

Inferring causation in epidemiology: mechanisms, black boxes, and contrasts

This is the final draft (September 2010) of a paper forthcoming in Causality in the Sciences, Phyllis Illari McKay, Federica Russo and Jon Williamson (eds), Oxford University Press, 2011.

Abstract

This chapter explores the idea that causal inference is warranted if and only if the mechanism underlying the inferred causal association is identified. This *mechanistic stance* is discernible in the epidemiological literature, and in the strategies adopted by epidemiologists seeking to establish causal hypotheses. But the exact opposite methodology is also discernible, the *black box stance*, which asserts that epidemiologists can and should make causal inferences on the basis of their evidence, without worrying about the mechanisms that might underlie their hypotheses. I argue that the mechanistic stance is indeed a bad methodology for causal inference. However, I detach and defend a mechanistic interpretation of causal generalisations in epidemiology as existence claims about underlying mechanisms.

1. Causal hypotheses in epidemiology

What does it take to establish a causal hypothesis in epidemiology? What standards need to be met? Or, if establishment comes in degrees, degrees of what?

The most obvious aspect of this problem concerns inferring causation in a particular study. A study reveals a statistical association between smoking and lung cancer, or a certain gene and obesity. Statistical analysis reveals a low *p*-value – a measure of the chance that the association is due to chance. Study design controls for confounding variables (what philosophers would call common causes of the putative cause and effect). Can it be inferred that, for *this* group, a causal relationship exists between smoking and lung cancer, or having that gene and obesity?

Oddly enough, this is not a question that epidemiologists like to answer. A single study would not normally be considered a sufficient basis for a causal inference. Replication is a guiding epidemiological principle. From a methodological point of view this is extremely interesting. Epidemiologists' credence in a causal hypothesis about Study Group A increases when the effect is replicated in Study Group B. Explaining (or, I suppose, refuting) this attitude is a central task for any methodological analysis.

A second difficulty concerns the inference from a study, or a collection of studies, to a wider population. Epidemiologists are centrally concerned with extrapolating from the people they study to people they have not studied. Replication is important here too, because one way to argue that differences between the population studied and the target population are causally irrelevant is to replicate the study among people who are drawn from the target population. However, replication cannot solve the problem of generalisation. Often the study group will *already* be drawn from the target population: for example, when generalising from the Whitehall studies to the population of Britain.¹ Differences between those studied and those not studied will always remain; the difficulty is working out when these differences make a difference. On other occasions, studies on a subset of the target population may be impractical: for an obvious example, consider future populations. Quite generally, a central purpose of epidemiology is to get more for less: to learn something not only about those who *have* been studied, but about those who *have not*.

Epidemiologists make efforts to be precise about the scope of their claims, by explicitly stating whether they are intended to apply to the group studied, or to a wider population, and if the latter what conditions the wider population are to meet. For example, instead of saying “genetic influences cause paediatric obesity”, they might say:

Genetic influences on BMI and abdominal adiposity are high in children born since the onset of the pediatric obesity epidemic.

(Wardle et al. 2008, 398)

However, even this is an incomplete specification. The children studied were British, but the obesity epidemic affects Europe and America too. Are these results evidence for high genetic influences on BMI and adiposity in children in Britain only, or in Europe and America too? The incompleteness of the specification is not necessarily a failing of the authors of the study. It reflects the genuine difficulty of deciding how to generalise.

Note that replicating the study in European or American children is a way to *avoid* the question, not to answer it. Replication cannot tell us whether a generalisation from this study to American children would be *warranted*, only whether it would be *correct*. In

¹ The Whitehall studies concerned various aspects of social status and health among British civil servants (Marmot 2006). The aim was to identify social determinants of health, especially the role of purely social differences in explaining differences in health, by studying a group of people whose absolute wealth was such that their basic biological needs were met.

circumstances where we can replicate, that may be the best strategy; but for reasons I have already given, we cannot universally substitute replication for generalisation.

Thirdly, there is a difficulty interpreting general claims, even when their scope is fixed. A great deal of philosophical attention has been directed towards singular causal claims, such as “Jones’ smoking caused his lung cancer”. But epidemiologists are almost exclusively concerned with general claims, such as “smoking causes lung cancer”. Does the latter express a relationship between smoking and lung cancer, or is it a generalisation over individual causal relationships – along the lines of “*in X% of cases, smoking causes lung cancer*”?

This difficulty is a relatively familiar one to philosophers, but it is perhaps not the most pressing one for epidemiologists. In practice, epidemiological hypotheses are explicitly exception-ridden. Accordingly they are framed not as universal generalisations, but as measures of the *influence* of one factor on an outcome, or measures of the *strength* of an association, or of the *proportion* of an effect that is due to a particular factor or group of factors. These sorts of claims raise what is fundamentally the same problem, but in a slightly different way. For example, saying “Genetic influences... are high” makes it clear that the generalisation is not *exceptionless*. But it still does not make clear exactly how the degree of influence is to be interpreted. Is the claim that, in each individual, the genetic influence is high? – This would amount to a universal generalisation attributing a certain genetic influence to each individual. Such an interpretation is hard to make sense of on either the effect side or the cause side. On the effect side, obesity might be absent in some of the individuals studied. Genetic factors cannot then be said to influence it. Switching from a qualitative property (obesity) to a quantitative one (such as bodyweight) will not always be straightforward: the absence of effects such as lung cancer, diabetes, and suicide are hard to interpret as zero degrees on any quantitative scale. Similarly on the cause side, it makes little sense to attribute some degree of influence to a factor that is absent. This is not clear in the example I have picked, since “genetic influences” are always present in people, but it is obvious when we consider single-gene studies. When somebody lacks the gene but has the trait in question, it makes no sense to attribute the trait’s presence to the influence of the absent gene in *any* degree.

Another interpretation would see measures of influence, proportion, strength of association, and similar, as measures of the *proportion of cases in which a factor is causal*. (This interpretation is akin to the generalisation-over-singular-causation view of universal

causal generalisations.) This view is easy enough to make sense of, but there is a case that it reflects metaphysical commitments rather than epidemiological evidence. Take a measure such as heritability, which is the proportion of a given trait in a given population that is due to genetic factors. The idea that causes are either present or absent, strictly speaking, and the view that we cannot quantify the contribution of a *particular* cause to a *particular* effect, are widespread among philosophers from John Stuart Mill to David Lewis (Mill 1843; Lewis 1973). On such a view, heritability expresses the proportion of the *population* in which the trait is caused by genetic factors – but in each *individual*, the trait either is or is not caused by genetic factors. But this interpretation is not stable, because on this metaphysical picture, causation is not exclusive. Saying that a trait is caused by genetic factors in an individual is compatible with saying that it is caused by non-genetic factors: events have many causes. In the study I have been using as an example, the aim is to measure the *contribution* of genetic influence to obesity in a population. To interpret this as a claim about the proportion of individuals in whom genetic factors cause obesity would be bizarre, since genetic factors are part of the causal history in 100% of cases of obesity. Indeed, every trait is both 100% genetic and 100% environmental, on this interpretation (Rothman and Greenland 2005, S146). Better, we could see it as a claim about the proportion of individuals in whom genetic factors *make the difference* between being obese and not. This interpretation might be made to work; but that would be no trivial philosophical achievement. The interpretation of heritability is a topic of considerable dispute (e.g. Schonemann 1997; Sesardic 2005).²

Two lines of response to this bundle of difficulties may be discerned in the contemporary methodological-epidemiological literature. These lines of response are in tension. One is the *mechanistic stance*: the view that causal inference in epidemiology aims at discovering mechanisms: that discovering mechanisms is necessary and sufficient for establishing a causal hypothesis. The other is the *black box stance*: the view that epidemiology is primarily concerned with statistical analysis of associations, and only incidentally concerned with uncovering mechanisms. In §2 and §3 I will describe and evaluate each stance, and in §4 I will propose a resolution.

