

Causal webs in epidemiology

Federica Russo

Philosophy, Kent

Draft of 30 October 2009

To appear in *Paradigmi* – Special issue on the Philosophy of Medicine.

Abstract. The notion of ‘causal web’ emerged in the epidemiological literature in the early Sixties and had to wait until the Nineties for a thorough critical appraisal. Famously, Nancy Krieger argued that such a notion isn’t helpful unless we specify what kind of spiders create the webs. This means, according to Krieger, (i) that the role of the spiders is to provide an explanation of the yarns of the web and (ii) that the sought spiders have to be biological and social. This paper contributes to the development of the notion of causal web, elaborating on the two following points: (i) to catch the spiders we need multi-fold evidence—specifically, mechanistic and difference-making—and (ii) for the eco-social to be explanatory, the web has to be mechanistic in a sense to be specified.

1. Causal webs

Epidemiology studies the distribution of diseases within and across populations and looks for the causes of such distributions. This way of pitching epidemiological research is underpinned by an explicit causal stance. At times, however, terminology adopted by epidemiologists is much less explicit, replacing ‘cause’ and ‘effect’ with less metaphysically-laden terms such as ‘risk factors’ or ‘determinants’. Whether or not epidemiology ought to adopt an explicit causal stance is certainly an important debate both for the foundations of epidemiology itself and for the implications for public health. However, I won’t pursue this line of argument here. There are two other aspects of epidemiology worth emphasising in order to set the tone of the paper.

The first aspect concerns the status of epidemiology as a discipline mid-way between social science and medicine; this aspect that has been stressed by a number of epidemiologists (to name just a few: Mackenbach (1998), Vineis (1998), Savitz (2003), Vineis (2003), Susser (2004)) and, according to some (most famously, Susser and Susser (1996)), is the distinctive character of modern epidemiology. The second aspect concerns the advancements made in epidemiology thanks to, especially, a better and more widespread use of statistical methods. The prominent use of statistics in epidemiology shouldn’t create a false impression, though. Epidemiology does aim to go beyond mere association and, arguably, statistics in epidemiology is driven by a causal goal, be it to provide a causal explanation of a disease or to inform actions such as public health policies.

Since the early Sixties, epidemiologists introduced in their theoretical framework the notion of ‘causal web’. Yet, it had to wait until the early Nineties to have the first significant critical appraisal of this notion. The locus classicus, or better, the first systematic discussion of the ‘causal web’ is Nancy Krieger’s 1994 paper. Krieger reminds us that the firsts to use the notion of ‘causal web’ were MacMahon, Pugh, Ipsen (1960), as a reaction to the prevalent notion of chain of causation, which failed to account for the complexity of the genealogy of the antecedent of the chain and for the possible partial overlap between different factors. Thus, since the groundbreaking work of MacMahon, Pugh and Ipsen (1960), *multi-factorial* aetiology of disease gained ground. Krieger’s endorsement of a ‘multiple causation’ view or a ‘multi-factorial aetiology’ testimonies this trend. However—and here starts Krieger’s critique of the notion of ‘causal web’—MacMahon, Pugh and Ipsen (1960) did not go far enough in challenging another widespread view at the time, namely biomedical individualism. Biomedical individualism, in a nutshell, is the view according to which causes of disease can always be traced down to individual biological factors. Echoing a distinction introduced by Rose (1985), Krieger indicts MacMahon, Pugh and Ipsen (1960) to not differentiate between causes of cases and causes of incidence. But there are other aspects that Krieger criticises. Notably, there is an important omission in this story about the causal web—this is a story about the *spider*. Krieger is worried that until the spiders, that is the creators of the webs, are provided, there is no *explanation* of the yarns of the web. Moreover, in looking for (causal and explanatory) spiders, we ought not to forget that spiders aren’t just biological. In other words, Krieger is pleading for an *eco-social* view where agent and environment are simultaneously considered to explain disease causation.

Such a strong support for the ‘eco-social’ has been further echoed by eminent names in epidemiology such as Mervyn Susser (1998, and commentators in the same volume). Susser’s vision of multilevel epidemiology, is a steady step beyond MacMahon, Pugh and Ipsen’s ‘web of causation’. Susser develops the idea of Chinese boxes and stresses interconnections between the levels rather than exclusivity of explanation at each level. He distinguishes three levels: social epidemiology, that is the macro level; risk factor epidemiology on behaviours and exposure, that is the individual level; molecular epidemiology, that is the micro level. So here the web of causation acquires a further dimension next to the eco-social: epidemiological explanations of disease may refer to different levels of analysis of disease causation.

Discussions on the shift from the mono-causal model so successful in the 19th century to the multi-factorial eco-social model that emerged since the 1960s, recur in the most recent literature too. For instance, Broadbent (forthcoming) investigates the tension between the mono-causal and the multi-causal model and reaches conclusions much similar to those of Krieger’s: it is not enough to claim that a multi-factorial model has to be used—a multi-causal account has to provide an explanation of disease causation, what is more, a *causal* explanation. In other words, Broadbent too hunts for the spiders. Broadbent’s critique hinges upon the distinction between a ‘*bare* multi-factorial’ model and a ‘*contentful* multi-factorial’ model. Whilst the former simply denies the correctness of the mono-causal model and asserts that

diseases have many causes, the latter, in addition, makes, or ought to make, explicit positive claims about the structure of disease. Broadbent goes on saying that much of the epidemiological literature has in fact supported the *bare* multi-factorial model, thus leaving open and unresolved the complex issue of disease causation. The problem Broadbent points to is not new at all among epidemiologists. For instance, Rizzi and Pedersen (1992), talking about the emergence of multi-factorial model, depict it as

[...] a systematic approach to the selection and ordering of causal factors, and as a consequence, a taxonomical terminology that serves this purpose and also purports to determine general and singular causation in a particular field of medicine (p.253).

Nevertheless, this looks more like a declaration of intents rather than a fully developed account of disease causation. Thus the question remains open. What does the multi-factorial model really amounts to?

The real work, i.e. the *pars construens* of the paper, starts here. In particular, I want to expand on two aspects that the literature, and Krieger in particular, already touched on, but without sufficient depth. *First*, to get the spider(s) we need multiple sources of evidence: notably, mechanistic and difference-making evidence. I will draw a distinction between ‘evidence’ and ‘concept’, and based on this distinction I will reiterate and hopefully better state claims made elsewhere. *Second*, for the eco-social to be explanatory, the web has to be mechanistic. In particular, I will develop the idea of ‘mixed mechanism’ that allows to include in the same mechanistic schema causal factors of very different ontological nature. Also, I will support the notion of ‘mechanistic web’ on the grounds that the mechanistic aspect is desirable in order to take the appropriate action, such as in public health policy.

2. Preliminaries

Before getting into the heart of the contentful multi-factorial model, a number of preliminary remarks are in order. The *first* is about the distinction between evidence and concept. This remark is in order because the philosophical literature on causality has exploded since the 1980s. The accounts on the market are numerous and each one points to different features of causation or to different concepts causation is related to. Yet, what exactly be the status (metaphysical, epistemological, methodological) of such concepts is too often left in a foggy background. The *second* is about the relation between epistemic causality (one of the theories of causality on the market, and the one I actually endorse) and disease causation. This remark is in order because although I agree, by and large, with Broadbent’s critique of the ‘bare multi-factorial’ model, it seems to me that, here again, metaphysical and epistemological/methodological issues conflate. The metaphysical structure of disease causation and questions about methods for causal inference, I shall argue, are perpendicular problems. The *third* is about different meanings of the levels of causation. This remark is in order because the issue often recurs both in the philosophical and scientific literature in the form of side comments rather than in the form of a topic deserving attention on its own.

2.1 Evidence and concept: epistemology/methodology vs. metaphysics

Philosophers showed a renewed interest in causality since the 1980s. After hesitating for several decades as a consequence of the well-known hard attacks launched by Russell (1912-13), Mach (1905) and Pearson (1911), they decided to tackle back the thorny issue of causation. However, the strategy of attack has not been planned carefully enough, resolving in many extraordinarily insightful ideas but muddled as to their scope.

