



Author's Reply

Formalism or pluralism? A reply to commentaries on 'Causality and causal inference in epidemiology'

Alex Broadbent,^{1,*} Jan P. Vandenbroucke² and Neil Pearce³

¹Department of Philosophy, University of Johannesburg, Johannesburg, South Africa, ²Leiden University Medical Center, Department of Clinical Epidemiology, Leiden, The Netherlands and Department of Clinical Epidemiology, Aarhus University Hospital, Aarhus, Denmark and ³Department of Medical Statistics and Centre for Global NCDs, London School of Hygiene and Tropical Medicine, London, UK and Centre for Public Health Research, Massey University, Wellington, New Zealand

*Corresponding author. Department of Philosophy, University of Johannesburg, PO Box 524, Auckland Park 2006, Johannesburg, South Africa. E-mail: abbroadbent@uj.ac.za

Accepted 13 September 2016

Introduction

We thank the commentators on our paper.^{1–5} We have learned a great deal from this exchange. These papers do much to clarify previous statements, and show that there is much that we agree on. They also show that there remain important points of disagreement. In this response we do not expect that we can resolve all these points of disagreement; rather, our intention is to clarify them and our position in relation to them.

We approach these commentaries by identifying what we take to be an authoritative version of the Potential Outcomes Approach (POA), which we believe that most of our interlocutors would accept, given what they have written. In addition, we define the Restricted Potential Outcomes Approach (RPOA), which (as we have indicated from the outset) is characterized by a further commitment concerning the role of interventions (whether or not humanly feasible). We have concerns about certain aspects of the POA, but we regard the RPOA as more seriously mistaken.

Our approach in what follows is to identify, as clearly as we can, our points of agreement and disagreement on the substantive elements of the POA and RPOA. Having dealt with these 'hard' theoretical issues, we seek to

relate these to the various 'soft' downstream issues relating to pedagogy, emphasis, prioritization of study designs and neglect of important exposures. These 'soft' aspects of the debate are by no means less important—in fact they may even be more important for the characteristics that epidemiology takes on as it develops. But if clarity on the substantive points of agreement and disagreement can be achieved before discussing the 'soft' downstream issues, then those discussions will be more fair and fruitful.

The version of the POA that we believe emerges from the exchange has the following elements.

- (i) Counterfactual dependence of E on C is not necessary for C to cause E, but it is sufficient (POA's Basic Metaphysical Stance).
- (ii) Sufficient evidential conditions currently exist for attributing the counterfactual dependence of E on C, but necessary conditions currently do not; the POA identifies some (but not all) of these sufficient conditions (POA's Basic Epistemological Stance).
- (iii) Causal inference includes two distinct aspects: causal identification, in which the truth value of a claim of the form 'C causes E' is determined; and quantitative causal estimation, in which a numerical value n is

estimated for a claim of the form ‘C has n effect on E’ (the Identification/Estimation Distinction).

- (iv) Adequately well-defined counterfactual contrasts are necessary for giving meaning to quantitative estimates of causal effect (POA’s Semantic Stance on Estimation).

These four elements characterize what we regard as a ‘standard’ version of the POA. Different authors emphasize different elements of these in their responses to our original paper and in their writings preceding this discussion. However, we believe that most of our interlocutors accept these elements. In his considered and especially clear response,⁴ Tyler VanderWeele touches on almost all of these, and is particularly at pains to emphasize that counterfactual dependence may be sufficient for causation even if it is not necessary, part of (i), and that there is an important difference between causal identification and quantitative estimation, as stated in point (iii). Hernán has emphasized (iv) in particular, as well as point (v) below.⁶

Perhaps because of the POA’s resonance with the theoretical basis of randomized controlled trials, some have taken it implicitly to involve interventions. This is not a logically essential element of PO theory (as we pointed out in our original paper), and it is certainly not inherent in the counterfactual approach (as Pearl and others make clear⁷). This gives rise to an additional fifth element which, taken with (i)-(iv), defines the Restricted Potential Outcomes Approach (RPOA).

- (v) Counterfactual contrasts are adequately well-defined if and only if we can specify a corresponding adequately well-defined intervention on the putative cause, by which the counterfactual contrast would be (or would have been) brought about. (RPOA’s Restriction to Interventions).

This additional clause imposes both a necessary and a sufficient condition on adequately well-defined counterfactual contrasts. Hernán is particularly closely associated with (v), having made statements to the effect that causal questions are well-defined when interventions are well-defined,⁶ and having critiqued the work of epidemiologists who do not define interventions.⁸ However, because the difference between (i)-(iv) and (i)-(v) is not generally appreciated, all of our interlocutors have at some stage assumed elements (i)-(v): that is, they have assumed or committed themselves to the RPOA approach, as we show in Section B1 below.

It is the additional element involving interventions that has given rise to the most debate, since ‘intervention’ has not been defined. In particular, there has been discussion about whether interventions must be humanly feasible events or not. However, there are other problems with (v),

even setting aside the matter of human feasibility (that is, even accepting the recent insistence that ‘intervention’ does not imply ‘humanly feasible’). We were clear on this in our original paper and we reiterate this below.

