Causation: Many Words, One Thing?

Lorenzo CASINI

ABSTRACT: How many notions of cause are there? The causality literature is witnessing a flourishing of pluralist positions. Here I focus on a recent debate on whether interpreting causality in terms of inferential relations commits one to *semantic* pluralism (Reiss 2011) or not (Williamson 2006). I argue that inferentialism is compatible with a 'weak' form of monism, where causality is envisaged as *one, vague* cluster concept. I offer two arguments for this, one for vagueness, one for uniqueness. Finally, I qualify in what sense the resulting form of monism is 'weak'.

Keywords: causality; pluralism; monism; inferentialism.

RESUMEN: ¿Cuántas nociones de *causa* hay? Hay un resurgir de posiciones pluralistas en la bibliografía sobre causalidad. Me ocuparé aquí de un debate reciente sobre si la interpretación de la causalidad en términos de relaciones inferenciales nos compromete con el pluralismo semántico (Reiss 2011) o no (Williamson 2006). Defiendo que el inferencialismo es compatible con una forma 'debil' de monismo, donde la causalidad se presenta como un único concepto vago y de agrupación (*cluster*). Ofrezco para ello dos argumentos, uno sobre vaguedad y otro sobre unicidad. Finalmente, matizaré en qué sentido esta forma de monismo es 'débil'.

Palabras clave: causalidad; pluralismo; monismo; inferencialismo

As the reader may have noticed, the title of this paper is reminiscent of another, well-known paper in the causality literature, namely Nancy Cartwright (2004)'s 'Causation: One Word, Many Things'. There she defended a metaphysical thesis, inspired by Anscombe (1971), which is a leitmotif in her writings on causality: "causal" is a label which is attached to a plurality of essentially different relations. What is perhaps less known is that in the same paper as well as later work Cartwright has also advanced – although not fully articulated – another, interesting thesis on the *meaning* of causality, namely that causality is a sort of 'inference license'.

Can one, in the face of pluralism, explain *why* the *same* label, "causal", is attached to all these relations? And need an inferentialist approach to the meaning of causality commit one to semantic pluralism (Reiss 2011) or not (Williamson 2006)?

This paper aims to answer exactly these questions. If we believe that the fact that seemingly different relations are all labelled "causal" needs a philosophical explanation, I claim that inferentialism *can* provide the desired explanation, in terms of the inferences that causal claims license.

In particular, I will be concerned with clarifying how, if at all, inferentialism

can be used to defend a sort of 'weak' semantic monism, where causality is envisaged as *one, vague* cluster concept. I offer two arguments for this, one for vagueness, one for uniqueness. Finally, I qualify in what sense the resulting form of monism is 'weak'.

1.   The monist's challenge

The term "causal pluralism" has been recently associated with a variety of positions, all sharing the idea that there are distinct kinds of causal relations, and no single feature that makes all of them causal. From which it follows that no monistic account of causality as, e.g., regular association, probability raising, counterfactual dependence, mechanism, etc. is possible.[1] But underlying this one, shared idea is a plurality of positions.

     Here I focus on two specific kinds of pluralism, namely *evidential* pluralism (Williamson 2006, Russo and Williamson 2007) and *semantic* (or conceptual) pluralism (Reiss 2009, 2011). An evidential pluralist is pluralist in the minimal sense that he acknowledges that evidence of more than one kind contributes to establish a causal claim. A semantic pluralist, instead, claims that there are various (perhaps 'fuzzy') notions of cause. Whilst evidential pluralism is in principle compatible with conceptual monism, semantic pluralism is clearly not.

     Against all pluralisms – except the evidential one – the monist (Williamson 2006, Russo and Williamson 2007) maintains that, although causal claims may be supported by distinct evidential criteria, there is just *one* notion of cause. This position is in line with the epistemic view of causality, according to which a causal relation is the inference relation drawn by an ideal, fully rational and informed agent (cf. Williamson 2005, 2006).[2] Since the epistemic view purports to say what causality 'essentially' is, it is a metaphysical view. However, insofar as it defines the meaning of causality in terms of inferences, it also counts as a position on the semantics of causal claims. Only the latter aspect, and not the details of the epistemic view, is relevant to the purpose of this paper. Here, I will only be concerned with the epistemic monist's objection against pluralism, in particular with its implications for Reiss' *semantic* pluralism. Accordingly, my argument will only concern the semantics, *not* the metaphysics of causality.[3]

---

[1] Here I am not concerned with *justifying* pluralism. Rather, I will *assume* that causality is diverse – no analysis captures its meaning exhaustively – and consider *how pluralist* one should be.

[2] To clarify: *contra* 'traditional' monistic analyses, the epistemic view is – minimally – pluralist as it doesn't erect any *test* condition to the status of *truth* condition; at the same time, it also maintains that there is something *all* causal relations have in common, namely their being (ideal) inference relations.

[3] This is not to say that inferentialism is incompatible with giving causal talk a referential value, or being 'realist' about causality. In this regard, it would be worth considering whether a case can be made for inferentialism providing something more substantial

Williamson's objection to pluralism is that it cannot explain why the same label "causal" is used to denote a variety of seemingly different – and differently characterised – relations. If one believes this is a philosophically interesting question, then one will demand that a theory of causality provide an answer to it. Notice that some deny that the question deserves a philosophical answer (see Psillos 2009, Reiss 2009). In §4., I endeavour to show that an interesting, philosophical story *can* be told, viz. an inferentialist story.