To unravel the causal puzzle, the first thing to do is to identify areas of investigation. Traditional accounts of causality touched on metaphysical, epistemological, and methodological issues. Those areas engage in quite distinct questions about causality. The metaphysics of causality is interested in what causality in fact is, or what the causal relata are. Epistemology, being concerned with knowledge, rather asks questions concerning inference and evidence. Finally, the methodology of causality is interested in developing and implementing successful methods for causal discovery and confirmation in the sciences. Agreed, the border line between epistemology and methodology is most often blurred, but arguably there is still a gain in keeping the distinction: some methodological problems are related to practical difficulties (e.g., because of the data or computer software at hand) rather than theoretical or conceptual issues. Very recently, Nancy Cartwright (2007a) has identified a fourth area in the philosophy of causality: the *use* of causal knowledge. According to Cartwright, it is a mistake to leave problems about using causal knowledge just to policy consultants. Metaphysical, methodological/epistemological, and ‘use’-questions need, instead, to be addressed together.

As I mentioned, the philosophical literature has produced a variety of accounts and of theories of causality where the scope and applications of the claims made are not always crystal-clear. I first offer a very rough course on the current leading contenders, classified according to the main aspect of causality they try to capture, and then point to the confusion they commit to. The interested reader may have a look at Williamson (2007), Dowe (2008), Hitchcock (2008), Schaffer (2008), Woodward (2008), Menzies (2008), Reiss (2009), Russo (2009a, ch.7), for more detailed presentations of those accounts. I will mainly focus on misunderstandings about metaphysical and epistemological/methodological issues, and leave for section 3.2 considerations about the use of causal knowledge as a corollary of the mechanistic character of causal webs.

Difference-making theories. Difference-making theories typically come in three variants: probabilistic, counterfactual, and manipulationist or interventionist theories. In *probabilistic theories* there are three main ideas: (i) positive causes raise the probability of their effect(s), i.e. $P(E|C) > P(E)$; (ii) preventatives, or negative causes, lower the probability of their effects, i.e. $P(E|C) < P(E)$; and (iii), as controversial as it may sound, if events A and B are correlated, then either A causes B , or B causes A , or there is a common

cause C that brings both A and B about.¹ In *counterfactual theories*², to say that an event C is a cause of E is spelled out as follows: had C not occurred, E would not have occurred either. Counterfactual conditionals allow us to pick out single-case causal relations by comparing what happens in the actual world with what happens in other possible worlds. *Manipulationist or interventionist theories*³ cash out causation in terms of the notion of ‘invariance under intervention’. They usually set causation within the framework of statistical modelling and Bayesian nets and claim that a variable C is a cause of variable E if, were we to change C , E would change accordingly and the relation between C and E would remain invariant. All three theories can be labelled ‘difference-making’ because what is specific about causes is that they make a difference to the occurrence or production of the effect.

Physical theories. Physical theories typically come in two variants: process and mechanistic theories. Process theories cash out causation in terms of the spatio-temporal physical process that links the cause-event C to the effect-event E . In the Salmon-Dowe theory⁴, in particular, causal processes are those in which there is a transfer of conserved quantities, such as momentum. Mechanistic theories⁵, instead, say that C causes E if there is a mechanism linking C to E . Mechanisms, currently undergoing a vivid debate, are made of entities and activities organised in a particular way. Both theories can be labelled ‘physical’ because the ‘secret connexion’, as Hume called it, is sought in the physical ground.

Epistemic theory. The epistemic theory⁶ analyses causality in terms of rational belief. According to the epistemic theory, causation is not out there in the world but in our head, so to speak. Causality is not to be reduced to any physical features of the world (such as spatio-temporal processes or mechanisms) but is the mapping from evidence and inference that leads us to form *causal* beliefs about the phenomena we investigate. The more evidence we collect and the more reliable our inference methods, the stronger our causal beliefs will be. Thus, the epistemic theory is antirealist in that it locates causation within the realm of the epistemic activities of a rational agent (here, the scientist) rather than within the basic ontology with which we purport to describe the world. Having *causal* beliefs, rather than just beliefs (albeit rational), carries advantages in many scientific activities such as explanation (we explain by invoking causes) or prediction (knowledge of causes allows more accurate predictive inferences).

¹ See, for instance, Reichenbach (1956), Suppes (1970), Eells (1991), and Arntzenius (2008) for an introduction to the highly controversial issue of the principle of common cause.

² Lewis (1973 and 2004).

³ Woodward (2003), Hausman (1997), Hausman and Woodward (2004), Woodward and Hitchcock (2003).

⁴ Salmon (1984 and 1998), Dowe (1992 and 2000).

⁵ Machamer, Darden and Craver (2000), Glennan (2002 and 2009), Bechtel and Abrahamsen (2005), Craver (2007).

⁶ Williamson (2005, 2006, 2007), Russo and Williamson (2007).

Let us now go back to the three areas of investigation within the philosophy of causality. Difference-making and physical theories have quite a natural metaphysical reading, as they try to capture a fundamental aspect constituting the very nature of causation. Difference-making theories capture the intuition that genuine causes, unlike spurious ones, are responsible for changes in the course of events. Physical theories, instead, capture the intuition that cause and effects must be somehow physically linked (either by a spatio-temporal physical process or by a mechanism). Of course, there also are more elaborated accounts that try to combine the difference-making and the production/physical aspect supporting a pluralistic stance⁷. Leaving aside for the time being issues about pluralism or monism, it is a fact that those accounts shed little if no light at all on inferential, evidential, and methodological questions.

But there is more pressing problem, I think, which I will in fact deal with thoroughly in section 3.1. It is quite uncontroversial that evidence, in any field, is complex. In spite of this complexity in establishing causal relations, we look for evidence showing that the cause made a difference to the occurrence of the effect and that the cause is somehow linked to the effect. The mistake made in difference-making and physical accounts is to reduce the *concept* (or the metaphysical content) of causality to one of its *evidential* components. The fallacious argument goes along these lines: causation requires an analysis in terms of difference-making or in terms of processes/mechanisms depending on the context, therefore causation must be a cluster of concepts rather than a monistic concept with a unique meaning. However, this is fallacious reasoning: from the fact that causation requires multi-fold evidence, it doesn't follow that the *concept* itself be multi-fold. The full argument is given in Russo and Williamson (2007), and in section 3.1 I will offer a slightly different defence, more centred on methodological arguments.

Manipulationist or interventionist theories can be read as providing identity conditions (metaphysics) or test conditions (methodology). The metaphysical reading would go as follows: if a relation between variables C and E is causal, *then* manipulating the cause-variable would change the effect-variable and the relation between C and E would show some kind of invariance. Thus, causation is cashed out in terms of invariance under manipulation or intervention. The methodological reading would instead go as follows: to test whether the relation between variables C and E is causal, one should manipulate the causal-variable in order to see what changes occur in the effect-variable and to see whether the relation would remain stable nonetheless. Thus, the possibility to test causal relations lies in the possibility to manipulate or intervene on the causal relata.

The kind of invariance required for the metaphysical or the methodological reading is, needless to say, not of secondary importance, but this is entirely beyond the point right now. The point at stake right now, instead, is that each reading has its share of problems. On the one hand, endorsing the metaphysical

⁷ For instance, see Glennan (2009).

reading raises the question of whether, reducing causation to invariance, an *epistemic* aspect of causality is again confused with a metaphysical one, which boils down to asking whether causal relations be *intrinsically* invariant (and arguably they don't need be). On the other hand, endorsing the methodological reading, the account falls short in providing test conditions for causal relations in purely observational contexts. In other words, to put the whole burden of the test in manipulability seems to be too strong a requirement, let alone a realistic one.

Finally, the epistemic theory is, at present, the only account that succeeds in drawing a clear line between evidence for a causal relation and the concept of causality (that is, what causality is). The question about the concept of causality pertains to metaphysics, whilst the question about the evidence to form causal judgements upon pertains to epistemology/methodology. The epistemic theory provides us with an account of how from certain types of evidence and from certain inferential methods we can draw causal conclusions. Difference-making, for instance in the form of probabilistic relations between variables, is *evidence* for a causal relations but is not what constitutes causation. Analogously, knowledge about plausible or confirmed links or mechanisms from the cause to the effect is *evidence* for a causal relation but is not what causation is reduced to. This distinction is important for the philosophy of medicine too, as I shall explain next.