In Part A of this paper, we examine each of (i)-(v). Roughly speaking, our views are as follows.

- We accept (i).
- We reject (ii) because we do not believe that there are logically sufficient circumstances for inferring causation. Causal inference is always inductive, not deductive; unless ‘sufficient’ means ‘inductively sufficient’, in which case we accept (ii) but regard it as misleading because we were capable of inductive causal inference long before the POA, or indeed science in general, were conceived.
- We accept (iii) for the sake of the current discussion, i.e. to be able to explore its consequences.
- We accept (iv) provided that ‘well-defined contrasts’ are understood in a suitably liberal way, and provided that it is not confused with the incorrect claim that well-defined contrasts are sufficient for causal estimation;
- We reject (v) as a mistake. ‘Intervention’ is a misleading term, suggesting human feasibility when this is apparently not intended, and still without a precise definition in the present discussion. Moreover, there is no good reason to focus on interventions (even non-human ones) when seeking to specify counterfactual contrasts. As indicated in our original paper: (a) interventions are not sufficient for well-defined counterfactual contrasts; (b) we may not always know in advance whether a contemplated intervention is well-specified; and (c) interventions are not necessary for well-defined counterfactual contrasts.

We focused on the RPOA in our original paper because we regard (v) as the most problematic. It is striking that, in their responses, our critical commentators have only responded to our discussion of interventions in relation to human feasibility. Having clarified that they never intended ‘intervention’ to be restricted to humanly feasible actions, they proceed to other matters. But we were clear in our original paper⁹ that there are problems with the role of interventions in the RPOA even after the assumption of human feasibility is dropped. We suspect that the difference between (i)-(iv) and (i)-(v) has simply not been fully appreciated and that some of our interlocutors slide unconsciously from one to the other, especially when pressed (for example, when asked for a definition of ‘intervention’). One of our hopes in this response is to bring this distinction out more clearly, enabling us to reconcile our relatively benign attitude to the POA (even though we have serious concerns about some aspects) with our rejection of

the RPOA which we think could restrict and damage epidemiology as a science.

In Part B of the paper we turn to more general considerations. In B1 we consider the Straw Man objection, which we have encountered frequently. At risk of seeming belligerent, we quote earlier passages in which VanderWeele and Hernán commit to the RPOA and interpret ‘intervention’ as ‘humanly feasible intervention’. We hesitate to do this, because ‘who said what and when’ is not strictly relevant to the substantive matters at issue. However, one of our central concerns with the RPOA concerns lack of clarity: there has been a failure to define key terms (specifically, ‘intervention’), and a failure to distinguish between very different assertions (between POA, RPOA, humanly feasible RPOA). Thus what has been said in recent letters and commentaries is not equivalent to what was said before we published our original paper. We believe that perceiving these differences is important for the conceptual development of the science.

In Section B2 we argue that science is a process and we reiterate our general view that the POA does not fit easily into the view of the nature of science that we favour, and from which we believe epidemiology as a whole has benefited and will continue to benefit.

A. Substantive issues

A1. The POA’s Basic Metaphysical Stance

We regard the POA’s Basic Metaphysical Stance (i) as reasonable. Regarding (i), epidemiologists are not concerned with metaphysics except insofar as they are forced to take a metaphysical stance in order to carry on with their work. We take this to be agreed on all sides in the present debate. The most obvious way for epidemiologists to proceed, given that metaphysics is not their day job, is to follow the broadly accepted consensus among the academic community whose day job is the metaphysics of causation, namely philosophers. Only if that consensus seems unhelpful for epidemiological purposes need epidemiologists disagree with the philosophical consensus.

The philosophical consensus agrees with (i) in saying that effects do not always depend on their causes,¹⁰ as when the president’s death does not depend on the assassin’s shot because there was a back-up assassin who would have killed the president if the actual killer had missed (or, had she not smoked, a smoker may have got lung cancer for other reasons). And the philosophical consensus agrees with Tyler VanderWeele in particular that if X and Y occur, and Y counterfactually depends on X, then X causes Y¹⁰ (even though VanderWeele’s use of laws in expressing this thought differs from the consensus view of the role of

laws in determining the truth conditions of counterfactuals^{11–14}). We too accept (i), for these purposes.

A2. POA’s Basic Epistemological Stance

VanderWeele identifies, as a key point, that ‘there are well established sufficient conditions for attributing causation and the potential outcomes framework provides one such set of sufficient conditions’.⁴ We understand and unpack the claim as follows. We can have certain kinds of evidence which are sufficient for us to attribute causation; and the POA represents or supplies one such kind (or set of kinds). Specifically, the POA provides sufficient conditions to attribute counterfactual dependencies, and by (i) evidence for a counterfactual dependency is likewise evidence for causation. We suppose that the route must be from evidence via counterfactual dependence to causation, and not directly from evidence to causation, because it is characteristic of the POA that it does not concern inferential routes from evidence to causation that do not specify anything about counterfactual dependence (although it need not deny the existence of such routes, as some of the commentaries have stressed, notably VanderWeele’s¹⁵). We take this to be accepted on all sides.