If one finds pluralism dissatisfying, one may want to resort to one or the other kind of monism. One, as I said, is epistemic monism. Another, as I argue in §4., is the (Sellersian-Brandomian) version of causal inferentialism I sketch in §2.

That inferentialism *may* have monistic implications is suggested by Cartwright herself who, after pointing to the possibility to interpret causal claims as inference licenses (cf. Cartwright 2007, p. 46), states: "What we should be looking for is a theory of causality, in much the same way as we have a theory of the electron" (Cartwright 2007, p. 52). It is now commonplace that there are no necessary and sufficient conditions that *define* theoretical terms; still, we have theories that give *one* story about them. Cartwright seems to think – here, at least – that the same reasoning applies to causality.

To the inferentialist camp belong, besides Cartwright, also Godfrey-Smith (2009) and Reiss (2011). Yet, they define their inferentialism differently – so that only Godfrey-Smith but not Reiss can be associated with the sort of weak, conceptual monism I am arguing for in this paper (see §4.2.).

Interestingly, what the advocates of both epistemic monism (Williamson 2005, 2006) and inferentialism (Reiss 2011) have in common is that they tend to characterise causality in terms of *inferential* relations (see §3.). Reiss, in particular, explicitly draws his pluralist conclusion from an inferentialist approach to the semantics of causal claims. So the question arises as to whether endorsing inferentialism need to commit one to semantic pluralism (Reiss) or not (Williamson). Before tackling this question in §4., I will briefly sketch the contours of an inferentialist account of causality.

2. *Causality* as *inference*

Inferentialism is an approach to semantics (Brandom 2000, 2007, Peregrin 2006, 2012), rooted in the pragmatist tradition (Sellars 1953, Wittgenstein 1956), and only recently advocated in philosophy of science (Brigandt 2010, Reiss 2011). In short, inferentialism purports to account for the meaning of linguistic expressions in terms of "what can both serve as and stand in need of

---

than a story on the 'expressivist' role of our conceptual activity (cf. Brandom). However, this would require a *separate* argument – an argument that I cannot give here. I leave this to future work. For the purpose of this paper, we may be anything from eliminativist (à la Psillos) to Anscombian pluralist (à la Cartwright) as regards the metaphysics of causality. This has no implications for the general point I wish to make.
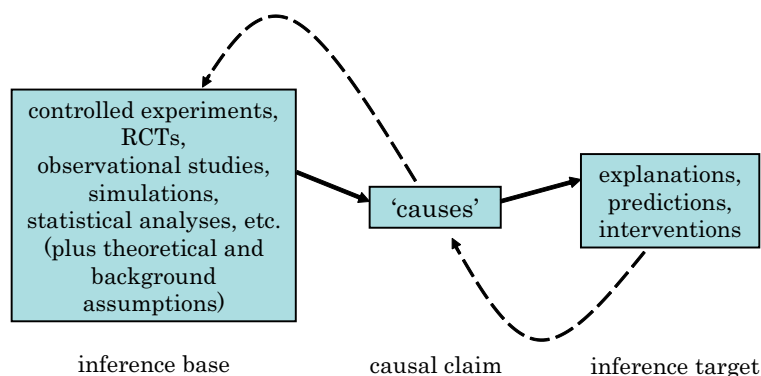
Figure 1: Causality as *inference*

reasons" (Brandom 2007, p. 654), that is, in terms of *inferential* relations between circumstances of appropriate application (*premisses* of inferences) and appropriate consequences of application (*conclusions* of inferences).

Expressions derive their meaning from the rules of inference they obey (see Brandom 2007, §1-§3). In particular, subsentential locutions derive their meaning not from their referential function but from the sentences in which they occur; and sentences, in turn, derive their meaning not from truth-valuations or truth-conditions but from their inferential role (cf. Sellars 1962).

As regards the meaning of a given *sentence* (e.g., a specific causal claim such as 'X causes Y'), its meaning-constitutive inferences can be divided into those that, respectively, *warrant* the sentence and *are warranted* by it. Following Reiss (2011), I call the sentences that warrant the claim 'inferential base' and those warranted by it 'inferential target'.

As regards the content of a *concept*, this is defined by Brandom (2007) in terms of inferences and incompatibilities.[4] We need not be concerned with the details of Brandom's semantics, only bear in mind that the inferentialist has a machinery to get the meaning of words (e.g.: 'causes'; 'causality') out of inferences.

Now, if one endorses inferentialism about meaning, what should one say about the meaning of causality? Arguably, the meaning of causal claims is constituted by and analysable in terms of inferential relations between evidence of experiments, results of RCTs, observational studies, simulations, statistical analyses, etc. (together with theoretical and background assumptions) on the one hand, and possible explanations, predictions and interventions on the other (see figure 1). But is this sufficient to say what 'causes' means?

---

[4] In particular, the class of inferences must comprise not just the logically correct ones but also the *materially* correct ones, and not only language-to-language inferences, but also inferences involving non-inferential circumstances of appropriate application (observations) and non-inferential appropriate consequences of application (actions) (Brandom 2007).