2.2 Epistemic causality and disease causation

Given that I just claimed superiority of the epistemic theory over the competing ones, the following issue immediately arises. To what extent does the epistemic theory offer an adequate account of disease causation? Disease causation is normally meant to be the account of how disease is caused. Thus, the most in vogue and accepted account of disease causation in the 19th century was the germ theory of disease, expressed by Henle-Koch's postulates. To put things simply, diseases are caused by microorganisms the causal efficacy of which is to be evaluated according to four criteria: (i) we must find the microorganism in large quantity in all or in the majority of organisms suffering from the disease, but not be found in healthy individuals, (ii) we must be able to isolate the microorganism from a diseased organism and grow it in pure culture, (iii) once introduced into a healthy organism, the cultured microorganism should cause disease, and (iv) the microorganism must be reisolated from the inoculated, diseased experimental host and identified as being identical to the original specific causative agent. This model could finally explain infectious diseases (such as tuberculosis) and provide directions for preventions. However, many other diseases cannot be explained in the same way, and in fact medical doctors and epidemiologists slowly but surely shifted towards a multi-causal model of disease that better suits chronic diseases such as cancer or diabetes. Yet, even multi-causal models have been the object of criticisms, for instance because of too strong an emphasis on biological factors to the expenses of social ones. Thus, although the multi-causal model certainly attracts more consensus and praise than the mono-causal model, it is far from being established as the ultimate and definite account for disease causation.

In the same spirit of the distinction between evidence and concept given in the previous section, debates around disease causation are possibly failing to put the issue in the right focus. Let us begin from what can reasonably be taken as well established. *First*, diseases are of very different nature. Some are infectious, others are chronic, for instance. *Second*, diseases have very different causes. There are at work biological (e.g., helicobacter pylori is a bacterium causing gastric ulcer) as well as psycho-socio-economic causes (e.g., stress is apparently the root cause of many diseases from cardiovascular and stomach troubles, to gum problems, to depression). *Third*, the study of disease may concern a particular individual or a population. Whilst the medical doctor wishes to point to *the* cause(s) that changed the patient's health status, an epidemiologist is rather interested in similarities and dissimilarities between groups of individuals having or not having the disease.

These three simple remarks about disease causation should immediately cast doubt on the models proposed so far. The reason is the following. What disease causation in fact amounts to is a distinct (albeit related) problem with respect to the criteria for causal attribution. In particular, whilst it may be argued that there isn't *a single* model for disease causation, general criteria to evaluate causal attribution in biomedical contexts may and ought to be provided. Recall, under the epistemic theory, 'causation' is cashed out in terms of (causal) belief. So causation is not reduced to the kind of causation described by either the mono-causal or the multi-causal model. Disease causation may be properly described by the mono-causal or the multi-causal model, but that depends on the disease at hand, not on the allegedly specific nature of disease causation. The epistemic theory offers an answer to the question of what *causality* is: it is the causal belief a rational agent, based on all the available evidence.

Consider now the question of what *causes* are. This is still a metaphysical question albeit very different from the former about *causation*. The former question asks what kind of entities the *relata* are. The advantage of the epistemic theory is that it allows different kinds of entities to be causes—that depends on the field, on the methodology, on the object of study, etc. Under the epistemic theory a virus can be rightly considered a cause, as much as a variable in an epidemiological study. For instance, diagnosis may reveal that the bacterium helicobacter pylori is the cause of my gastric ulcer; or, exposure to asbestos dust may be shown to be a strong risk factor for lung cancer in an epidemiological study. Thus the epistemic theory is not forced to endorse one or the other 'traditional' theories of disease causation, which have all been proved, strictly speaking, false. What is key in the epistemic theory, instead, is the *evidence* upon which causal beliefs are formed.

2.3 Levels of causation

The third issue concerns the levels of causation. The problem stems from the (seemingly innocuous) recognition that causal claims may refer to different levels. For instance, there is an intuitively easy

difference to grasp when we say that (i) *Harry's smoking caused him to develop lung cancer* and that (ii) *smoking is a cause of lung cancer*. Philosophers in the probabilistic tradition recognised this difference quite early and, mostly, they quickly identified the trick: even if it is true that, generally speaking, smoking causes lung cancer, it doesn't necessarily follow that Harry will develop cancer even if he smokes, or that he developed cancer because he smoked, since smoking is not a necessary nor a sufficient cause for developing cancer. Thus, the observation that causal relations may go in opposite directions at different levels, pushed some philosophers to argue that two distinct analyses for token or individual-level causation and for type- or population-level causation are to be developed. Typically, philosophers have held the position that population- or type-level causation is related to the frequency with which the putative cause raise the probability of the effect in a group, whilst individual- or token-level causation is related with chance or propensity of the cause to produce the effect in a particular case.⁸

Although such a move has some plausibility, it leaves completely unresolved a number of questions, notably about the *relations* between the type and token level. However, before even trying to get to the centre of the level-labyrinth, it will be worth pointing out that this is a paradigmatic case of dialogue of the deaf between philosophers and scientists. Whilst all call the problem with the same name, philosophers and scientists actually talk past each other because they have different questions in mind. Below I distinguish two senses in which we may talk about the levels of causation: by aggregation and by discipline.

By aggregation: generic / single-case. In the philosophical literature, the pairs of oppositions token/type-level and individual-/population-level are used interchangeably. This is, unfortunately, the source of misunderstandings with practising scientists. The source of confusion lies in the fact that statistical models used especially in quantitative social and biomedical science handle variables that are categorised according to whether they express aggregate or individual data. On the one hand, when scientists use the term individual level, they usually have in mind an analysis that relies on a data set containing information of individuals, e.g. for each observed individual in the study, information about smoking, dietary behaviour or socio-economic class is used. On the other hand, when scientists use the term population level, they usually have in mind analyses that rely on data sets containing information about populations, e.g., summary statistics about the percentage of smoking people, percentage of men, women, of certain age or socio-economic class, etc. This is patently very different from the use philosophers made of the individual/population-level distinction.

A way to dissolve the confusion is using the following distinction instead: generic vs. single-case. The generic/single-case distinction has to do with whether the causal relation is at least in principle repeatable

⁸ See for instance Cartwright (1979), Eells and Sober (1983), Sober (1984 and 1986), and Russo (2009a, ch. 2 and 6) for a discussion.

or whether it has been instantiated. Therefore, the generic/single-case distinction avoids confusions with the unit of observation. Thus, there may be population-level causal relations that are generic and others that are single-case. The distinction is analysed, for instance, in Russo and Williamson (2009) and in Russo (2009a), along the lines of Rizzi and Pedersen (1992), who stated the distinction specifically in the context of disease causation. Of course, many issues need further investigation, for instance what exactly is the relation between the generic level and the single-case; what is the exact metaphysical import of each type of causal claims. This is not the place to solve all these quarrels, but it was worth mentioning because, in the discussions on the causal web, some objected that causes of incidence (generic) are not causes of cases (single-case). I will get back to this point towards the end of section 3.2.

By discipline: mixed mechanism. The above characterisation of the levels suggests that there is a ‘vertical’ structure in causality going from the bottom single-case level to the top generic level. Yet, there is another characterisation of the levels to be made. This has to do with the field in which the causal analysis is carried out. Considering causal analyses of social, physical, biological, or medical phenomena as completely independent and apart is perhaps a heritage of the hyper-specialisation of the sciences. This trend is being opposed, in more recent times, by a renewed interest for interdisciplinary investigations. Yet, the shared idea of interdisciplinary analysis seems to stick to a vertical structure, where a coarse grained analysis at the physical level is integrated with a finer grained analysis at the chemical level. Or, to give another example, a coarse grained analysis of economic behaviour can receive a finer grained analysis in terms of neuronal activity. The kind of interdisciplinarity required for disease causation (and for many cases of social causation), however, may also assume a different facet, where levels are displayed ‘horizontally’. Most importantly, levels are not simply horizontally juxtaposed but they significantly *overlap*. In fact, the social and biomedical sciences model (or ought to model) causal relations that include causal factors of very different sorts, for instance socio-economic *and* biomedical factors. This is different from providing an explanation that goes ‘levels-down’ in the amount of detail provided. In this perspective, social and biological factors are on a par for explaining disease. The need for such ‘horizontal’ level-analysis is mirrored in the *eco-social* perspective that has been supported by many quarters, as we have seen in the previous section. However, the *explanatory* import of the eco-social needs further elaboration—this is the task I shall undertake in section 3.2 developing the notion of ‘mixed mechanism’.

