ambrose and Margaret Macdonald*. Amherst (NY): Prometheus Books.

Beauchamp, T. (1974). Hume's two theories of causation. *Archiv für Geschichte der Philosophie, 55*(3), 281–300.

Beauchamp, T., & Rosenberg, A. (1981). *Hume and the problem of causation*. Oxford: Oxford University Press.

Beebee, H. (2007a). Hume on causation: The projectivist interpretation. In H. Price & R. Corry (Eds.), *Causation, physics and the constitution of reality: Russell's republic revisited* (pp. 224–249). Oxford: Oxford University Press .

Beebee, H. (2007b). Hume on causation: The projectivist interpretation. In H. Price & R. Corry (Eds.), *Causation, physics and the constitution of reality* (pp. 224–249). Oxford: Oxford University Press .

Beebee, H., Hitchcock, C., et al. (2009). *The Oxford handbook of causation*. Oxford: Oxford University Press.

Brandom, R. (1994). *Making it explicit: reasoning, representing, and discursive commitment*. Cambridge (MA): Harvard University Press.

Brandom, R. (2000). *Articulating reasons*. Cambridge (MA): Harvard University Press.

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

Campaner, R., & Galavotti, M. C. (2007). Plurality in causality. In P. Machamer & G. Wolters (Eds.), *Thinking about causes: from Greek philosophy to modern physics* (pp. 178–199). Pittsburgh (PA): University of Pittsburgh Press.

Cartwright, N. (1989). *Nature's capacities and their measurement*. Oxford: Clarendon.

Cartwright, N. (2004). Causation: One word, many things. *Philosophy of Science, 71*(PSA 2002), 805–819.

Cartwright, N. (2007). Are RCTs the gold standard? *BioSocieties, 2*(2), 11–20.

Codell-Carter, K. (2003). *The rise of causal concepts of disease: Case histories*. Aldershot: Ashgate.

Gilbert, C. (2010). How to understand high food prices. *Journal of Agricultural Economics, 61*(2), 398–425.

Gillies, D. (2011). The Russo-Williamson thesis and the question of whether smoking causes heart disease. In P. McKay Illari, F. Russo, & J. Williamson (Eds.), *Causality in the sciences* (pp. 110–125). Oxford: Oxford University Press.

Godfrey-Smith, P. (2009). Causal pluralism. In H. Beebee, P. Menzies, & C. Hitchcock (Eds.), *Oxford handbook of causation* (pp. 326–337). Oxford: Oxford University Press.

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005

Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. Paul (Eds.), *Causation and counterfactuals* (pp. 225–276). Cambridge (MA): MIT Press.
Hitchcock, C. (1995). The Mishap and Reichenbach fall: Singular vs. generic causation. *Philosophical Studies, 78*(3), 257–291.
Hitchcock, C. (2007). How to be a causal pluralist. In P. Machamer & G. Wolters (Eds.), *Thinking about causes: From Greek philosophy to modern physics* (pp. 200–221). Pittsburgh (PA): University of Pittsburgh Press.
Illari, P. (2011). Mechanistic evidence. disambiguating the Russo-Williamson thesis. *International Studies in the Philosophy of Science, 25*(2), 139–157.
Kincaid, H. (2004). Contextualism, explanation and the social sciences. *Philosophical Explorations, 7*(3), 201–218.
Lewis, D. (1973). Causation. *Journal of Philosophy, 70*(8), 556–567.
Lie, R., & Miller, F. (2011). What counts as reliable evidence for public health policy: The case of circumcision for preventing HIV infection. *BMC Medical Research Methodology, 11*(34), 1–7.
Longworth, F. (2006). Causation, pluralism and responsibility. *Philosophica, 77*(1), 45–68.
McCloskey, D. (1998). *The rhetoric of economics*. Madison (WN): University of Wisconsin Press.
Medina, J. (2001). Verification and inferentialism in Wittgenstein's philosophy. *Philosophical Investigations, 24*(4), 304–313.
Norton, J. (2003). A material theory of induction. *Philosophy of Science, 70*(4), 647–670.
Pearson, K. (1911). *The grammar of science*. London: Walter Scott.
Price, H., & Corry, R. (2007). *Causation, physics, and the constitution of reality: Russell's republic revisited*. Oxford: Oxford University Press.
Reiss, J. (2007). Do we need mechanisms in the social sciences? *Philosophy of the Social Sciences, 37*(2), 163–184.
Reiss, J. (2008). *Error in economics: Towards a more evidence-based methodology*. London: Routledge.
Reiss, J. (2009). Causation in the social sciences: Evidence, inference, purpose. *Philosophy of the Social Sciences, 39*(1), 20–40.
Reiss, J. (2011). Third time's a charm: Wittgensteinian pluralisms and causation. In P. McKay Illari, F. Russo, & J. Williamson (Eds.), *Causality in the sciences* (pp. 909–927). Oxford: Oxford University Press.
Russell, B. (1913). On the notion of cause. *Proceedings of the Aristotelian Society, 13*, 1–26.
Russo, F., & Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science, 21*(2), 157–170.
Russo, F., & Williamson, J. (2011). Generic versus single-case causality: The case of autopsy. *European Journal of Philosophy of Science, 1*(1), 47–69.
Sellars, W. (1953). Inference and meaning. *Mind, 62*(247), 313–338.
Spohn, W. (1993). *Causal laws are objectifications of inductive schemes. Philosophy of probability*. Dordrecht: J.-P. Dubucs, Kluwer. 223–252.
Spohn, W. (2006). Causation: An alternative. *British Journal for the Philosophy of Science, 57*, 93–119.
Spohn, W. (2012). *The laws of belief: Ranking theory & its philosophical applications*. Oxford: Oxford University Press.
Vandenbroucke, J. (2004). When are observational studies as credible as randomised trials? *Lancet, 363*, 1728–1731.
Williamson, J. (2005). *Bayesian nets and causality: Philosophical and computational foundations*. Oxford: OUP.
Williamson, J. (2006a). Causal pluralism versus epistemic causality. *Philosophica, 77*(1), 69–96.
Williamson, J. (2006b). Dispositional versus epistemic causality. *Minds and Machines, 16*, 259–276.
Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.
Worrall, J. (2002). What evidence in evidence-based medicine. *Philosophy of Science, 69*, S316–330.

Please cite this article in press as: Reiss, J. Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.05.005