Problems may arise here (see §3.), having to do with identifying the meaning-constitutive inferences of, respectively, a *particular* causal claim and 'causality' *simpliciter* (i.e., 'causes' as occurring in *all* causal claims). At the root of these problems is a traditional trouble for inferentialism, viz. semantic holism: If the meaning of linguistic expressions is inference, and inferences are all related to one another, then the meaning of *any one* linguistic expression depends on *all* inferences (see Quine 1951). How can one isolate the meaning-constitutive inferences from the non-meaning-constitutive ones?

In the case of causality, the problem may be particularly acute. Intuitively, holism seems easier to address in the case of observational statements and expressions[5], which are 'low-level', viz. closer to the observational contexts in terms of which their meaning can be made explicit. But 'causality' is (often) 'high-level', and related to observational contexts only 'indirectly', through many other inferences. All the more reason to believe that 'causes', whether in a specific claim or *simpliciter*, is so entrenched in our conceptual apparatus that its meaning-constitutive inferences are too 'distributed' to be isolated.

Now, *whether* the inferentialist approach is *in general* successful depends on whether a *general* response to holism is successful.[6] This is an open issue which I need not address here. Rather, I will assume that one such response *is* available and, as a result, semantic holism *may* be contained, at least for a class of concepts, or for some concepts better than others. Even on this assumption, we still need an *additional* argument to the point that causality is among those concepts for which meaning *can* be specified. Providing such an argument is apparently harder because of the centrality of the notion of causality in our conceptual apparatus. Still, I want to argue that there is room for a weak form of semantic monism.

I will now present in some more detail evidential and semantic pluralism and their connections with inferentialism, before discussing in §4. whether and to what extent they follow from the above-sketched inferentialist semantics.

3.   *Evidential vs semantic pluralism*

Should we endorse only *evidential* pluralism Williamson (2006), Russo and Williamson (2007) or also *semantic* pluralism (Reiss 2009, 2011)? For the evidential pluralist, traditional monistic accounts offer at most evidence for causal relations, i.e., test conditions, *not* exhaustive analyses, i.e., truth conditions. Provided one rejects such accounts, this claim is quite straightforward and uncontentious. More contentious is whether these different analyses identify distinct notions of cause. The debate has an obvious bearing on whether the attempt to provide a conceptually monistic account of causality is well-founded.

---

[5] Yet, it seems possible to define inferentially *logical* vocabulary (see Peregrin 2006).

[6] To deal with holism, the inferentialist might resort to, e.g., Sellars (1948) or Brandom (1994) (cf. Brandom 2007, §4).

On the one hand, in line with his epistemic view, Williamson claims that different kinds of evidence for causal claims do not presuppose different concepts of causality. The evidence gathered by using the various criteria, which 'typically' apply together, is evaluated depending on how well it conduces to successful uses – viz. successful predictions, interventions and explanations – so as to *univocally* constrain rational belief in what causes what.

On the other hand, Reiss suggests the following identity condition for the meaning of 'causes':

> 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 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 (Reiss 2011, pp. 923-924).

This condition relativises the meaning of a token causal claim to both *kind* and *number* of criteria the claim is inferentially related to – where 'criteria' stands here for both *test* conditions (base criteria) and *use* conditions, or purposes (target criteria). One thing, in fact, is to say what 'causes' *simpliciter* means, quite another thing to say what 'X causes Y' means when applied to population P_1 as opposed to population P_2, since, e.g., mechanisms or probabilities by which the effect is brought about as well as criteria used to establish the claim may be different in P_1 and P_2.

For this reason, the meaning of any causal claim – e.g., in epidemiology, 'exposure to aflatoxin causes liver cancer' – is population-relative:

> (...) 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 (Reiss 2011, p. 917).

To this consideration, one may add that identity depends not only on sameness of *kinds* of sentences, or (formal) criteria, but also on sameness of *contents* of sentences, or 'propositions'. This further condition seems implied by what Reiss himself elsewhere states: "[t]here is a definite set of propositions with which any causal claim is inferentially related" (Reiss 2011, p. 924).[7] Notice that this applies even when the same relata, hence the same causal claim type (as in, e.g., 'X α-causes Y' and 'X β-causes Y'), are concerned. That is, meaning depends not just on base and target criteria but also on the specific content of base and target sentences. Indeed, the content of the sentences need not coincide with the content of the criteria. For instance, one thing is to say how a randomised control trial (RCT) is inferentially related to a causal claim, i.e.,

---

[7] Admittedly, the *content* of base and target sentences may not count, for Reiss, as part of the identity conditions for 'causes'. Indeed, that he intends to impose this further condition is denied by him (private communication). Still, *pace* Reiss, his argument for the unspecificity of 'causes' *presupposes* this condition (see §4.1.).

how it is to be conducted to provide support for the claim. Another thing is to specify how the meaning of 'exposure to aflatoxin causes liver cancer' – whether applied to *different* populations (e.g., mice and humans) or to the very *same* population – depends on the particular way the RCT is conducted, i.e., the specific individuals on which the RCT is performed (each, arguably, instantiating slightly different mechanisms), the result of the randomisation in a specific trial, etc. So, even on the assumption that the criteria are the same, meanings may end up different depending on the different circumstances in which tests are performed and on the different actions taken – described, respectively, in base and target.