² Heritability is not to be confused with heredity. An individual can inherit a trait from her parents, for example her eye colour. The disposition to develop blue eyes given a certain broadly congenial environment is hereditary; favourite colour probably is not. The heritability of a trait is defined with reference to a population, and makes no sense applied to an individual. It is meant to be a measure of the *relative* contribution of genes and environment, and is not fixed for a given trait. Hence Wardle's interest in showing that obesity is *still* highly heritable despite an increase in the availability of calories.

My terms “stance”, “line of thought”, and similar are intended to avoid commitment on the question of whether any actual epidemiologist wholeheartedly asserts any of the views discussed. I rather doubt that any does. Nonetheless, these are not straw men: these stances are evident in the methodological writings of actual epidemiologists, and there is value in seeking to draw them out into the light for explicit evaluation, even though – indeed, partly because – nobody would endorse these views when stated explicitly and taken to their logical conclusions.

2. Mechanisms

There has been a surge of interest in mechanisms in recent philosophy of science. One well-known definition is this:

A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations.

(Glennan 2002, S344)

The word “mechanism” is less common in epidemiology than it is in some other biological sciences. Nevertheless, it would be interesting if a similar idea could be identified in actual epidemiological practice.

I think it that it can. Epidemiologists, like neuroscientists, use the term “mechanism”, and they do so in revealing ways, as I shall argue. But in addition, I suggest that it is plausible to take the common epidemiological phrase “causal pathway” as referring to a mechanism. Perhaps a causal pathway is not exactly the same thing as a mechanism: for one thing, it may be longer, and include several “mechanisms” in the sense intended by Glennan. Nevertheless, causal pathways probably do meet the proffered criterion for mechanisms, since they typically will be postulated to explain a “behaviour”, and will plausibly constitute a “complex system” whose parts interact according to “direct, invariant, change-relating generalizations”. Moreover, identifying a mechanism in neuroscience, and identifying a mechanism in epidemiology, satisfy the same goal: they both explain *how something works*. For these reasons, it is plausible to see the search for causal pathways in epidemiology as a search for what Glennan and others call “mechanisms”, even if the causal pathways identified in epidemiology are not terribly similar to the mechanisms in neuroscience.

The notion of mechanism neatly captures one methodological story that may be discerned in the epidemiological literature. That story has two parts. Initially, associations

are identified between variables and health outcomes. By a sort of process of elimination, it is established that it is very unlikely that the association is due to chance, or to confounding variables (common causes, to philosophers) or other biases. A variable for which such an association has been established is called a *risk factor*. This first part of the process provides a good reason to think that the risk factor is causal, but a second stage is required for a conclusive case. The *mechanism* for the operation of the risk factor must be identified. Perhaps not immediately; but if in the fullness of time no mechanism is identified, the credibility of the hypothesis will suffer. Moreover establishing a mechanism is sufficient for proving that a causal hypothesis is correct: showing *how* A actually causes B is conclusive proof that A *does* cause B.³ The *mechanistic stance* is the methodological position that identifying the underlying mechanism is both necessary and sufficient for warranted inference to a causal hypothesis.

As a description of contemporary epidemiological methodology, the mechanistic stance is appealing. It has plenty of illustrations. Genetic epidemiology is an extremely good fit. Consider this extract from the introduction of another high profile clinical study in the genetics of obesity:

The genetic contribution to body weight has been established through family studies, investigation of parent-offspring relationships, and the study of twins and adopted children... As is the case for height... environmentally driven changes in body weight in the population occur against a background of susceptibility to weight gain that is determined by genetic factors. Thus, genetic approaches can be applied to understand both the molecular and physiological mechanisms involved in human obesity.

(Farooqi and O'Rahilly 2006, 710)

The rationale appears to follow exactly the lines of the mechanistic stance I have described. There is considerable evidence for a causal link between genetics and weight; the next step is to understand the mechanisms involved.

A similar sentiment is evident at what might be seen as the other end of the epidemiological spectrum, among those who work on social determinants of health. They, too, see it as crucial to identify plausible pathways for the determinants they

³ Showing that a given mechanism is how A *actually* causes B must not, of course, be confused with exhibiting a mechanism by which A *could* cause B. The entailment goes through only when what has been shown is how A *actually* causes B.

identify. Here is an extract from another introductory rationale, this time from a chapter by two leading proponents of social determinants of health research:

Psychosocial factors and their influences on health are active areas of research... There is now enough evidence to suggest that this is an important field for those concerned with improving public health... Plausible mechanisms linking psychosocial factors to health are described in the first half of this chapter. We then look to the evidence from both human and animal literature to illustrate the ways in which social organisation can influence our biology and, therefore, the health of individuals and populations.

(Brunner and Marmot 2006, 8)

This passage mixes the mechanistic stance as a purely methodological stance with the natural companion view that identifying mechanisms is a good way to improve public health. Setting that aside, the mechanistic stance is clearly discernible. There is already evidence that psychosocial factors influence health; the purpose of the chapter is to work out how, by postulating “plausible mechanisms” and presenting “evidence from both human and animal literature” in support. Whether or not this is the authors’ intention, they certainly give the impression that providing a mechanism a key element in the case for the hypothesis that social status is a determinant of health.

In each of these cases, the identification of mechanisms is seen as important: important enough to devote an entire paper to. Why? It *could* be simple scientific curiosity: the desire to find out how things work, just for the sake of it. But I think there is more to it than that. Identifying mechanisms is presented as important not only to complete the scientific picture, but also to seal the case for existing causal hypotheses. This is especially clear in the social determinants of health literature, where the identification of pathways is seen by proponents and detractors alike as crucial for the case that the socioeconomic factors identified as health determinants really do have the effects claimed for them. Michael Marmot’s Whitehall studies provide evidence for a causal link between social status and health, among British civil servants – a population whose basic biological needs (food, water, shelter) are amply met. Marmot’s efforts to confirm this hypothesis have not focused merely on replicating the results in different populations (though that is of course one area of activity). A considerable amount of effort has also been devoted to identifying mechanisms by which social status might affect health.