3. ‘Contentful’ causal webs

If causal webs are to be of some utility in better understanding and dealing with disease causation, we have to go beyond the mere denial of the mono-causal model and beyond the formulation of a ‘bare’ multi-factorial model—thus Broadbent (forthcoming) advice. As anticipated, I intend to contribute to this project by giving content to causal webs. The content they need, I shall argue, has two aspects: (i) a pluralistic stance concerning evidence, and (ii) a clarification of the explanatory import of the eco-social.

3.1 Evidential pluralism

Multiple evidence to establish causal claims

The position I defend is evidential pluralism. It is worth emphasising that this is not pluralism about the *concept* of causality but about the pieces of *evidence* needed in order to establish causal claims. Conceptual pluralism, simply put, is the view according to which causality is a cluster of concepts, composed of, say, a mechanistic and a difference-making component. Alternatively, conceptual pluralism may be understood as the view according to which different areas require different concepts of causality, and consequently causation means different things in different contexts. The position defended here can be classified as conceptually monistic and evidentially pluralist. The monistic concept of causation is given by the epistemic theory, as discussed in the previous section. In the following, I focus on the evidential aspect.

To put things straight, the claim I defend is the following. Causal relations are to be established on the basis of difference-making and of mechanistic evidence. Simple as it may look at first sight, this claim hides and hinges on, instead, the whole complexity of the distinction between evidence and concept and raises questions about the interrelations between different types of evidence. The status of the claim is at once theoretical and, as much as possible, normative. A number of arguments may be given to support this claim. One, for instance, exploits descriptive methodology: as a matter of fact, methods used by medical doctors and epidemiologists to establish causal claims are based on the evaluation of those two types of evidence (mechanistic and difference-making). Another argument considers case-studies in the history of medicine: causal claims have not come to be established until both mechanistic and difference-making evidence could be (more or less explicitly) identified. Yet another argument develops on theoretical considerations, in particular on the relations between evidence and concept as discussed in the current debate on pluralism: conceptual pluralists not only commit a fallacy in their argument, but also, by adopting one or another version of conceptual pluralism, they face more problems than the epistemic conceptual monist. Such arguments have been given in Russo and Williamson (2007), where the thesis about evidential pluralism has been first stated, and in Russo and Williamson (2009), where the claim has been refined and investigated in the case of causal inference in autopsy.

In the following, I want to offer a slightly different argument.⁹ But before, let me just linger on a last aspect related to evidential pluralism and epistemic conceptual monism. I said that the claim argued for in this section is normative. But this does not mean that it gives a checklist for causal relations: if mechanistic and difference-making evidence are both present, infer causation, and don't otherwise. Mechanistic and difference-making evidence are not 'monolithic blocks' that are either present or absent. Causal analysis is far from being a schematic and 'algorithmic' activity with step-wise procedures and

⁹ A similar argument is offered in Russo (2009b) to show that, although the methodology of epidemiology *is* essentially difference-making (or variational) in character, a difference-making metaphysics is not a viable option.

checklists to tick. Understanding what kind of evidence is needed in order to establish causal relations is tantamount to understanding what is the relevant information needed in order to draw sensible and reliable conclusions from the data available. This, most often, more resembles making ‘educated guesses’ on the basis of partial and incomplete knowledge than extracting ready-made causal relations from algorithmic machines for torturing data. It is possibly in this sense that Charlton (1996) talks about an epidemiologic mosaic, rather than chain. Charlton (1996 p.106) says that the mosaic—the *causal* mosaic—is “based on evidence but not dictated by evidence”. So it seems that *evidence* is what indicates whether the mosaic is causal or not. Later in this section I discuss how complex multi-fold evidence inform causal attribution.

The argument from methodology

As I mentioned in the opening of the paper, present-day epidemiology makes extensive use of statistical modelling. From this, it would be fairly easy to make the point that statistical modelling handles difference-making evidence. What I am interested in is showing that statistical modelling handles more than merely difference-making evidence. In other words, statistical modelling does not exhaust the evidential component of causation.

Let us begin with presenting the general features of statistical modelling in epidemiology. Savitz (2003, p.35) notices that epidemiology is primarily interested in establishing statements such as ‘the risk of disease is x times greater among exposed persons than unexposed persons’. Such claims contribute to establishing causal relations through *comparative statements*, which are in fact the bulk of a difference-making or variational methodology (for details, see also Russo 2009b). This idea can be found already in Susser (1973, p. 3), where he says that epidemiology is all about comparing and interpreting group exposure and response. Notably, comparisons involve establishing whether factors make or do not make a difference, that is whether distributions of disease differ conditional on the exposure, whether relative risks are greater for exposed individuals than for non exposed individuals.

Epidemiology, as it happens, is most often concerned with observational data rather than experimental data. It is quite uncontroversial that the methodology of experimentation, e.g. in randomised clinical trials, is difference-making because the idea is just this: to see what changes occur by making certain interventions. In fact, in experimental studies we estimate certain predetermined outcomes of a well defined intervention which is deliberately administered to certain individuals, and then we compare results with outcomes in individuals that have not been administered the intervention.¹⁰ Nevertheless,

¹⁰ No doubt the most direct tests for causation would be experimentation and intervention; yet, even randomised controlled trials do not ‘prove’ causation because of practical and methodological pitfalls. A thorough discussion of this issue falls beyond the scope of the paper and I will not reiterate arguments given elsewhere (see e.g., Elwood (1988), Cartwright (2007b), and Glasziou et al. (2007)). Let me just point out that Glasziou et al. (2007) defend the

observational studies rely on a difference-making methodology too. Timmerick (1994, p. 326) expresses this idea very clearly:

Observational method of study is based on the concept that changes which are observed in one trait or variable can cause changes in another characteristics or variable, and those changes occur without the event being altered by the epidemiologist or without intervention by a researcher.

That observational methods in epidemiology are difference-making in character is also clear from the fact that they are all *comparative*. Cohort studies compare individuals exposed to the putative cause with individuals that have not been exposed. Case-control studies compare individuals with the disease with individual that do not have the disease. In cross-sectional studies data is collected at a specific point of time and comparisons are made for that specific moment. It is not my goal here to evaluate strengths, weaknesses, or applicability of those type of studies. Instead, it is my intent to stress their *comparative* aspect: the goal of epidemiological studies is to establish whether and how the distribution or the risk of disease changes according to the exposure and to other medical and social characteristics of individuals.

This isn't breaking news, of course, and it is actually a complaint oft made against statistical methods—that difference-making evidence coming from statistics cannot prove causation in any possible means. The worry, differently put, is that all a statistical model can prove is a correlation between variables. But this is a wrong way of setting the problem. It was stated above that mixed evidence, that is evidence about difference-making *and* mechanisms is needed. This is a shared view amongst epidemiologists too. Although the issue has never been specifically addressed in the epidemiological literature, the need of mechanistic evidence besides difference-making evidence comes out in passing in a number of authors and, mostly, it belongs to generally accepted criteria for causal inference, such as the classic Bradford Hill's (Hill 1965) or the Preamble of the International Agency for Research on Cancer (IARC 2006). So how can we reconcile those *prima facie* contradictory views?