Now, since meaning of both criteria and contents is to be defined inferentially, it is easy for pairs of inferential webs to differ somewhere. No surprise that Reiss takes inferentialism to entail conceptual pluralism: the concept of cause is "*unspecific* rather than *ambiguous*" (Reiss 2011, p. 914), and should be replaced by a plurality of well-defined – or well-defin*able* – concepts.

Who is right? In spite of the disagreement, both Williamson and Reiss maintain that causality has to do with inference. Shouldn't we expect from inferentialism itself to deliver an answer on the monism-vs-pluralism debate? Well, things are not so easy. Two readings, in fact, are available.

According to a first reading, clearly adopted by Reiss, 'causes' is ambiguous and differs from claim to claim, depending on the particular base and target. That is, tokens of 'causes' used in claims established by different methods and licensing different inferences have different meanings.

A second reading, instead, has it that the premises from which causal claims are entailed, and the conclusions that causal claims entail, only differ as to the *weight* of the different kinds of inferences which, respectively, warrant the claims and the claims warrant. Causality is only 'moderately' diverse: criteria employed in different circumstances, although weighted differently, are similar enough to legitimate a strong clustering among the various tokens of 'causes'. This second reading, for which causality is *one* 'cluster concept', is closer to Williamson's position.

The issue is all the more relevant insofar as, according to one's favourite reading of the inferentialist take on meaning, evidential and semantic pluralism give two different answers to the question of what, if at all, is the philosophically interesting fact that explains the use of "causal" to denote all kinds of causal relations. As mentioned, for Reiss there is no fact of the matter: whether different criteria 'typically' coincide is an empirical not conceptual matter, "much like discovering that various symptoms of a disease typically co-occur" (Reiss 2009, p. 33) (cf. Psillos 2009); "why we have come to call the different kinds of relationships causal is a matter of historical, not philosophical, inquiry" (Reiss 2009, p. 37). For Williamson, instead, there is *one* fact of the matter: different criteria typically coincide because they all provide evidence for the *same* (kind of) relation; we have come to call all these relations 'causal' because they *all* share the role of licensing inferences. Which answer is correct?

The issue largely depends on whether 'causes' is to be regarded as *vague*, as

the monist can maintain, or *unspecific*, as Reiss claims. If 'causes' is vague, there can be *one* cluster. In this case, vagueness is *semantic*.[8] If 'causes' is unspecific, instead, there are irreducibly *many* notions. On this view, the 'vagueness' is *epistemic*, i.e., it concerns ignorance on which of many distinct notions is employed in each case: "True, we might not always have a very clear idea of what [the meaning-constitutive sets of propositions] are. But this is a question of epistemology, not of semantics" (Reiss 2011, p. 924).

One way to tackle the issue is to look at the foundations of the inferentialist project itself. As mentioned, granting that holism can be contained, we still need an argument to show that 'causality' is liable to be analysed inferentially.

More specifically, holism generates two further problems: (i) the *(in)stability of conceptual contents* under change of belief and commitment to the properties of various inferences; and (ii) the *(im)possibility of communication* between individuals who endorse different claims and inferences. But, then, how can a change in one belief not result in a change in all other beliefs? And how can two speakers' whose beliefs differ somewhere ever talk about the same thing?

When it comes to causality, the inferentialist needs an argument to guarantee, in the face of the holistic nature of meaning, (i) the relative stability of conceptual content of 'causes' and (ii) the possibility of successful communication on causal claims. In §4., I show that, if these two goals can be achieved, holism can be contained in a way that makes room for 'causes' being *one* and *vague* rather than unspecific. I offer two arguments to this point, – one for the vagueness of the notion of cause – the 'argument from (in-)stability of content' – another for its uniqueness – the 'argument from communication'.

4.   *Causality* as *one, vague concept*

4.1.   *Argument from (in-)stability of content*

To begin with, notice that Reiss' argument for semantic pluralism trades on an ambiguity between (i) analysing meaning in terms of different *sets* of (token) sentences and (ii) analysing it in terms of sets of different *kinds* (or types) of sentences within those sets.

On the first interpretation, strong semantic pluralism follows straightforwardly. To each inference base and target of (token) sentences there correspond a different meaning. However, this would make it strictly speaking impossible for linguistic expressions to share the *same kind* of meaning – unless some similarity criterion is allowed, that is, unless one opts for the second interpretation. Intuitively, it must be possible and legitimate to group linguistic expressions depending on their similar inferential roles. Analogy of inferential roles could then be used to make meaning of kinds of linguistic expressions explicit. Otherwise one would be stuck with (meaning of) single-case claims and unable to

---

[8] My argument does not rely on a philosophical interpretation of vagueness. An intuitive notion of 'vague' as 'non-fully-specifiable' suffices here.

formulate general claims.