This suggests a methodological thesis: that identifying an underlying mechanism is both necessary and sufficient for establishing a causal hypothesis.⁴ I am not suggesting that Marmot or anyone else endorses this thesis: indeed it may be that no actual epidemiologist would sign up to it, in that blunt form. But identifying mechanisms clearly *is* relevant to establishing causal hypotheses in epidemiology, and setting up a somewhat extreme but clear stance this may be a useful technique for exploring that relevance.⁵

It is not hard to see why identifying a mechanism might be considered necessary and sufficient for establishing a causal hypothesis, at least in epidemiology. An argument for its being necessary might appeal to the fact that epidemiology is clearly not a fundamental science. Causal associations identified in epidemiology presumably exist in virtue of the way that things are organised at a more fundamental level. The existence of causal associations at the population level is naturally seen as arising from certain regularities in the way that the members of that population are organised, and regularities in their environment. Identifying a mechanism is just identifying the “events and activities” of the population’s members and environment which together give rise to the population-level association. If no mechanism can be identified, then the status of the population-level causal association remains mysterious. Failure to identify such events and activities is not always evidence that they do not exist, of course, and perhaps this is why Austin Bradford Hill famously urges that “biological plausibility” be treated with caution (Hill 1965). Nevertheless, a hypothesis for which no mechanism is *remotely* plausible, or for which no mechanism is discovered after a long period of time, remains at best tentative.

It is similarly obvious why identifying a mechanism underlying a causal association might be considered sufficient for establishing the corresponding causal hypothesis. If the events and activities giving rise to an association are identified, then it presumably follows that the variables which the hypothesis asserts are causally linked at least *can* be causally connected. There is an interesting twist here, however. Showing that a mechanism exists by which, say, stress can cause poor health, does not bear directly on

⁴ Where there is more than one distinct mechanism underlying a given association, presumably all the underlying mechanisms would have to be identified to satisfy the spirit of this requirement.

⁵ The mechanistic stance is *not* supposed to be the view that identifying mechanisms is the goal of epidemiological research: only of causal inference. Causal inference may have other goals, such as public health intervention, or indeed informing further causal inferences in an iterative process.

the claim that stress *does* cause poor health in any particular population. The mechanisms identified by Marmot are ways in which the results of the Whitehall studies *might* have come about. This explains his two-part strategy outlined in the excerpt quoted previously, of first identifying mechanisms and then arguing for their actual operation in humans and animals. So identifying a mechanism is not on its own sufficient for establishing a causal hypothesis; a further inference is required to the claim that the identified mechanism is indeed the explanation of the causal association asserted by the hypothesis in question. What is sufficient, then, to establish a causal hypothesis, is the identification of the mechanism *actually* responsible for the association, not merely a mechanism which physically *could* be responsible for it. I take it that the fundamental motivation for this stance is that it follows from the claim that A causes B in a particular way, that A causes B simpliciter. Showing *how* A causes B is only possible if, or in other words entails that, A does in fact cause B.

The mechanistic stance responds as follows to the difficulties facing causal inference in epidemiology which I identified in §1. In answer to the question when we are justified in inferring causation for a studied group, the mechanistic answer is presumably, “Not until a mechanism has been identified.” A single epidemiological study will not usually identify a mechanism, perhaps explaining why single studies typically provoke replication and further research, rather than a causal inference. But if epidemiology seeks to uncover mechanisms, then there may be another, more subtle reason that replication is important. Replicating in Study B an association observed in Study A provides evidence that the same mechanism underlies *both* associations; it seems *prima facie* less likely that two different mechanisms gave rise to the two associations, than that one did.⁶ This explains why replication by Study B can confirm a hypothesis about Study A: because Study B provides further reason to think that some mechanism underlies both studies, and thus provides further reason to think that some mechanism underlies the association first observed in Study A. And showing that some mechanism underlies an association is sufficient for showing that the association is indeed causal, on the mechanistic stance. This methodological stance therefore explains and vindicates the reluctance of epidemiologists to make causal inferences on the basis of individual studies, and the importance they attach to replication.

⁶ This is an application of what is sometimes called the Common Cause Principle.

The mechanistic stance also provides useful clarification of the other difficulties we identified, the problem of generalising from studies to population, and the question of how to interpret general causal hypotheses in epidemiology. Generalising to a wider population is typically safer when the mechanism underlying a causal association has been identified, because knowledge of the mechanism yields detailed knowledge of what differences are relevant to the association. For example, our knowledge of the mechanism underlying the analgesic effect of paracetamol enables us to identify the circumstances relevant to this effect, and thus to say whether the effect will hold in a very wide range of circumstances. Moreover, it also gives us a lot of other useful knowledge, for example, about *other* associations – other effects of taking paracetamol. These uncontroversial facts motivate a methodological idea: that the generalisation of an association observed in a particular study to a wider population is only really warranted when the mechanism underlying the association has been identified. Generalisation before a mechanism has been identified may of course be required in some circumstances, but until the mechanism has been identified, generalising remains a sort of guessing, according to the mechanistic stance: because we can't be sure exactly *when* our causal generalisations will hold, until we know *why* they hold.

This suggests an interpretation of general causal claims, such as “smoking causes lung cancer” or “genetic influences on BMI and abdominal adiposity are high in children born since the onset of the pediatric obesity epidemic”. Such claims are to be interpreted, first, as asserting the existence of a mechanism – or perhaps several – linking the causal and effect variables, and second, as claiming that the operation of this mechanism(s) is (are) in fact what underlies the causal association between the two variables. The mechanism may or may not be known. When it is, this interpretation is very natural. “Paracetamol causes pain relief” is naturally precisified, not by specifying more exactly the probability that an analgesic effect will be observed in various circumstances, but rather by saying more about the way paracetamol works. Detailed description of the mechanism *yields* information about whether an analgesic effect will be observed in a wide range of circumstances, in an efficient way. When a mechanism is not known, the causal generalization is a sort of stand-in: a claim that some unknown mechanism does link the variables.

The exception-ridden nature of epidemiological generalizations is unproblematic, on this view. Suppose a causal generalization is just a claim that a mechanism underlies the association between two variables *C* and *E*. That is to say, in individual cases where an

individual c leads to an individual e , a mechanism connects c and e . The exception-ridden nature of the generalization reflects the fact that sometimes, C -events occur, but the mechanism by which they give rise to E -events is absent, and so E -events fail to occur. Likewise, the fact that C -events can by a certain mechanism give rise to E -events does not preclude E -events from coming about some other way. Causal generalizations are claims about the actual linking of individual C -events and E -events by actual instances of a mechanism; such claims are entirely compatible with C -events occurring without the E -producing mechanism, or E -events occurring as the result of some other mechanism. The content of the causal generalization is that a certain mechanism *is responsible for the causal association between the identified variables*; it is not a claim about the association itself, but about how the association arises. Further support for this interpretation of exception-ridden causal generalizations derives from one of the points argued in the last paragraph, that when a mechanism is known, it is natural to precisify a causal generalization by specifying the underlying mechanism, rather than offering more statistical detail about the circumstances in which the association holds. Whether or not it is a model for other sciences, it seems to be a good fit for epidemiology.

The one trouble previously identified which the mechanistic stance does not seem to handle is the worry about measures of proportion of influence, such as heritability. It can handle causal generalisations amounting to associations of less than 100%, such as “smoking causes lung cancer”, in the manner already indicated. But a claim such as “pediatric obesity is over 70% heritable” does not submit to a mechanistic interpretation, since presumably mechanisms linking genes to weight exist in every human being. I am, however, inclined to think that this reflects badly on the concept of heritability, and similar attempts to apportion causal responsibility; such concepts are hard to make sense of in any analysis. It may be that they have no clear sense, or that they need further clarification (cf. Lewontin 1974; Schonemann 1997). Accordingly I propose to set them aside, and focus on interpreting those epidemiological hypotheses that clearly do make sense.