There is a short and a longer answer to this question. I'll give here the short answer and direct the reader to the places where the longer answer is articulated. Mechanistic information is and ought to be used at the model building stage and at the model testing stage. Let me spell out this idea further. In talking about statistical modelling, two sorts of models ought to be distinguished. On the one hand, associational models merely aim at correctly describing co-variations between variables, by keeping statistical

idea that observational evidence can indeed support causal claims, and then provide a significant number of historical examples where convincing causal inferences have been performed without resorting to randomised trials. Timmerick (1994) goes as far as claiming that observational studies provide many more insights into the effects of diseases. The reason is that they deal with population groups, whereas experimental studies deal with individuals or smaller treatments or experimental groups and therefore the inference to relations in large populations is limited. Of course, issues remain concerning the problem of confounding or the use of frequentist rather than Bayesian statistics, but those are quite perpendicular to point I am concerned with.

assumptions to a minimum and avoiding, as much as possible, any causal and/or metaphysical assumption. On the other hand, *causal* models are much richer than associational models in that they have supplementary sets of assumptions that participate in the guarantee of the causal interpretation of results (see Table 1). For instance, some assumptions about the non-correlation between the errors and the covariates in the equations have specific causal import, as assumptions about additivity do. Causal models also include assumptions that are not statistical in character and that belong to the background knowledge of the epidemiologist who is building and testing the causal model. Those assumptions typically concern any relevant information and knowledge of the possible or plausible mechanism(s) explaining the phenomenon at hand. Such knowledge informs and advises on the choice of variables, the structure of the equations, the statistical tests to be performed and the way to interpret the results of tests. The moral so far is that statistical models in epidemiology do handle mechanistic evidence, albeit in the form of background information to be used at various stages of causal analysis.

But there is more than mechanistic evidence in the background of causal models. Causal modelling, in fact, ought to be conceived as the *modelling of mechanisms*. This, however, does not introduce a vicious circle but rather a virtuous one. The reason is that science, and epidemiology in this case, does not establish immutable and eternal causal truths. The whole modelling activity in science—and the biomedical sciences are no exception—is a *va et vient* between established knowledge and knowledge yet to be established. Mechanisms (or partial knowledge thereof) are at once what we base (or ought to base) our causal models and what we want to (or should) establish. The reader interested in the long answer may want to have a look at Russo (2009a, 2009b, 2009c), and Mouchart and Russo (2010), where these ideas are developed specifically in the field of quantitative causal modelling.

Associational Models	Causal Models
Statistical Assumptions	Background knowledge/Causal context
	Statistical Assumptions
	Extra-Statistical Assumptions
	Causal Assumptions
	Hypothetico-deductive Methodology

Table 1: Associational Models vs. Causal Models

What this argument from methodology shows, it seems to me, is that there is no sharp dichotomous distinction between difference-making and mechanistic evidence. Distinguishing the two aspects may be useful for other reasons, for instance because difference-making evidence is most needed (but not exclusively) for prediction whilst mechanistic evidence is most needed (but not exclusively) for

explanation, but it does not follow that these evidential components be completely apart. If they aren't, then what is needed is an account of how the two aspects intertwine and interrelate. Which is the task I undertake next.

The tangle of difference-making and mechanistic evidence

So far I claimed that causal relations are to be established on the basis of difference-making and mechanistic evidence; I backed this argument showing that epidemiological methods, which essentially deal with difference-making evidence, also involve, albeit in less direct ways, mechanistic evidence. There are also other senses in which difference-making and mechanistic evidence are not completely apart. In particular, I am interested in highlighting three ways in which the two evidential components are entangled.

First, difference-making evidence need not be necessarily probabilistic or statistical, that is expressed in a quantitative form. Qualitative difference-making evidence, as the one used in counterfactual reasoning, qualifies as difference-making too. This is what happens, for instance, when coroners reason about causes of death. Coroners appeal to qualitative difference-making considerations to establish whether and to what extent, say, a wound caused death (see Russo and Williamson (2009) on this point). More generally, descriptions of mechanisms may involve qualitative difference-making considerations, for instance when we say that had *helicobacter pylori* not been able to survive in an acid environment such as the stomach's walls, it could not cause gastric ulcer.

Second, as mentioned before, we model mechanisms to back up causal claims (for explanatory, external validity, and other inferential purposes). So what we need, in a first instance, is a *plausible* mechanism to be tested and (dis)confirmed afterwards (a point also made by Gillies (2010)). This means that mechanisms are not given as such, so to speak. The whole modelling activity is meant to figure out what the mechanism is. In order to test and (dis)confirm mechanisms, causal modelling employs statistical strategies that are essentially difference-making. It is worth noting, moreover, that when we talk about mechanistic evidence this does not necessarily refer to 'ready-made' mechanisms. Mechanistic evidence in causal modes is any information that refers to the functioning or *modus operandi* underlying the phenomenon under investigation. Such mechanistic information is most often used *implicitly* in randomised clinical trials (RCTs) that, instead, pretend to extract causal relations from purely statistical hats (on the grounds that randomisation meets the gold standards of causal inference, whatever happens). But this is unfair play on the side of RCTs supporters, because they do use information about known or plausible mechanisms, their parts or functioning, when they set up and evaluate the trials.

Third, as a corollary of the first two points, mechanistic and difference-making evidence cannot be completely disentangled. On the one hand, 'plausible' mechanisms are part of our background knowledge and thus mechanistic evidence, although in quite an indirect form, is used to establish causal relations in

‘difference-making’ procedures such as statistical modelling. Conversely, difference-making evidence is also part of mechanisms and of modelling mechanisms. In Russo (2009, ch.6) I argue that we model mechanisms based on which variational or difference-making relations between variables hold. According to others (e.g., Craver 2007), counterfactuals, which are difference-making epistemic devices *par excellence*, are test conditions for mechanisms or even part of the definition of mechanisms themselves (e.g., Glennan 2009).

3.2 Explanatory eco-social webs

If epidemiologists are so enthusiastic about the eco-social turn, it may sound surprising that so much space be devoted to debating the plausibility of eco-social webs. Unfortunately, such a peaceful consensus is a mirage. This is for two reasons. The first is that the ‘eco-social’ turn has been much less smooth than it might be thought. The second is that the turn has been only partial—so far the ‘mechanistic’ component of the eco-social has been overtly neglected.

The eco-social turn in epidemiology

The eco-social is not such a well-established perspective as it may look at first sight. In his article on the features of traditional and modern epidemiology, Pearce (1996) recalls and discusses some of the tensions that occurred in the past between the two main paradigms in epidemiology: traditional and modern epidemiology. In the 19th century the pioneering works of, for instance, John Snow—who carried out remarkable investigations on the cholera epidemics in London in 1849 and 1854—raised debate and frictions concerning public health measures that, he urged, had to consider *social* factors too. It is worth noting that at that time epidemiology was a branch of public health and consequently focused on disease prevention at the population level, unlike the clinical sciences more focused on disease pathology and individual treatment. Methodological developments for the measurement of disease occurrence gradually shifted the focus from the population to the individual thus giving rise to the new paradigm of *modern epidemiology*. In the 19th century, then, the development of germ disease theory marked an important step towards epidemiological investigations on the individual rather than the population, thus focusing on individual risk factors. Biological agents were seen as the ‘real’ causes (according to the Henle-Koch postulates, microorganisms—that is ‘real and existing’ entities—cause diseases) thus displacing the interest of ‘traditional epidemiologists’. Epidemiology had to wait until the mid 20th century for a renewed interest in *social* causes of disease. However, this has been inspired more by the denial of the mono-causal model in favour of a multi-causal perspective, rather than by a deep and serious recognition of the causal effectiveness of social factors.

If Pearce’s reconstruction of the history of epidemiology is correct, then it may sound naïf of Krieger (2001) to say that the eco-social should be a truism, “since social as well as biological processes inherently shape population health” (p.693). Krieger is perfectly aware of the highly controversial character and of the of important conceptual consequences of such a statement. This is why Krieger insists on the

specificity of social epidemiology and provides a glossary explaining how the ‘social’ plays (or should play) an active causal role in traditional topics investigated by epidemiologists. What is key in social epidemiology, unlike traditional approaches, is that social determinants of health are not just in the background, but are *active* part of the causal story. The matter is precisely taking the social out of the background and putting it into the model as an *active and real* causal factor. It is not enough to include in the model social factors *qua* categories to partition the population, in order to minimise the action of confounding variables. Social factors, in other words, have to be regarded as ‘real’ causes. Consequently, a deeper understanding of how this has to be practically done and of how this conceptually works is urgently needed.