The second interpretation is surely more attractive. Here, strong semantic pluralism follows only on the assumption of a 'double standard' for the semantics of 'causes' and the ancillary notions (e.g., 'depends', 'produces') that help make 'causes' explicit. That is, if one wants to distinguish between *kinds* of sentences – which is plausible – and still draw the same pluralist conclusion on the meaning of causal claims, one must have a similarity criterion that legitimates the clustering of sentences in base and target of the ancillary notions but prohibits the clustering of sentences in base and target of 'causes'. For instance, the criterion will dictate that the circumstances that entail and are entailed by 'probabilistically depends' are similar enough to make one cluster, whereas those that entail and are entailed by 'causes' are not. As a result, 'probabilistically depends' counts as a cluster concept, whereas 'causes' is an unspecific term that subsumes a plurality of specific concepts. But this is problematic, as I am now going to show.

For the inferentialist, language is essentially *dynamic*. This means that

> any codification or theoretical systematization of the uses of (...) vocabulary-kinds by associating with them meanings that determine which uses are correct will, if at all successful, be successful only contingently, locally, and temporarily (Brandom 2008, p. 5).

Drawing on this insight, the thesis of the argument from (in-)stability of content (S) is that the respect in which different meanings differ cannot be made fully explicit not only for epistemic but also for semantic reasons:

S1. Communication can be unsuccessful, not only because of the ignorance of the speakers, but also because language is a *dynamic* network of concepts.

S2. If so, then the meaning of 'x-causes' can be specified only *on-the-fly* by successful use.

S3. But, at any time, *also* the meanings of the ancillary notions that should make 'x-causes' semantically explicit is specified on-the-fly.

SC. Hence, the meaning of 'x-causes' can be – semantically – only *partially* specified, or specifi*able*. That is, it is not only semantically unspecific but also *semantically vague*.

To claim otherwise would mean to accept the double standard: meaning of ancillary notions is vague whereas meaning of 'causes' is unspecific. But how could this be justified?

True, the meaning of ancillary notions may be easier to make explicit. In inferentialist terms, their tokens are easier to group as belonging to the *same kind* on the basis of their inferential role. For instance, the meaning of 'probability raising' or 'counterfactual dependence' can be formally fixed in

terms of necessary (target) and sufficient (base) conditions. However, holism applies to such concepts, too: when it comes to applying the formally-defined concept to *non-formal* circumstances, it is still left to us to decide whether the formal notion applies or not – whether 'X raises the probability of Y' or 'Y counterfactually depends on X', etc. Strictly speaking, *no* concept can be *totally* isolated from the other concepts.

And if one applies to the ancillary notions the same strict identity condition which is applied to 'causes', one ought to conclude that distinct tokens of the ancillary notions (e.g., 'α-probabilistically raises' as occurring in 'X α-probabilistically raises Y' as opposed to 'β-probabilistically raises' as occurring in 'Z β-probabilistically raises W') have distinct meanings, since their inferential role is fixed by distinct sets of sentences. Token sentences in base and target are different for *each* claim – whether involving the concept of cause or other concepts.

So Reiss' argument for unspecificity generalises in principle to the ancillary notions: one cannot say there are many distinct, unspecified concepts of causality without at the same time saying that there are many distinct, unspecified concepts of probability raising, counterfactual dependence, etc. From which it follows that if we want the meaning of the ancillary notions to be semantically vague on *similarity* grounds, then we should allow causality to be vague as well, on the same grounds.

Admittedly, one could object to this analysis that the vagueness of 'causes' is of a different kind from that of the ancillary notions, 'causes' being more like a multiply-realisable concept (e.g., 'bird', realised by distinct species of bird), the ancillary notions more like *non*-multiply-realisable concepts (e.g., 'bald'). But even granting that this is so, we don't deny that there is one, legitimate concept of 'bird-ness'.[9] So why deny that there is one concept of 'causality'?

Now, a serious problem may arise, having to do with the possibility of meaningful communication. If the identity of meaning depends on the inferences in terms of which the meaning is analysed, and is only partially fixed/fixable, how can meaning be stable enough to allow us to discuss about the same things? In particular, how can meaning be stable enough so that we can make it explicit in communication, by clarifying the inferential presuppositions and implications of what we say, so as to resolve controversies on what claims to endorse?

Clearly, due to holism, meaning cannot be *totally* transparent to the speakers (see above). Yet, much of what we say, in particular many of the implications of our commitments, *must* be transparent to us – otherwise communication wouldn't be possible. Granted that holism can be contained, there must be a 'semantic level' at which meaning of many notions that we use in communication is stable enough. Arguably, that is the level at which meaning is *vague enough* for the speakers to take it as *one* and treat its possible applications as

[9] To reiterate a point made earlier (see fn. 3), I am not concerned here with metaphysical issues, such as whether or not there is an object-type the concept refers to, whether bird-ness is a 'natural' kind, whether causal relations are 'real', etc.

similar enough. That is the level at which, I claim, there is a *unique* concept of causality. But here we need an argument to the point that we can legitimately talk of *one* vague notion of causality, as opposed to *many* vague notions.

4.2. Argument from communication

Before presenting the argument, it is crucial to point out the source of the difference between Reiss and Williamson in the unity-vs-disunity debate, namely the different weight they ascribe to base and target.

On the one hand, Williamson places more weight on the target. All kinds of evidence are required, in principle, to establish a causal claim (Russo and Williamson 2007). This is because, ultimately, their role is helping maximise the target's success. The meaning-constitutive inferences are (only) claim-to-target inferences (see Williamson 2006, p. 78). And since the target is a unique class of claims, namely explanations, predictions and claims about results of interventions, there is *one* notion of cause. That is, the unity of purpose(s) is what fixes the meaning of 'causes' and blocks the fragmentation. Here, inferentialism is used to explain philosophically the uniqueness of meaning.