The mechanistic stance has a lot to recommend it, then: it provides a neat interpretation of several tricky features of actual epidemiological practice, and thus vindicates that practice. The idea that epidemiologists ought to first identify risk factors, and then look for the mechanisms underlying them, sounds like both sensible methodological advice and a fair description of actual epidemiological methodology. In

the next section, however, I will identify and amplify some doubts that thoughtful epidemiologists have expressed about the mechanistic stance.

3. Black boxes

Notwithstanding the foregoing, there are some reasons to doubt that the mechanistic stance provides a good methodology for causal inference in epidemiology. Both the necessity and the sufficiency of discovering a mechanism for inferring causation are open to criticisms of principle, and these criticisms can be illustrated by actual episodes in epidemiology.

Let us start with the idea that discovering underlying mechanisms is necessary for the inference of general causal hypotheses. The motivation identified in the previous section for this methodological principle is that the existence of a mechanism is necessary for the existence of a general causal relationship. On the mechanistic interpretation, a causal relationship exists at the population level between two variables only if the particular instances of those variables are related by a mechanism. I suggested that the general causal claim could be interpreted as a claim about the existence of an underlying mechanism. If so, it is natural enough to require that we *discover* this mechanism when we assert a general causal claim. So an inference to a general causal claim requires, as warrant, the discovery of an underlying mechanism – sooner or later.

This “sooner or later” indicates a difficulty, though. Our failure to discover an underlying mechanism does not mean the mechanism doesn’t exist. Moreover, when we are entirely ignorant of a mechanism, it seems that we are often quite incredulous about the existence of one. Combined with the view that a mechanism must underlie causal associations, this native incredulity can lead, and on occasion has led, to an unreasonable prejudice against hypotheses when we do not see how the hypothesised causal link might work. The view that mechanisms need to be identified for a causal inference to be secure can colour our assessment of the evidence for new hypotheses or against accepted hypotheses. I will suggest that this is because it is not really a methodological principle: it tells us what to aim at (discovery of mechanisms) but not how to achieve that goal. As a consequence, it tilts the balance in favour of existing knowledge, and inhibits what it recommends – the discovery of mechanisms about which we do not yet know.

Several well-known episodes from the history of epidemiology illustrate these claims. Perhaps most famously, the miasma theory of the nineteenth century offered a mechanism for the transmission of disease, based on the movements of “miasms”

(roughly, bad air). The fact that the theory purported to give a mechanism for disease transmission was its principle virtue. Nineteenth century epidemiological heroes such as John Snow and Ignaz Semmelweis were criticised for failing to identify plausible mechanisms for their causal hypotheses, leading to unnecessary loss of life in both cases. Snow argued, on the basis of incredibly careful door-to-door inquiries, that a causal connection existed between water supply and cholera (Cameron and Jones 1983). Semmelweis that differential childbed fever rates between two wards in a Vienna hospital were caused by the dirty hands of medical students, who worked in one ward but not the other (Carter 1994). Both hypotheses were resisted by the authorities and many doctors (although Snow was somewhat more persuasively successful than Semmelweis), and the principal reason given was that no plausible mechanism for the transmission of disease along these vectors had been identified. It was not until some decades later that a mechanism which might plausibly underlie their respective hypotheses was identified, in the shape of microbial theory (Carter 2003 esp. Ch 3, 4).

J.P. Vandebroucke argues that Snow had a “contagionist” hypothesis about the mechanism of disease transmission, and thus that this is not an example of “black box” epidemiology (Vandebroucke 1988). Nevertheless, it is doubtful that Snow had what we now view as a *correct* theory of cholera transmission (and contagionism is certainly false as a *general* theory, since it is possible to contract diseases other than from another diseased person, e.g.: tetanus; puerperal fever). It is important to distinguish psychological claims about what leads scientists to their theories, from methodological claims about what justifies or otherwise warrants those theories. Snow’s results are impressive *to modern epidemiologists*, even though we now believe he was wrong about the mechanism of disease transmission. His proposed mechanism cannot, therefore, be part of the warrant which we *now* accept for his conclusions (whatever he thought); yet we still regard those conclusions as warranted to a high degree by the evidence he procured. Semmelweis also hypothesised a disease mechanism, namely the resorption of animal-organic matter leading to the decaying of the blood (Gillies 2005). But similarly, acceptance of Semmelweis’s resorption/decaying blood mechanism surely cannot be part of *our* reason for thinking that any of Semmelweis’s causal hypotheses were well-founded, since we believe the resorption/decaying blood mechanism is not what underlay the associations he identified, for instance between disinfecting hands and reductions in differences in mortality between the two wards. Yet even though we reject his account of the underlying mechanism, this does not prevent us from accepting Semmelweis’s causal

hypothesis that disinfecting hands caused the reduction. Vandenbroucke's argument stumbles on exactly the point I am trying to make, confusing the discovery of mechanism as a *goal* of causal inference, with the discovery of mechanism as *method*. I am suggesting that it is a good goal, but a lousy method. It is not their theories of the mechanism of disease transmission which have elevated Snow and Semmelweis almost to hero status in the eyes of many modern epidemiologists; and it is not the failure of their contemporaries to appreciate these mechanism-theories that is so often lamented. Rather, it is the way they procured evidence to support their causal hypotheses and to refute counter-hypotheses, which is so widely admired; and the way that evidence was ignored that is lamented.

Thus these episodes cast doubt on the usefulness of a methodological principle stating that discovery of mechanism is necessary for warranted causal inference. At least three doubts are distinguishable. First, the obvious logical point that has already been mentioned: that the inference from no known mechanism to no mechanism is a tricky one. Second, demanding that a mechanism be identified before a causal inference is accepted simply seems to be an unreasonable position, because it seems possible to have excellent evidence for a causal link, without understanding how the link works. Even if the mechanistic *interpretation* is plausible, and general causal claims are to be interpreted as existence claims about underlying mechanisms, it does not follow that a general causal claim is only warranted when the underlying mechanism is identified. It is possible, in an epidemiological context, to know that there is a causal link – and therefore that a mechanism must exist, on the mechanistic stance – yet not know what that mechanism is. The opponents of Snow and Semmelweis are generally considered to have been *unreasonable* to doubt the extremely convincing evidence for a causal link; this is difficult to explain if warrant for causal inference requires the identification of a mechanism.

Of course, in these cases, a mechanism *was* eventually discovered. But (and this is the third doubt) the discovery came *after* the causal hypothesis was well-established. The mechanistic methodology gets things the wrong way round, in these cases. The discovery of a mechanism *can* of course help to confirm a causal hypothesis, but a causal hypothesis can also be solidly confirmed well before the underlying mechanism is known. Therefore discovery of an underlying mechanism is not a necessary condition on warranted causal inference.