To claim that the social is a ‘real’ cause of disease certainly raises important metaphysical issues. Nevertheless, I would like to leave for another occasion metaphysical quarrels and to pursue here again a line of argument axed on methodology and epistemology rather than metaphysics. Krieger says that the way forward for social epidemiological methods is multi-level analysis. The reason is that

[...] if guided by well developed conceptual models clearly specifying which variables are to be studied at which level, these analyses can potentially assess whether individual’s health is shaped by not only ‘individual’ or ‘household’ characteristics (for example, individual or household income) but also ‘populational’ or ‘area’ characteristics. (p.695).

The multilevel approach is also advocated by Pickett and Pearl (2001) and, before, by McPherson (1998) or Vineis (1998), to name just a few. For the sake of history, Vineis traces the interest in the eco-social even back to Rudolf Virchow who supported the ‘social medicine’ already in the mid-19th century. So the need to integrate social factors into disease causation is not just a new trend of modern epidemiology but has remote roots in the very early stages of epidemiology.

If the reader too sympathises with the eco-social perspective, she will certainly rejoice in learning about this support for the eco-social on epidemiologists part. Nevertheless, I would like to play the devil’s advocate and indict epidemiologists of ‘conceptual naivety’. The multilevel or multi-causal model they so strongly support is not pushed far enough. What is missing is not only the spider, as Krieger (1994) said, but exactly *what kind* of spider. In the remainder of this section I will plead for a spider that creates a *mechanistic* eco-social web.

The mechanistic turn in philosophy

To understand what such mechanistic eco-social web amounts to, we must take a step back and start from the very recent philosophical literature on mechanisms.

Philosophers only just turned the attention to mechanisms. The leading accounts developed by Machamer, Darden and Craver (2000) and by Bechtel and Abrahmsen (2005), albeit in slightly different forms, endorse a hybrid metaphysics made of entities and activities showing some kind of organisation or structure. Glennan (2002), moreover, stresses that mechanisms are always ‘mechanisms for behaviour’ thus emphasising the functional import of mechanisms and of mechanistic explanation.

In spite of the nuances, there is a point on which the leading contenders all agree—mechanisms are not to be reduced to Salmon/Dowe-like processes. Roughly put, whilst mechanisms are, in slightly different ways, characterised as an organised assembly of entities interacting with each other in one way or another, processes are more akin to world-lines à la Russell (1948). Thus, the mechanistic turn may be seen as progress with respect to process theories of causation, according to which, roughly, ‘ A causes B ’ means that there is a spatio-temporal physical process linking A to B . However, many issues remain controversial and debated as of today. Those concern the hybrid metaphysics of entities and activities, the concepts of function and of organisation, of multi-, intra-, inter-level mechanisms, etc. Much discussion is also devoted to epistemological and methodological issues, that is about how we come to know about mechanisms. For instance, the works of Craver (2007) and of Darden and Craver (2002) offer a story about how we come to know about mechanisms in the natural—and particularly in the biomedical sciences—and Russo (2009a) and Mouchart and Russo (2010) tackle the issue in the context of quantitative causal modelling in social science.

What I am most interested in, though, is the core of agreement that can be extrapolated by abstracting from the metaphysical debate. Thus, the issue is quite perpendicular to metaphysical concerns about the entity-activity ontology. This agreement, it seems to me, concerns the explanatory power of mechanisms. In fact, independently of the precise metaphysical characterisation of mechanisms, they can be taken as ‘epistemic devices’ to produce successful explanations. This is the sense in which Machamer, Darden and Craver (2000), and Darden and Craver (2002), Machamer (2004) talk about ‘mechanisms schemata’. Or the sense in which Bechtel and Abrahmsen (to appear) oppose the recur to laws in psychological explanation.

To be more precise, the explanatory power of mechanisms lies in providing more understanding of the phenomenon by adding information, that is by maximising the quantity and quality of information. Such information concerns the elements of the mechanism (entities and activities) as well as their organisation. The more we can tell about the elements and the functioning, the better our understanding of the phenomenon. It is in this sense that Mouchart and Russo (2010) claim that statistical models in social science are explanatory insofar as a joint probability distribution over the initial set of variables is *decomposed* into marginal and conditional components, and the *decomposition* is interpreted as representing a mechanism. Also, it is in this same sense that Bechtel and Cory (forthcoming) and Bechtel and

Abrahamsen (forthcoming) say that psychology explains by *decomposing* a mechanism into parts—this amounts to ‘vivisection’ the mechanism to understand its parts and inner functioning.

Merging the eco-social and the mechanistic turns

Now we have the eco-social turn in epidemiology on one side and the mechanistic turn in philosophy on the other side. The task is to merge them into a coherent account. This can be done in the following way.

Eco-social webs have to be explicitly characterised and modelled as *mixed mechanisms* where both social and biomedical variables play explicit causal roles. This poses a noticeable metaphysical problem. Social and biomedical variables have very different metaphysical status. Usually, social variables express concepts or conceptual constructs, such as ‘education’, that can be measured in a variety of ways, or ‘socio-economic status’, that doesn’t exist ‘as such’ but its measure depends on how it is defined, etc. Conversely, biomedical variables do have ‘physical counterparts’: a bacterium is either physically present or absent in the body, patients ‘physically’ smoke, etc. Combining social and health variables into the same mechanism is not a challenge just for social epidemiology. Demographers have been facing the same problem too. John Caldwell (1979), for instance, suggested that, in developing countries, variations in levels of maternal education explain variations in child survival rates. However, Caldwell’s thesis acquired theoretical plausibility and depth only when Mosley and Chen (1984) developed an analytical framework in which indirect paths from social factors to health factors were made explicit in *mechanistic* terms. Agreed, Mosley and Chen (1984) did not explicitly formulate the analytical framework in mechanistic terms, but this is the most straightforward and cogent interpretation of their model. The interested reader may have a look at Russo (2009a, ch.6) where the case study is discussed at length.

Let us go a step further. The reason why we need to make mixed paths explicit is to maximise the quantity and quality of knowledge. Adding information about intermediate paths increases understanding of the phenomenon exactly because it fills the details of the *functioning* of the mechanism. And this is what gives the mixed mechanism metaphysical plausibility, in spite of the hybrid status of the figuring causal factors.

Eco-social webs and levels of causation

Krieger (1994) criticises the ‘received’ notion of causal web, among other reasons, on the grounds that it does not differentiate between determinants of health in individuals and in populations, and, perhaps more seriously, it gives too much an emphasis on individual risk factors. This, says Krieger, is clearly a consequence of the underlying framework of *biomedical individualism*. It is not my intention to fight a battle against biomedical individualism here, but I’d just like to make a few remarks that hopefully clarify some claims made in the epidemiological literature.

Krieger puts the whole problem in terms of individual- vs. population-level risk factors (or determinants) of health. She refers to the distinction made by Rose (1985) between the ‘causes of cases’ vs. the ‘causes of incidence’. Rose distinguishes between two aetiological questions. (1) Why do some individuals have hypertension? Such a question underlies epidemiological studies on the characteristics of *individuals*. (2) Why do some populations have much hypertension, whilst in others it is rare? Such a question, on the contrary, underlies epidemiological studies on the characteristics of the *population*. Whilst the first hunts for individual risk factors, that is causes of cases, the second hunts for population-level risk factors, that is causes of incidence. Rose notices that determinants of incidence are not necessarily causes of cases. Although the intuition is clear enough, the causes of cases vs. causes of incidence distinction needs some further elaboration. And the elaboration it needs comes from the most recent philosophical literature on the levels of causation, which hopefully will shed further light on the quarrel on biomedical individualism.

Instead of using the couple of concepts individual- vs. population-level, let us try with the following: single-case vs. generic level. Single-case causal relationships are instantiated at a particular time and space. Thus, the cause of *Harry's* hypertension counts as a single-case factor. Generic causal relations are instead repeatable, at least in principle. Thus, individual characteristics, such as being overweight or obese, are generic causes as much as population-level characteristics, such as dietary habits (e.g., northern European countries consume fatter food and less olive oil).