On the other hand, Reiss places more weight on the base: "whereas the meaning of an expression is given by its inferential connections (...), its method of verification determines what these inferential connections are" (Reiss 2011, p. 923). And since we have many ways to establish a causal claim, i.e., many appropriate base-to-claim inferences, we have *many* concepts of cause as well (see Reiss 2011, p. 924). Nor does an appeal to purposes (target) help unify. Rather, purposes *dis*unify, by pulling apart the role of causal claims. So, an epidemiologist may be interested in explaining whether the population-level correlation between aflatoxin exposure and liver cancer is due to the carcinogenicity of aflatoxin; a policy maker in knowing whether controlling aflatoxin is an effective way to reduce mortality; someone exposed to aflatoxin in predicting whether this exposure will result in a higher chance of liver cancer. To the three purposes there correspond different criteria to establish the causal claim, hence different meanings of cause (see Reiss 2011, §3-§5) – even within the very same discipline (cf. Reiss 2009, §6). The fact that different relations are all called 'causal' has no philosophical significance. Inferentialism is only invoked to describe the fragmentation, not to explain the way this came about, which is a contingent, historical matter.

Here, it is worth contrasting Reiss' position with Godfrey-Smith (2009)'s. For Godfrey-Smith, a cluster concept is an "amiable jumble" of criteria, such that its criteria sometimes pull apart but 'typically' don't. Causality is peculiar because the jumble is "cantankerous", not amiable (Godfrey-Smith 2009, p. 331). By this, he means that causality is partly an "essentially contested concept" (ECC) and partly low-level, or uncontentious. It is ECC, because

> it is not just hard to work out when the conditions of application are met, but (...) the conditions for application themselves are, given the concept's role, permanently susceptible to being challenged and renegotiated (Godfrey-Smith 2009, p. 335)

However, it has also low-level, uncontentious uses, where disputes with respect to boundaries and criteria for application *can* be resolved. Terms with both ECC and low-level uses acquire their role when

> their successful application has significant downstream consequences, but their domain is complex in ways that involve the absence of sharp boundaries that function as attractors to usage. [In the case of causality], an accepted set of exemplars and a sense of shared *purpose* behind diverse uses prevent a fragmentation into distinct concepts. These ideas might be linked to tools developed in recent 'inferentialist' philosophy of language (...) (Godfrey-Smith 2009, p. 336)

In virtue of such a shared sense of purpose, unity prevails over disunity. This conclusion, partly reached by invoking inferentialism, is in obvious disagreement with Reiss, for whom an appeal to purpose does *not* prevent fragmentation.

How should one decide between evidential and semantic pluralism? To answer this question, partly drawing on Godfrey-Smith's distinction between ECC and low-level uses, I employ the argument from communication (C):

C1. For distinct concepts to exist there have to be different communities using the *same* word with *different* meanings, and yet either (i) *never* successfully communicating by using the word, viz. never agreeing whether implicitly or explicitly on its rules of 'correct' application; or (ii) communicating by using the word but *never* discovering the disagreement on such rules.

C2. The second option is clearly implausible, as evidenced by the explicit, semantic disagreements among both scientists and philosophers.

C3. The first option is implausible, too: true, often different communities don't discuss subject-specific causal claims; but sometimes they do; plus, sometimes they engage in high-level semantic reflections; in both cases, they *can* reach semantic agreement.

CC. Therefore, 'causes' *can* be – in a sense to be qualified – *one* concept.

This argument agrees with Williamson that – in a sense – purposes contribute to unify, whilst conceding to Reiss that – in another sense – they don't. What exactly are the two senses can be illustrated in terms of different *semantic levels*, i.e., different contexts in which purposes play their role in fixing the meaning (see below). Depending on the level, we may have either unity (uncontentious uses) or disunity (contentious uses). In both low- and high-level cases, a shared purpose – together with the taking of our linguistic expressions as *committing* us to certain consequences[10] – tends to generate a shared commitment on meaning. That is, several speakers/communities that given

---

[10] This is usually referred to by the inferentialist as the 'normativity of meaning' (see Peregrin 2012).

some accepted base agree on endorsing a claim *ought to* agree on the claims that follow from it and the claims that are incompatible with it – that is, they ought not deny (respectively, endorse) the claims that follow from it (are incompatible with it), once they are made aware of them.

First, I agree with Godfrey-Smith that unity is possible at the subject-specific, low level of *token*, 'C causes E' claims. This may not be the case when, say, distinct claims (e.g., 'X causes Y' and 'Z causes W') are concerned. Here, different purposes can easily produce low-level variability. However, take the *same* claim, e.g., 'exposure to aflatoxin causes liver cancer', and *distinct* communities holding some commitment (not necessarily the same) towards the claim, e.g., community A endorsing the claim in the light of an RCT in mice, community B not endorsing it, in the light of an observational study in humans. Here, disunity is possible only provided findings of one community about the claim, obtained by using one criterion and ensuing in certain consequences, once communicated to the other community do *not at all* affect their commitment to the claim. But this sounds implausible. For instance, it is implausible that A's evidence for the causal connection between aflatoxin and liver cancer in mice has *no* bearing on B's belief that aflatoxin causes liver cancer in humans. Reasoning from analogy, especially on evolutionary grounds, is common in science. And the possibility of an RCT in mice and not in humans doesn't seem to be reason enough for the meaning of 'causes' in the two claims to be of different *kinds*.