It must be remarked that drawing morals from historical episodes is a delicate business, because it is possible for different commentators to see different lessons in the same episode. For example, while most commentators (of a moral-drawing sort) would agree that something went wrong in the Semmelweis episode, Federica Russo and Jon Williamson see the episode differently. They insist that, to establish a causal claim, it *is* necessary to identify an underlying mechanism. But they do not adopt the mechanistic stance as I have outlined it, because they do not think that identifying a mechanism is sufficient for causal inference: they hold that it is also necessary to provide what they variously refer to as “statistical” and “probabilistic” evidence. In support, they cite episodes where causal hypotheses were supported by evidence of one kind but not of the other, and were rejected on that basis. Thus they cite the Semmelweis episode to illustrate the claim that “the relation between contamination and puerperal fever... was not accepted until backed up by mechanistic evidence, i.e., until the germ theory had been developed” (2007, 11), as an instance of their more general thesis that identifying mechanisms is necessary (though not sufficient) for causal inference. They go on to propose a theory of causation which is intended to fit this methodological picture.

This line of argument makes me uneasy, because I am unsure whether it is meant as a descriptive account of causal inference (then and now), or as a normative account of the standards which *ought* to be used when deciding whether to infer causation. Suppose we grant (for the argument) Russo-Williamson’s descriptive claim, that Semmelweis’s contemporaries rejected his theories because of a lack of “mechanistic evidence”, which I take to mean a lack of any then-acceptable theory about what the underlying mechanism for the proposed causal association might be. In this sense, indeed, Semmelweis did fail to establish his various causal hypotheses: he failed to provide evidence which was *in fact* compelling, as demonstrated by the fact that his evidence did not compel many of his contemporaries. But in another sense he clearly *succeeded* in establishing (at least some of) his causal hypotheses: he provided evidence which, in the eyes of most modern epidemiologists, his colleagues *ought* to have taken more seriously. In particular, the evidence for the efficacy of his proposed intervention – disinfecting (not merely washing) hands – is extremely strong. And replication would have made it stronger, without necessarily advancing knowledge of underlying mechanisms.

Unfortunately a purely descriptive reading of Russo and Williamson’s claim renders it largely irrelevant from a methodological point of view, and does not justify or explain why they themselves treat it a motivation to seek a theory of causation, apt for the health

sciences. If the claim is merely descriptive, then we may conclude that Semmelweis's contemporaries were simply *wrong* to insist that a mechanism be identified before they accepted any causal connection. (And so we do not need a new theory of causation that would be compatible with this insistence.) Suppose, then, that Russo and Williamson intend their claim to be normative. Then they are making a normative claim that that mechanistic (as well as statistical) evidence is necessary for *good, rational, warranted* causal inference. On this reading, the lesson Russo and Williamson draw from the Semmelweis episode is that Semmelweis's theories *ought* not to have been accepted until knowledge of underlying mechanisms was obtained. (This explains why they offer a theory of causation intended to justify this stance.)

But if this is indeed what Russo and Williamson are claiming, then I fear they need a far stronger argument than any they supply. For then they are committed to the startling view that, had germ theory not come along and the underlying mechanism remained a mystery, we today would be rational to dismiss Semmelweis's work, no matter how much evidence we had gathered in the meantime about the efficacy of disinfecting hands. That is not a view which many modern epidemiologists would share. Modern epidemiologists set very high store in some methods, such as the randomized control trial, which involve no requirement to identify underlying mechanisms. Empirical evidence suggests that this view would cost lives if it were adopted by modern epidemiologists, and that indeed it did cost lives if it was in fact the reason that Semmelweis's claims were not accepted more promptly. In short, the descriptive Russo-Williamson thesis does not (without further argument) support any claims about causation, nor bear on questions about how causation *ought* to be inferred; while the normative thesis is false by the methodological standards prevalent in epidemiology, and also, arguably, in light of the empirical evidence from episodes such as the Semmelweis case itself.

In fact, on the mechanistic interpretation, it is not surprising that a causal hypothesis can be well confirmed before the underlying mechanism is discovered. On the mechanistic interpretation I have suggested, a causal generalisation asserts that there *is* some mechanism that is responsible for the association in question. We could have good reason to believe that there is some mechanism, yet not know what it is.⁷ Moreover, this

⁷ This is an instance of the more general fact that we can know *that* something exists without knowing what it is. For example, I know that there *is* something holding the floor up, but I don't know what it is; you may know that there *is* something in the cave, but not what; etc.

order of events suggests a plausible story about how we discover mechanisms when we previously had no idea about them. We make warranted inferences to causal generalizations; these generalizations imply the existence of underlying mechanisms; and we then conduct further research to find the mechanisms. We *know* where to look. It seems, then, that the mechanistic metaphysics does not after all motivate the corresponding methodological principle that the discovery of underlying mechanism is necessary for warranted causal inference.

Stepping for a moment beyond the confines of the methodology of epidemiology, there seems to be little intuitive support for the idea that causal inferences require knowledge of mechanisms as warrant. There are many everyday cases where we take ourselves to have knowledge of causal relations, without having the slightest idea about the mechanisms underlying them. I know that the clear Turkish liquor, raki, goes cloudy when water is added, but I don't know what the underlying mechanism is. Presumably someone does, and it might be suggested that I can have a warranted causal belief just as long as I have recourse to an expert who can explain the mechanism to me. But the Turks have known for centuries that raki goes cloudy when water is added. It is absurd to insist that, for centuries before the underlying mechanism was known, the Turks did not know that mixing water and raki in roughly equal quantities caused the cloudiness they routinely witnessed.

Moreover, to insist as a general matter that underlying mechanisms must be identified before a causal inference is warranted raises a dilemma. As we go on uncovering underlying mechanisms, either we will come up against causal relations for which no underlying mechanism can be discovered, or we will not. If we do, then none of our causal knowledge will be secured, because we will have reached a point where we are unable to discover underlying mechanisms and therefore unable to obtain the warrant we sought for our higher-level causal inferences.⁸ If we do not, then we will never finish uncovering underlying mechanisms, and thus again we will never obtain the warrant we seek for our causal inferences. As a general methodological principle, then, the

⁸ Might it be objected that warrant of a higher-level causal hypothesis can be conferred by knowledge of an underlying mechanism, even if that underlying mechanism contains causal links for which yet lower-level underlying mechanisms are not known? No: because (on this view) we are not warranted in believing that the causal links in the mechanism underlying our higher-level hypothesis, until we have identified the mechanisms underlying them in turn. Without that warrant, we do not know (on this view) that they are causal links: for all we know (on this view), our putative mechanism may be a coincidental dance of its parts.

requirement that underlying mechanisms be identified, before a causal inference is warranted, is a guarantee that we can never have causal knowledge. If it has any applicability then it must be confined to particular domains, such as causal inference in epidemiology; but historical episodes previously alluded to suggest that it does not work even so confined.

These famous historical episodes have more recent echoes. In an influential (but not uncontroversial) report for the US government, Richard Doll and Richard Peto argued that many environmental causes of cancer could be identified from careful analysis of epidemiological evidence (Doll and R Peto 1981). This approach suggests that epidemiologists should treat diseases as “black boxes” (Richard Peto 1984), and that the identification of a causal mechanism is not necessary for a warranted causal inference. Accordingly, epidemiologists need not concern themselves with the discovery of mechanisms, but can directly attack causal questions without worrying about the mechanisms underlying the hypotheses they generate. As in the cases of Snow and Semmelweis, the pragmatic benefits of this approach are evident. If Doll and Peto are correct, then labouring to uncover mechanisms may well prove to be a waste of time and money, from a public health point of view. Especially where environmental causes are concerned (smoking being the best-known example), refusing to make a causal inference until a mechanism is known can be seriously detrimental to public health. Doll and Peto’s recommendation typifies what I will call the *black box stance*.