The take home message for supporters of causal webs is that the critical target is not biomedical individualism *as such*, but the fact that, contrary to what biomedical individualists hold, individual risk factors participate in the explanation of a generic phenomenon—e.g., hypertension—as much as population-level factors. However, individual risk factors do not explain *yet* causes of cases, such as *Harry's* hypertension, because a particularised diagnoses is needed to this end. Thus, it is misleading, if not incorrect to say, as McDowell (2008) does, that risk factors explain at the individual level, if by individual level is actually meant the *single-case* level. But McDowell (2008) is definitively right in saying that for public health purposes many levels of investigation are needed. Not just investigation at the individual and population level, but also investigation involving different disciplinary areas. Which leads me to the last topic of ‘contentful’ causal webs.

Eco-social mechanistic webs for action

Eco-social mechanistic webs are not just important for explanation, as suggested above, but also for planning public health policy. Let me point up this idea by means of an example. The HIV infection epidemic well illustrates how various level components are involved at the same time: biomedical and social, individual and population.

The first thing worth noting is that HIV epidemic is considered as *global*, that is involving not just populations but many populations. Although the ultimate goal is to reduce the epidemic at the population

level, ultimately this can be done if we protect and cure *individuals'* health. Thus, it seems necessary to have an understanding of both the determinants at the population level as well as of individual risk factors.

There is also another aspect of the HIV epidemics that is directly concerned with the issue raised by eco-social webs. Fee and Krieger (1993) present and discuss different paradigms of disease that influenced and shaped understanding of AIDS. The first paradigm is that of infectious disease. According to this view AIDS was considered as a 'gay plague', a devastating epidemics due to individual characteristics of infected people. Consequently, no interest whatsoever was paid to social factors related to AIDS. The focus was solely on the mechanisms of the disease, especially at the cellular level and genetic determinants. The second paradigm is that of AIDS as chronic disease on which to intervene over the long run. In this paradigm there was a shift of focus from aetiology to pathology and from prevention to potential therapies. Fee and Krieger (1993) criticise the two paradigms because both are underpinned by a biomedical model that is typically reductionist in that it claims the primacy of disease aetiology on explanation, it focuses on disease mechanisms, and it views social factors as secondary if not irrelevant—all these are the well-known premises of biomedical individualism.

The alternative paradigm to endorse considers AIDS as at once a biological and social disorder: it is a collective chronic infectious disease and a persistent pandemic. First, considering it as a collective disorder shows the inadequacy of individualistic perspective and puts emphasis on how class, race, and gender may affect people's living conditions and therefore their health status. Second, considering it as a chronic disease puts emphasis on the potentially prolonged duration and therefore on prevention rather than therapy. Third, considering it as an infectious disease highlights the fundamental aspects of aetiology, modes of transmission, and possibilities of prevention. Finally, considering it as a persistent pandemic puts emphasis on the long time frame of the disease and global impact.

If all these dimensions—biomedical and social, individual- and population-level—are considered, then we can have a different regard on interventions. In fact, interventions are themselves at different levels. Pearce (1996) endorses a 'multilevel' view of analysis and points out that public health actions are not to be taken just at the population level. And this is what happens in the case of the HIV epidemics. There are in fact interventions at the 'level of discipline/area': some are biomedical (e.g. female controlled methods such as microbicides and diaphragm, HIV vaccines), and others are rather behavioural (e.g. induced changes in sexual and drug-use behaviour). But there also are interventions at different levels of the hierarchical structure of society: prevention programs may target individuals, families, communities, societies, or any combinations of these.

The moral to be drawn from the example of HIV infection is that we can plan better actions—that is make better *use* of causal knowledge, as Cartwright (2007a) would say—if the webs are at once mechanistic and eco-social. This, notice, does not imply that no action is possible unless full knowledge of the mixed mechanisms is known, but just that we should do our best to provide the most detailed explanation of disease causation to help setting up efficacious policies. For policy purposes, it is important to spell out the details of the functioning to decide where exactly to intervene. Consider again the case discussed above about the influence of maternal education on child survival in developing countries. Knowledge of the indirect and mixed paths from the social factor of education is important because it may take too much time to improve women’s education in order to reduce child mortality. It may be more effective, time-wise, to intervene on biological factors in the short run and to develop social policy for education in the long run. But all these decisions are possible, for the best, only insofar as mechanistic eco-social webs are modelled.

Conclusion

Slowly but surely, the metaphor of the causal web gained ground amongst epidemiologists. As many epidemiologists and philosophers noticed, this metaphor has certainly done a great job in opposing the dominant and in many ways too simplistic idea of the chain of causation. The ‘chain’, in fact, was at once meant in a temporal sense and in an ‘unidimensional’ metaphysical sense (the mono-causal model). Nowadays, it is fairly uncontroversial to claim that this picture does not faithfully represent disease causation. Thus, the advent of the ‘multi-causal’ models, as opposed to the ‘mono-causal’ model of disease causation, has been generally well accepted. However, as Broadbent (forthcoming) lucidly noticed, adopting a ‘multi-causal’ model that does not specify the explanatory import of causal relations is not good enough. It is in this very same sense that Nancy Krieger, in her 1994 pioneering article on the web of causation, denounced that the spider—i.e., the creator of the web—was still missing.

Given this background, I attempted to expand on Krieger’s notion of ‘causal web’ and on Broadbent’s notion of ‘contentful multi-causal’ model. After some preliminary distinctions on the metaphysical, epistemological or methodological scope of different issues in the philosophy of causality, I argued for the need of multi-fold source of evidence—notably mechanistic and difference-making evidence—reiterating and clarifying arguments offered in previous joint and sole-authored works. In particular, I tried to give details about how the two evidential components intertwine, and how mechanistic evidence can be encapsulated into difference-making evidence and vice-versa. Finally, I elaborated further on the explanatory power of the eco-social arguing that what is needed is an explicit *mechanistic* characterisation of the web. The mechanistic turn in the philosophy of causality, I suggested, can be of much help in giving the eco-social strong explanatory power. But the virtues of mechanistic eco-social webs do not exhaust themselves for explanatory tasks, but are also needed for action, in particular for setting up public health policies.

Acknowledgements. I am grateful to the British Academy for funding this research. I am also grateful to Lorenzo Casini, David Corfield, George Darby, Jon Williamson, and Phyllis McKay Illari in particular, for extremely stimulating and helpful comments.

References

- Arntzenius F. (2008), Reichenbach's Common Cause Principle, *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2008/entries/physics-rpcc/>>.
- Bechtel W. and Cory D. (forthcoming), What is psychological explanation?, in P. Calvo and J. Symons (eds), *Routledge Companion to philosophy of psychology*, London: Routledge.
- Bechtel W. and Abrahamsen A. (2005), Explanation: a mechanist alternative, *Studies in the History and Philosophy of the Biological and Biomedical Sciences*, 36, 421-441.
- Bechtel W. and Abrahamsen A. (forthcoming), Dynamic mechanistic explanation: computational modelling of circadian rhythms as an exemplar for cognitive science, *Studies in History and Philosophy of Science Part A*.
- Broadbent A. (forthcoming), Causation and models of disease in epidemiology, *Studies in History and Philosophy of Biological and Biomedical Sciences*.
- Caldwell, J. C. (1979). Education as a factor in mortality decline: an examination of Nigerian data. *Population Studies*, 33(3), 395-413.
- Cartwright, N. (1979). Causal laws and effective strategies. *Noûs*, 13(4), 419-437.
- Cartwright N. (2007a), *Hunting causes and using them: approaches in philosophy and economics*. Cambridge: Cambridge University Press.
- Cartwright N. (2007b), Are RCTs the gold standard?, *Biosocieties*, 2, 11-20.
- Charlton B.G. (1996), Attribution of causation in epidemiology: chain or mosaic?, *Journal of Clinical Epidemiology*, 49 (1), 105-107.
- Craver C. (2007), *Explaining the brain*. New York: Oxford University Press.
- Darden, L. and Craver, C. (2002). Strategies in the interfield discovery of the mechanism of protein synthesis. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 33(1), 1-28.
- Dowe, P. (1992). Wesley Salmon's process theory of causality and the conserved quantity theory. *Philosophy of Science*, 59, 195-216.
- Dowe, P. (2000). *Physical causation*. Cambridge: Cambridge University Press.
- Dowe, P. (2008), Causal Processes, *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/causation-process/>>.
- Eells, E. (1991). *Probabilistic causality*. Cambridge: Cambridge University Press.
- Eells, E., Sober, E. (1983). Probabilistic causality and the question of transitivity. *Philosophy of Science*, 50, 35-57.
- Elwood, M. (1988). *Causal relations in medicine. A practical system for critical appraisal*. New York: Oxford University Press.
- Fee, E. and Krieger, N. (1993). Understanding AIDS: Historical interpretations and the limits of biomedical individualism, *American Journal of Public Health*, 83, 1477-1486.
- Gillies, D. (2010), The Russo-Williamson Thesis and the Question of whether Smoking Causes Heart Disease, in P. McKay Illari, F. Russo and J. Williamson (eds), *Causality in the Sciences*, Oxford University Press.
- Glasziou P., Chalmers I., Rawlins M., McCulloch P. (2007), When are randomised trials unnecessary? Picking signal from noise, *British medical journal*, 334, 349-351.
- Glennan, S. (2002). Rethinking Mechanistic Explanation, *Philosophy of Science*, 69 (September 2002), S342-S353.
- Glennan, S. (2009). Mechanisms, causes and the layered model of the world. *Philosophy & Phenomenological Research*, forthcoming.
- Hausman, D. (1997). Causation, Agency, and Independence. *Philosophy of Science*, 64(4 Supplement), S15-S25.
- Hausman, D. and Woodward, J. (2004) Manipulation and the Causal Markov Condition. *Philosophy of Science*, 71(5), 846-856.
- Hill, B. (1965). The environment of disease: association or causation?, *Proceedings of the Royal Society of Medicine*, 58, 295-300
- Hitchcock, C. (2008), Probabilistic Causation, *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/causation-probabilistic/>>.
- IARC (2006), *Preamble to the IARC Monographs*, <http://monographs.iarc.fr/ENG/Preamble/index.php>
- Krieger, N. (1994). Epidemiology and the web of causation: has anyone seen the spider? *Social Science and Medicine*, 39, 887-903.