On the contrary, whenever scientists coming from different backgrounds interact, for instance in interdisciplinary projects, they must agree on the interpretation of their results as well as accept the bearing of each others' methods on such results. The outcome of their research is very often the formulation of causal claims. When urged by philosophers to say what they mean by 'C causes E', scientists may well disagree – this largely depends on how their training shapes their methods and purposes. However, they also need to come up with a coherent story. After all, the causal claim is the result of a collaborative effort, shared methodology, assumptions, results' interpretation, etc. There must be something they all mean by 'C causes E', at least in that context.

One instance of such interdisciplinary projects is EnviroGenoMarkers.[11] This project is driven by the success of past epidemiological studies in identifying the risk of environmental exposures (e.g., air pollution and passive smoking) with respect to the onset of chronic diseases (e.g., cancer and coronary artery disease). The current project, whose methodology comprises both experiments on biological samples and observational studies, aims at "the identification of both biomarkers of exposure (e.g., dietary components, environmental pollutants) and of markers of early damage (e.g., early disease-specific metabolic changes), notably for carcinogenesis" (Chadeau-Hyam et al. 2011, p. 84). The team includes chemists, molecular biologists, epidemiologists, statisticians, etc. Due to such a diverse composition, one would expect a clash of intuitions

---

[11] For more information on the project, see `http://www.envirogenomarkers.net/`.

on what counts as causal. In spite of this, there is agreement in interpreting previous studies on environmental exposures as providing evidence for causal claims: "biomarkers (...) have contributed to make the association more plausibly causal" (Vineis et al. 2009, p. 54). And with regard to their own study, the team maintains that "the finding that preclinical biomarkers shown to be related to particular exposures in prospective studies are also elevated in certain subclasses of disease would strengthen causal links between exposures and disease" (Chadeau-Hyam et al. 2011, p. 85). At the same time, the connection between purposes is implicitly acknowledged: "intermediate biomarkers can provide important mechanistic insight into the pathogenesis of environmental diseases" (Vineis et al. 2009, p. 54). That is, a successful prediction provides at least some explanation. This supports my unity-over-disunity claim: to the extent that there is some connection between purposes, successful communication is conducive to low-level unity, not fragmentation.

Secondly – here I part with Godfrey-Smith – unity may be possible also at the non-subject-specific, high level where speakers discuss the meaning of a word, e.g., 'causes' (now treated as a *type*), and agree on the formal criteria that are 'typically' associated with it. In science, such 'abstract' discussions may not take place frequently, but they do take place sometimes. Instead, in areas of discourse such as philosophy they take place on a regular basis. Now, *whenever* they take place beween any two speakers A and B, provided the (counter-)examples used by A, who adopts certain rules for the use of some expression, are regarded by B as informative on the correctness of B's rules of application for that expression, the possibility of some unique concept – that is, a shared core of rules – that A and B are referring to is presupposed.

This largely depends on the existence of a stable class of interrelated purposes. In particular, the class must be *stable enough*, so that discussions on the appropriate tools to achieve the purposes, in the light of past consequences of the tools' application, can take place. And the purposes must be *interrelated enough*, so that the success with regard to one purpose, driven by the application of one tool, affects to some extent the other purposes as well. These two conditions seem both well satisfied in the case of causality. Here the class of purposes, i.e., prediction, intervention and explanation, is traditionally very stable. And, although meaning is in principle flexible, such a class seems – although this is just my best guess on the way our linguistic practice will evolve – very unlikely to change. Also, the purposes are strongly interrelated. So, it is hardly the case that a good explanation (say, a mechanistic explanation of how aflatoxin causes liver cancer) has *no* bearing on successful prediction, when the appropriate circumstances are in place (when the individual instantiates the mechanism); and it is unlikely that a successful prediction (a preclinical biomarker predicting, in the presence of some exposure, the onset of the disease) tells us *nothing* on possible interventions (either directly, if the biomarker is itself a cause, or indirectly, if disease and biomarker are effects of a common cause); etc. If the above reasoning is correct, then making explicit the meaning of this unique, high-level concept is in principle possible. Causality need not

be *essentially* contested. Unity *can* be achieved and, although it is ultimately dependent on variations in tools and purposes, is likely to remain stable.

Conclusion

Why is the same label "causal" applied to many, seemingly-different relations? Answer: their sharing the feature of licensing certain inferences. *What* inferences? This may depend on the context as well as the purpose of the enquirer.

So, in what sense *can* 'causes' have the same meaning – *across* contexts and purposes? In the sense that, as far as our tools and purposes are stable enough (so that we can make them explicit) and related to one another (so that there is something to make explicit), there is, as a matter of fact, one, vague cluster of criteria that helps us best achieve those purposes. 'Causes' may have, after all, *one*, *vague* meaning.