A great deal of contemporary research is, I think, undecided between the mechanistic and the black box stances. On the one hand, epidemiological research largely proceeds by identifying associations and applying various statistical tests and methodological principles to form a view about whether these associations are causal. On the other hand, the explosion of identified risk factors has not produced a corresponding increase in the scope of our understanding of the conditions studied, nor has it been accompanied by a corresponding explosion in public health or medical interventions. This is not just a case of technology lagging: it is also due to the fact that the causal hypotheses in question do not seem terribly reliable. For example, studies seemed to show that hormone replacement therapy reduced risk of heart disease, and public health policies were implemented on this basis, before subsequent studies found the opposite effect (Rutter 2007). A slightly more subtle problem is that directly translating a causal hypothesis into a public health intervention may have unintended consequences (an instance of the generalisation problem). For example, beta blockers administered after surgery appeared

to reduce the risk of death by heart attack, but subsequent studies showed that they increased the risk of death overall by increasing risk of death by stroke and other conditions (Deveraux 2008). This, of course, is grist to the mill of the mechanistic stance, because one of the chief benefits of discovering the mechanism underlying an association is that it often comes with information about *other* associations, and so makes unintended consequences of this sort less likely.

Requiring that a mechanism be identified before a causal hypothesis is accepted may be too strict; but it does at least have the merit of clarity. It is easy to tell *whether* a mechanism has been postulated; moreover testing the hypothesis by replicating the *mechanism* may sometimes be a more straightforward (because lab-based or clinical) business than replicating the association itself in a large observational study. Whereas analysing the methodology of a published study in order to form a view as to the security of its results is devilishly difficult. Indeed it may be impossible, depending on the accuracy and completeness of the published methods section.

One solution to this tension, then, would be to require the discovery of mechanisms as an admittedly too strong necessary condition on causal inference. In pragmatic contexts where great harm appears to be a real possibility, but where a mechanism cannot be identified, some other decision-theoretic principle, such as the precautionary principle, might be appealed to. This stubborn mechanistic stance is, however, a last resort, because of the difficulties we have been discussing, and because identifying sound principles for decision-making under uncertainty is itself a difficult task. Even erring on the side of safety is not straightforward. The HRT/heart disease example shows that we can be wrong, not only in our causal inferences, but also in deciding which way it is safe to err.

So far I have been focusing on the view that mechanism discovery is necessary for causal inference. What about the claim that it is sufficient? Here, the mechanistic stance might appear to be on stronger ground: showing *how* A causes B seems to entail that A *does* cause B.

Unfortunately, that does not mean that “look for mechanisms” is a good methodological principle. It has at least two serious weaknesses. First, it is too vague. Ironically, it does not tell us *how* to look for mechanisms. It states a goal, but gives no indication how to get there. This, I suggest, explains why the mechanistic stance has on occasion led to a bias towards existing knowledge. Mechanisms we already take ourselves

to know about satisfy the methodological directive; but the same directive doesn't help us find mechanisms we don't know about.

Second, the search for mechanisms can mislead, because it can allow us to believe we have achieved an understanding of something when we have not. Showing *how* A causes B indeed entails showing *that* A causes B. But giving information about an event's causes is not sufficient to explain that event, nor to allow you to devise an effective intervention. If I explain my late arrival by telling you that I was born, I am citing a cause of my late arrival: but I am probably not providing a good explanation of my late arrival. Similarly, identifying a mechanism by which, say, hydrochloric acid leads to ulceration of the stomach lining is indeed sufficient for showing *that* the presence of acid causes ulceration. But there is an explanation for the presence of excessive acid, in the case of many sufferers, namely the presence of bacteria, *Helicobacter pylori*. Showing that stomach acid causes ulceration by identifying the mechanism is a good method for proving causation; but epidemiologists (like other scientists) are typically interested in more than cataloguing the causes of the phenomena they study. They are interested in explaining them, and intervening to change them. Each of these goals plausibly requires the identification of causes. But not just *any* causes. The mechanistic methodology misleads because it provides a sufficient criterion for causation, but no guidance on whether the "right" causes have been identified, or what the "right" causes are.

This philosophical point is also well-illustrated by famous historical episodes. I have already mentioned the discovery of *H pylori*. Thoroughly documenting the mechanism by which hydrochloric acid causes stomach ulceration did not lead to the discovery of *H pylori*. Moreover, the hypothesis that peptic ulcer might in many cases be an infectious disease was initially treated with considerable scepticism – because it was thought that bacteria could not survive in such an acidic environment as the stomach. What led to the discovery of *H pylori* was initially chance observation of unknown bacteria in patients with peptic ulcer, followed up by observational studies, clinical work on the bacteria, and a dramatic piece of self-experimentation (for a summary see Angel 2008, Ch 2).⁹ Of course, the discovery of the mechanism by which *H pylori* causes ulcer is also an important feature of this episode. I don't mean to deny that discovering mechanisms is important and useful: only that the directive to do so is not a reliable guide for causal inference in epidemiology. This is not because it does not establish causation, but

⁹ One of the discoverers, Barry Marshall, drank a solution containing the bacteria and developed gastritis, then took antibiotics and recovered.

because it does not identify the *right* causal associations. Identifying causes, *any* causes, is not enough; there may be *other* causes – like *H pylori* – which better explain the phenomena in question, or offer readier foci for intervention. *Mere* causal inference is not all it is cracked up to be.

What makes the *H pylori* case such a neat illustration is that acid *does* play a role in ulceration. The bacteria cause excessive acid production, which is what directly causes ulceration. So there *is* a mechanism there. In the cases of Snow and Semmelweis, on the other hand, the mechanisms identified we would now regard as non-existent. So in those historical cases, a defender of the mechanistic stance might argue that, had the *real* mechanism of disease transmission been believed rather than miasma theory, the hypotheses of Snow and Semmelweis would have been better received. That may well be so. But this argument confuses reality with our grasp on it. A sound methodological principle cannot rely on our already knowing what we are trying to find out. In the case of miasma theory, many medical scientists *thought* they knew the mechanism of disease transmission. Partly as a consequence, they failed to properly appreciate the evidence before them. This is not simply a case of scientists being convinced of something and failing to give due weight to disconfirming evidence: it is a case of scientists believing they understand *how* something happens, and rejecting causal hypotheses that appear incompatible with this mechanism.

This problem seems to be at least partly what Doll and Peto have in mind when they advocate the black box stance. They argue that epidemiological evidence can warrant many causal inferences, without the underlying mechanisms being known. This sounds like a sort of call to arms for epidemiologists, a rallying cry for them to have faith in the methods of their discipline, and in particular to pay attention to the causal associations revealed by the evidence, without worrying about how the causal associations might work. Nevertheless, the black box stance does not offer much in the way of positive methodological recommendations. And it does not entirely dispense with the fundamental motivation of the mechanistic stance: that if we want to *really* explain – or control – something, we need to know how it works. In the next section I will propose a reconciliation.

4. Contrasts

It may be tempting to see the contrast between the mechanistic stance and the black box stance as a disagreement about the goals of epidemiology. One epidemiologist puts it like this:

...the epidemiologist who tries to explain, and if possible eliminate, variations in disease occurrence without much regard for mechanisms, stands in contrast to the laboratory scientist who prefers to disentangle the mechanisms first.