- Krieger, N. (2001). A glossary for social epidemiology, *Journal of Epidemiology and Community Health*, 55, 693-700.
- Lewis, D. (1973). Causation. Reprinted with postscripts in D. Lewis, *Philosophical Papers Vol. II* (pp. 159–213). Oxford: Oxford University Press.
- Lewis D. (2004). Causation as influence, in John Collins, Ned Hall and L. A. Paul (eds), *Causation and Counterfactuals*, The MIT Press, 75-106.
- Mach, E. (1905). *Knowledge and error*. Dordrecht: Reidel. 1976.
- Machamer P. (2004). Activities and Causation: The Metaphysics and Epistemology of Mechanisms. *International Studies in the Philosophy of Science*, 18(1), 27-39.
- Machamer, P., Darden, L., and Craver, C. F. (2000), Thinking about Mechanisms, *Philosophy of Science*, 67, 1-25.
- MacMahon B., Pugh, T. F. and Ipsen, J. (1960). *Epidemiologic methods*. Boston: Little Brown and Co.
- Mackenbach, J. (1998). Multilevel ecoepidemiology and parsimony, *Journal Epidemiology and Community Health*, 52, 614-615.
- McDowell, I. (2008). From risk factors to explanation in public health, *Journal of public health*, 30(3), 219-233.
- McPherson, K. (1998). Wider “causal thinking in the health sciences”, *Journal Epidemiology and Community Health*, 52, 612-613.
- Menzies, P. (2008). Counterfactual Theories of Causation, *The Stanford Encyclopedia of Philosophy* (Fall 2009 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2009/entries/causation-counterfactual/>>.
- Mosley, W. H., Chen, L. C. (1984). An analytical framework for the study of child survival in developing countries. *Population and Development Review*, 10 (Supplement), 25–45.
- Mouchart, M. and Russo, F. (2010). Causal explanation: recursive decompositions and mechanisms, in P. McKay Illari, F. Russo and J. Williamson (eds), *Causality in the Sciences*, Oxford University Press.
- Pearce, N. (1996). Traditional epidemiology, Modern epidemiology, and public health, *American Journal of public health*, 86(5), 678-683.
- Pearson, K. (1911). *The grammar of science*. London: A. and C. Black.
- Pickett, K., Pearl, M. (2001). Multilevel analysis of neighbourhood socioeconomic context and health outcomes: a critical review. *Journal of Epidemiology and Community Health*, 55, 111–122.
- Reichenbach, H. (1956). *The direction of time*. Berkeley: University of California Press.
- Reiss J. (2009). Causation in the social sciences: evidence, inference, and purpose. *Philosophy of Social Sciences*, 39, 20-40.
- Rizzi, D. A. and Pedersen, S. A. (2002). Causation in medicine: Towards a theory and terminology. *Theoretical Medicine*, 13, 233-254.
- Rose, G. (1985). Sick individuals and sick populations, *International Journal of Epidemiology*, 14, 32-38.
- Russell, B. (1912–1913). On the notion of cause. *Proceedings of the Aristotelian Society*, 13, 1–26.
- Russell, B. (1948). *Human knowledge*. Simon and Schuster, New York.
- Russo, F. (2009a), *Causality and causal modelling in the social sciences. Measuring variations*, Springer, New York.
- Russo, F. (2009b), Variational causal claims in epidemiology. *Perspectives in biology and medicine*, 52(4), 540-554.
- Russo, F. (2009c), Causal arrows in econometric models, *Humana.Mente*, 10, 25-40.
- Russo, F. and Williamson, J. (2007), Interpreting causality in the health sciences, *International studies in the philosophy of science*, 21(2), 157-170.
- Russo, F. and Williamson, J. (2009). Generic vs. Single-Case Causality: the Case of Autopsy. PhilSci Archive.
- Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world*. Princeton, NJ: Princeton University Press.
- Salmon, W. C. (1998). *Causality and explanation*. New York: Oxford University Press.
- Savitz, D. A. (2003), *Interpreting epidemiologic evidence: strategies for study design and analysis*. Oxford: Oxford University Press.
- Schaffer, J. (2008). The Metaphysics of Causation, *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/causation-metaphysics/>>.
- Sober, E. (1984). Two concepts of cause. *PSA: Proceedings of the Biannual Meeting of the Philosophy of Science Association*, 1982, vol. 2, 405–424.
- Sober, E. (1986). Causal factors, causal influence, causal explanation. *Proceedings of Aristotelian Society*, 60, 97–136.
- Suppes, P. (1970). *A probabilistic theory of causality*. Amsterdam: North Holland.
- Susser, M. (1973). *Causal thinking in the health sciences. Concepts and strategies of epidemiology*. New York, London: Oxford University Press.
- Susser, E. (2004). Eco-Epidemiology: thinking outside the black box. *Epidemiology*, 15(5), 519-520
- Susser, M. (1998). Does risk factor epidemiology put epidemiology at risk? *Journal Epidemiology and Community Health*, 52, 608-611.
- Susser, M. and Susser, E. (1996). Choosing a future for epidemiology: II. From black box to Chinese boxes and eco-epidemiology, *American Journal of Public Health*, Vol. 86, Issue 5 674-677.
- Timmerick, T. C. (1994). *An introduction to epidemiology*. Boston, London: Jones and Barlett Publishers.
- Vineis, P. (1998). Epidemiology between social and natural sciences. *Journal of Epidemiology and Community Health*, 52, 616-617.
- Vineis, P. (2003). Causality in epidemiology. *Social and preventive medicine*, 48, 80-87.

- Williamson, J. (2005), *Bayesian nets and causality: philosophical and computational foundations*. Oxford University Press, Oxford.
- Williamson, J. (2006). *Causal pluralism versus epistemic causality*. *Philosophica*, 77, 69-96.
- Williamson, J. (2007). *Causality*. In Dov Gabbay & F. Guenther (eds.), *Handbook of Philosophical Logic*, volume 14, Springer, 95-126.
- Woodward, J. (2003). *Making things happen: a theory of causal explanation*. New York: Oxford University Press.
- Woodward, J. (2008). Causation and Manipulability, *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2008/entries/causation-mani/>>.
- Woodward, J. and Hitchcock, C. (2003). Explanatory generalizations, Part I: A counterfactual account. *Noûs* 37(1), 1-24.