But *should* 'causes' have the same meaning? Yes, it 'should', in a qualified sense: we should strive towards unification – *not* fragmentation – to understand *to what extent* we are using the same concept, e.g., to what extent the endorsement of '*C* causes *E*' in one context should carry over to another context. In *this* sense, inferentialism can support a 'weak' form of conceptual monism, by promoting the identification of shared assumptions and implications behind our causal judgements. The aim is to make explicit the contexts which we implicitly take to warrant the inference to the claim, and the jobs that we implicitly want the claim to do for us, that is, the entitlements that we expect the claim to provide us with and the purposes that we expect the claim to help us satisfy.

Why is such a monism 'weak' then? Because the issue of what 'causes' means cannot be settled once and for all, by either scientists or philosophers. If our concepts are dynamic, only partly constrained by our practices and Nature's inputs and outputs to such practices, we can only try to interpret concepts on-the-fly. Their meaning – 'causes' included – must be called into question, made explicit and renegotiated, in a never-ending 'virtuous' circle. There is no ultimate court and no ultimate judge. In the case of causality, one (I) can only hope that meaning can be made explicit *enough*, core inferences identified, and vagueness contained.

Acknowledgements

REFERENCES

Anscombe, G. E. M. 1971. Causality and Determination. CUP Archive. Repr. in Sosa, E. and Tooley, M., editors, *Causation*, pages 88-104. Oxford: OUP, 1993.

Brandom, R. B. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment.* Cambridge, MA: Harvard University Press.

Brandom, R. B. 2000. *Articulating Reasons: An Introduction to Inferentialism.* Cambridge, MA: Harvard University Press.

Brandom, R. B. 2007. Inferentialism and Some of Its Challenges. *Philosophy and Phenomenological Research* 74: 651–676.

Brandom, R. B. 2008. Towards an Analytic Pragmatism. *Philosophical Topics* 36: 1–27.

Brigandt, I. 2010. Scientific Reasoning Is Material Inference: Combining Confirmation, Discovery, and Explanation. *International Studies in the Philosophy of Science* 24: 31–43.

Cartwright, N. 2004. Causation: One Word, Many Things. *Philosophy of Science* 71: 805–819.

Cartwright, N. 2007. *Hunting Causes and Using Them.* Cambridge: CUP.

Chadeau-Hyam, M., T. J. Athersuch, H. C. Keun, M. De Iorio, T. M. Ebbels, M. Jenab, C. Sacerdote, S. J. Bruce, E. Holmes, and P. Vineis. 2011. Meeting-in-the-middle Using Metabolic Profiling. A Strategy for the Identification of Intermediate Biomarkers in Cohort Studies. *Biomarkers* 16: 83–88.

Godfrey-Smith, P. 2009. Causal Pluralism. In *The Oxford Handbook of Causation*, edited by H. Beebee, P. Menzies, and C. Hitchcock, 326–337. Oxford: OUP.

Peregrin, J. 2006. Meaning as Inferential Role. *Erkenntnis* 64: 1–35.

Peregrin, J. 2012. Inferentialism and the Normativity of Meaning. *Philosophia* 40: 75–97.

Psillos, S. 2009. Causal Pluralism. In *Worldviews, Science and Us: Studies of Analytical Metaphysics. A Selection of Topics from a Methodological Perspective*, edited by R. Vanderbeeken and B. D'Hooghe. World Scientific Publishers.

Quine, W. V. O. 1951. Two Dogmas of Empiricism. *Philosophical Review* 60: 20–43.

Reiss, J. 2009. Causation in the Social Sciences. Evidence, Inference, and Purpose. *Philosophy of the Social Sciences* 39: 20–40.

Reiss, J. 2011. Third Time's a Charm: Causation, Science and Wittgensteinian Pluralism. In *Causality in the Sciences*, edited by P. Illari, F. Russo, and J. Williamson, 907–927. Oxford: OUP.

Russo, F., and J. Williamson. 2007. Interpreting Causality in the Health Sciences. *International Studies in the Philosophy of Science* 21: 157–170.

Sellars, W. 1948. Concepts as Involving Laws, and Inconceivable Without Them. *Philosophy of Science* 15: 287–315.

Sellars, W. 1953. Inference and Meaning. *Mind* 62: 313–338.

Sellars, W. 1962. Truth and Correspondence. *The Journal of Philosophy* 59: 29–56.

Vineis, P., A. Khan, J. Vlaanderen, and R. Vermeulen. 2009. The Impact of New Research Technologies on Our Understanding of Environmental Causes of Disease: the Concept of Clinical Vulnerability. *Environmental Health* 8: 54.

Williamson, J. 2005. *Bayesian Nets and Causality: Philosophical and Computational Foundations.* Oxford: OUP.

Williamson, J. 2006. Causal Pluralism versus Epistemic Causality. *Philosophica* 77: 69–96.

Wittgenstein, L. 1956. *Bemerkungen über die grundlagen der mathematik.* English Trans. *Remarks on the Foundations of Mathematics*, 1978, Oxford: Basil Blackwell.

Lorenzo CASINI is a PhD student in Philosophy at the University of Kent, working on theories of causality as applied to complex systems with focus on systems biology and computational economics.

ADDRESS: Department of Philosophy, University of Kent, Canterbury, UK. E-mail: lorenzodotcasini@gmail.com
Website: `http://www.lorenzocasini.com/`