(Vandenbroucke 1988, 708)

The mechanistic stance might be seen as a more properly scientific view, interested in deep explanation; while the black box stance might be thought of as a more pragmatic view, interested primarily in designing public health interventions by the most direct inferential route available.

This is a misunderstanding, in my opinion. The strengths and weaknesses of the mechanistic stance apply equally to the goals of explanation and intervention. Let me enumerate the principle strengths and weaknesses of the mechanistic stance, starting with the strengths.

- (i) Interpreting causal generalisations in epidemiology as existence claims about underlying mechanisms resolves gives a clear and plausible meaning to those generalisations, and helps us to understand the role of replication in epidemiology.
- (ii) How widely we can generalise from a known association seems to be directly linked to how well the underlying mechanism is understood.
- (iii) Discovering an underlying mechanism proves the truth of the causal hypothesis in question.
- (iv) A mechanistic explanation of a causal association increases our understanding of that association.
- (v) Knowledge of the mechanism underlying one causal association gives us, or at least can lead to, knowledge of *other* causal associations.

Each of these is a strength, whether the goal is scientific explanation or public health intervention. (i) is perhaps the most philosophical advantage, but it is surely of some importance to the scientist and the public health policy maker alike to have a good grasp

on the nature of the causal generalisations they employ. For the scientist, it yields greater understanding; for the policy-maker, the ability to avoid practical consequences of misunderstanding. (ii) is evidently of central interest to the policy-maker, since how widely the results of a given study apply is central to the question of what interventions it warrants. It is also of evident interest for someone who is interested primarily in explanation, since it bears on how much a given causal hypothesis might explain. (iii) is of interest for both explanation and intervention, since false causal inferences are neither explanatory nor reliable guides for intervention. (iv) is evidently of value for those interested primarily in explanation, but it is also of use to the intervention-focused. This is because when we understand a mechanism, we are often able to identify more than one point at which we might intervene on that mechanism (if we want to prevent it) or more than one point at which it might be vulnerable to breaking down (if we want to protect it). (v) is useful from an explanatory perspective, since knowledge of other associations is a symptom of explanatory power, as well adding to the grand total of our knowledge. And it is evidently useful from the point of view of intervention, because it enables us to avoid unintended consequences.

A similar exercise shows the weaknesses of the mechanistic stance to apply regardless of whether explanation or intervention is the goal. To summarise the weaknesses:

- (i) A general causal hypothesis can be warranted before the underlying mechanism is discovered (indeed, a warranted causal hypothesis is a great reason to look for a mechanism).
- (ii) The directive to seek mechanisms is non-specific, and does not tell us how to find mechanisms; on some occasions, this appears to have led to a bias in favour of known mechanisms and against causal associations for which no mechanism is yet known.
- (iii) A causal hypothesis may be deficient even though the underlying mechanism is understood, because it may lead us to believe that we have obtained a greater understanding of a phenomenon than we really have.

From an intervention-oriented point of view, (i) is a serious drawback of the mechanistic stance, because it shows that the mechanistic stance could lead to serious unnecessary delays on intervention. Less obviously, it is also a drawback from the explanatory point of view, because it leads to the unnecessary rejection of good explanations. From an explanatory point of view, the mechanistic stance is guilty of a why-regress fallacy. The

“why regress” is simply the fact that it is always possible to ask “Why?”, including on occasions when the explanation offered is a good one (Lipton 2004, 21-2). What I call a why-regress fallacy is the refusal to accept an explanation on the grounds that the explanation itself has not been explained. As a general rule, such grounds are fallacious, because explanations can be good as far as they go, without providing the entire causal history of the explanandum. Otherwise, epidemiologists would study the Big Bang. (ii) is clearly a problem for any application of the mechanistic stance on causal inference, regardless of motivation. (iii) applies to intervention- and explanation-oriented foci equally, since it means that we miss out on potentially more fruitful interventions or explanations respectively.

What, then, is the diagnosis of the tension between mechanistic and black box stances? I suggest it arises from a simple confusion of metaphysics and methodology. On the one hand, it does not follow from the strengths listed that discovering mechanisms is necessary for causal inference, nor that the discovery of mechanisms is a sufficient guide for explanation or intervention. None of these points directly motivate the mechanistic methodological stance on causal inference; what they motivate is *interpreting* causal hypotheses as existence claims about underlying mechanisms, and *seeking* these underlying mechanisms. Neither directive tells us anything about the *method* of causal inference; at most, they tell us about the goal. On the other hand, it does not follow from the weaknesses listed that we should abandon a mechanistic interpretation of causal hypotheses, or that we should abandon the search for mechanisms. It only follows that we should not set the identification of a mechanism as a necessary condition for causal inference; nor confuse the power of a mechanistic explanation for a given causal association with the explanatory power of the association itself, with respect to the goals of our investigative activities.

The advantages of a mechanistic interpretation of causal hypotheses in epidemiology were laid out in §2. My suggestion is that we endorse this interpretation, but resist the temptation to take the mechanistic methodological stance with respect to causal inferences. The mechanistic metaphysics is good, but its methodology is bad.

The obvious next question is: how should causation be inferred? I do not think there is an easy answer. There are just a bundle of methods, organised around a common theme of identifying patterns of differences and similarities (Mill 1843; Lipton 2004), but increasingly statistically sophisticated (Spirtes, Glymour, and Scheines 1993; Pearl 2000),

and (it is to be hoped) increasingly reliable. The point of this paper is not to contribute a new method of causal inference, but to identify and debunk a tempting bad method. That method is what I called the mechanistic stance. But at the same time, I hope to have shown that a mechanistic metaphysics for causal generalisations has a great deal to offer epidemiology. There is, therefore, no need to *choose* between the identification of causal associations at the population level, and the identification of underlying mechanisms. On the mechanistic interpretation, causal hypotheses at the population level *are* existence claims about underlying mechanisms. There need be no opposition between epidemiologists conducting observational studies, and those trying to “disentangle” mechanisms in a laboratory. They are studying the same thing.

It is not necessary to identify the underlying mechanism in order to have *warrant* for a causal hypothesis. But it is necessary in order to *explain* the association. A hypothesis can be perfectly warranted, without being understood. Conversely, a hypothesis may be well (mechanistically) understood, but may itself fail to provide a good explanation of the phenomenon in question for the purposes at hand. To prove a causal hypothesis, it is sufficient to identify an underlying mechanism. But identifying a mechanism is no guarantee of the explanatory power of the hypothesis itself – the explanatory power of the hypothesis that stomach acid causes ulcer, for example. Mechanism is sufficient for causation; but causation is not sufficient for explanation, or for purposes of intervention. Oxygen is a cause of every car crash, and I make that assertion not on the basis of any statistical information, but entirely on the basis of my knowledge of the mechanisms underlying internal combustion engines and human respiration. Yet oxygen offers good explanations of few, if any, car crashes; and controlling the oxygen supply is not a particularly promising avenue for policy makers to pursue.¹⁰

What I have not done is offer any detailed analysis of the notion of a mechanism. The goal of this paper is to see what the notion can offer causal inference in epidemiology, not to analyse that notion itself. Nevertheless, I would like to finish by sketching a view

¹⁰ Sometimes a distinction is drawn between distal and proximate causes of disease; and sometimes biological causes are also distinguished, which may be either distal or proximate. However, drawing and defending such distinctions is not easy, as a large literature in philosophy and also in jurisprudence shows (for references see, respectively: Broadbent 2008; Broadbent 2009a). A discussion of these distinctions would not be relevant here, since there is no particular reason to hope that the most explanatory cause in a given circumstance will be either proximate, or distal, or biological. Moreover, even if there were some reason for epidemiologists to favour one of these categories, a choice would still have to be made among causes within them.

of the relation between mechanistic explanation and other kinds of causal explanation. I am inclined to see mechanistic explanation as of a kind with causal explanation in general. I take it as widely accepted that causation is not sufficient for explanation, and that for a causal association to explain, it must amount to a difference between fact and foil (Lewis 1986; Lipton 2004 Ch 3). In a public health context, health provides a plausible contrast class, as I have argued elsewhere (Broadbent 2009b). To give a mechanistic explanation of an association is to tell a story about how events of one type cause events of another, by filling in the intervening steps in the causal chain, and specifying what conditions must hold. It is tempting to see mechanistic explanations as causal explanations, with associations as their explananda. Presumably the relevant contrast class is the failure of that association to hold. Identifying possible failures is not purely a hypothetical exercise: typically there will be plenty of actual failures, since epidemiological generalisations are typically exception-ridden.

If this sketch of a contrastive analysis is approximately correct, it would explain why mechanistic explanations are so useful in epidemiology and public health: because they provide information about what makes the difference between cause-events leading to effect-events, and not doing so. And intervening on links of this sort is the central purpose of public health policy. However, this sketch also illustrates why mechanistic explanation is not necessary for good causal inference: because a causal difference between fact and foil can be identified, even before the causal link between the proposed explanans and the fact is understood. We can know *that* drinking dirty water is the cause of the difference between people with cholera and those without, even if we don't know how drinking dirty water results in cholera. And the contrastive approach shows why mechanistic explanation is not sufficient for effective intervention or explanation. If a causal hypothesis fails to identify a causal difference between good health and ill, then it will offer neither explanation of nor intervention on that difference; and a mechanistic explanation of the causal hypothesis (e.g. miasma theory) cannot improve matters.

The correct analysis of mechanistic explanation is not, however, the purpose of this paper. To apply the moral reflexively, we can know that mechanistic explanation works, without knowing how it works. I hope to have shown the strengths of a mechanistic

interpretation of causal hypotheses in epidemiology, and to have illuminated the pitfalls of a tempting but erroneous companion methodology.¹¹

References

- Angel, Katherine. 2008. *Causality and Psychosomatic Histories in Contemporary Anglo-American Biomedicine*. PhD, University of Cambridge.
- Broadbent, Alex. 2008. The Difference Between Cause and Condition. *Proceedings of the Aristotelian Society* 108: 355-364.
- . 2009a. Fact and law in the causal inquiry. *Legal Theory* 15: 173-191.
- . 2009b. Causation and models of disease in epidemiology. *Studies in History and Philosophy of Biological and Biomedical Sciences* 40: 302-311.
- Brunner, Eric, and Michael Marmot. 2006. Social organization, stress, and health. In *Social Determinants of Health*, ed. Michael Marmot and Richard Wilkinson, 6-30. Second. Oxford: Oxford University Press.
- Cameron, D., and J.G. Jones. 1983. John Snow, the Broad Street pump and modern epidemiology. *International Journal of Epidemiology* 12: 393-396.
- Carter, K. Codell. 1994. *Childbed Fever: A Scientific Biography of Ignaz Semmelweis*. Greenwood Press.
- . 2003. *The Rise of Causal Concepts of Disease*. Aldershot: Ashgate.
- Deveraux, P.J. 2008. Effects of extended release metoprolol succinate in patients undergoing non-cardiac surgery (POISE trial): a randomised controlled trial. *The Lancet*. doi:10.1016/S0140-6736(08)60601-7. <http://www.thelancet.com/>.
- Doll, R, and R Peto. 1981. *The causes of cancer*. Oxford: Oxford University Press.
- Farooqi, Sadaf, and Stephen O'Rahilly. 2006. Genetics of Obesity in Humans. *Endocrine Reviews* 27, no. 7: 710-718.
- Gillies, Donald. 2005. Hempelian and Kuhnian approaches in the philosophy of medicine: the Semmelweis case. *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 159-181.
- Glennan, Stuart. 2002. Rethinking Mechanistic Explanation. *Philosophy of Science* 69: S342-S353.
- Hill, Austin Bradford. 1965. The environment and disease: association or causation? *Proceedings of the Royal Society of Medicine* 58: 259-300.
- Lewis, David. 1973. Causation. *Journal of Philosophy* 70: 556-567.
- . 1986. Causal Explanation. In *Philosophical Papers, Volume II*, 214-241. Oxford: Oxford University Press.
- Lewontin, Richard. 1974. The analysis of variance and the analysis of causes. *American Journal of Human Genetics* 26: 400-411.
- Lipton, Peter. 2004. *Inference to the Best Explanation*. Second. London and New York: Routledge.
- Marmot, Michael. 2006. Health in an unequal world: social circumstances, biology, and disease. *Clinical Medicine* 6, no. 6: 559-572.
- Mill, John Stuart. 1843. *A System of Logic, Ratiocinative and Inductive*. 8th ed. New York and Bombay: Longman's, Green, and Co.
- Pearl, Judea. 2000. *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press.

¹¹ I am grateful to Kevin Brosnan, Brendan Clarke, Donald Gillies, Stuart Glennan, Jeremy Howick, the audience at a conference on Causality and Mechanisms held in Kent 2009, and an anonymous referee, for their useful comments and criticisms.

- Peto, Richard. 1984. The need for ignorance in oncer research. In *The Encycolpedia of Medical Ignorance*, ed. R Duncan and M Weston-Sraim, 129-133. Oxford: Perganon Press.
- Rothman, Kenneth J, and Sander Greenland. 2005. Causation and Causal Inference in Epidemiology. *American Journal of Public Health (Supplement 1)* 95, no. 1: S144-S150.
- Russo, Federica, and Jon Williamson. 2007. Interpreting Causality in the Health Sciences. *International Journal of the Philosophy of Science* 21, no. 2: 157-170.
- Rutter, Michael. 2007. *Identifying the environmental causes of disease: how should we decide what to believe and when to take action?* The Academy of Medical Sciences.
- Schonemann, Peter H. 1997. On models and muddles of heritability. *Genetica* 99: 97-108.
- Sesardic, Neven. 2005. *Making Sense of Heritability*. Cambridge: Cambridge University Press.
- Spirtes, P., C. Glymour, and R. Scheines. 1993. *Causation, Prediction, and Search*. New York, N.Y.: Springer-Verlag.
- Vandenbroucke, J.P. 1988. Is 'The Causes of Cancer' a Miasma Theory for the End of the Twentieth Century? *International Journal of Epidemiology* 17, no. 4: 708-709.
- Wardle, Jane, Susan Carnell, Claire M.A. Howarth, and Robert Plomin. 2008. Evidence for a strong genetic influence on childhood adiposity despite the force of the obesogenic environment. *American Journal of Clinical Nutrition* 87: 398-404.