The view represented by (ii) is expressed clearly by VanderWeele when he says that there are well-established sufficient conditions for attributing causation. It seems that this view is widely shared among those impressed by the POA. We suspect that there may be frustration at our apparent reluctance to accept it, which may even be construed in some quarters as methodological Ludditism, or as the last gasp of an outmoded way of thinking. Nonetheless we do not accept the view expressed in VanderWeele’s remark and in (ii) above.

The meaning of VanderWeele’s point depends on the notion of ‘sufficient’ in play. One interpretation might be that logical sufficiency is intended. If that is the case, we deny that logically sufficient conditions exist either for the attribution of causation or for the attribution of counterfactual dependence, because all such attributions are the result of an inductive inference, either directly or indirectly.

Nancy Cartwright has been at great pains to warn proponents of evidence-based policy of the dangers of forgetting that a conclusion is only as well supported as the assumptions under which the method of reaching it operates.^{16,17} The method itself may be deductive, given the assumptions, and in this sense the causal conclusion is deduced from the evidence—but the assumptions underpinning the method are themselves inductively supported. These assumptions have very broad scope, and include the integration of all knowledge from the broader set of

evidence that is necessary to set up different studies, the assumption that the data are accurately recorded in the first place, assumptions about honesty and sobriety of investigators, and many more besides.

One well-known illustration of this point is the much greater probability of finding a positive effect (or finding a larger effect) in randomized controlled trials where the investigators have a conflict of interest, either because the trial is funded by pharmaceutical companies or for other reasons.¹⁸ This may often be the result of tweaking the ‘inputs’ of the trial rather than of outright fraud; this tweaking is not fraudulent precisely because the data underdetermine the very many decisions that trial designers have to make (meaning that more than one non-equivalent set of these decisions is compatible with the data).¹⁹ These decisions include choice of participants, measurement instruments (such as questionnaires), time periods and many other factors. Less transparently, scientists make multiple humdrum decisions on a daily basis about what techniques to employ, what ‘leads’ to follow, when to look harder and when to stop; and these, too, can influence the outcome.²⁰ So the trial might yield a causal conclusion, via a deductive inference, given a raft of assumptions; but no matter how huge and sophisticated the trial, the security of its conclusion is no better than the security of its assumptions. Furthermore, the possibility that findings are due to chance can never be ignored, even at *P*-values below some level whose basis may reflect a social consensus² among scientists in that discipline (in epidemiology, 0.05; in physics, often 0.001); any such level is ultimately arbitrary from an objective point of view, and lower *P*-values reduce but do not eliminate the possibility of a chance finding.

These familiar points show that randomized controlled trials do not constitute sufficient conditions for good causal inference: the full sufficient condition includes a raft of assumptions, which cannot be satisfied by specifying the design of the trial, no matter how exactly. These points apply equally to any proposed set of sufficient conditions for the analysis of data from observational trials. In particular, they apply to the use of the general approach advocated by the POA and the particular techniques associated with it. We take the point of Daniel *et al.* that being forced to state your assumptions is helpful,² but we are not convinced that the methods of the POA do in fact force epidemiologists to state all the relevant assumptions. The assumption that they suffice to enumerate the assumptions could itself be dangerous, since it obscures the fact that other assumptions remain unspecified.

Nonetheless we acknowledge that making even a subset of our assumptions clear can be very useful. Our problem is not with that exercise, but with the temptation to think

that by stating some of our assumptions more clearly, we have successfully formalized the entire inferential process. Daniel *et al.* say that: ‘Objective science eventually calls for a formal theory and approach. We view the [POA] as precisely offering formal tools to investigate cause-effect relationships.’² They concede that inference to the best explanation (IBE) guides the application of these methods, by which they appear to mean that IBE helps us decide which questions to investigate. However, they believe that intuition breaks down under certain circumstances, and thus: ‘There is no question, in our opinion, that a formal theory is needed to guide data analysis.’² Science may indeed seek objectivity, and for this reason a deductive method for causal inference is indeed highly desirable. But this does not mean that it is possible: we cannot have one just because we decide we need one. Causal conclusions do not follow deductively from data without a strong set of auxiliary assumptions, and (as just discussed) these assumptions are themselves not deductive consequences of the data. A formal method may indeed be extremely helpful, provided that its significance is not misunderstood and its dependence on supporting assumptions not forgotten.

We are also concerned about policy makers who may misunderstand what scientists say. If it is claimed that causal inference has been formalized and it is not explained that the formalism, powerful as it may be, is only as good as the assumptions that support it, then causal conclusions will look surer (‘more objective’) than they really are.

We suspect that VanderWeele and many others would agree that the POA does not amount to a fully logically sufficient set of conditions for inferring causality from empirical data. The second sense in which ‘sufficient’ might be intended is the informal sense, the more casual, normal usage: something like ‘enough, in the circumstances, to warrant a causal inference’, or perhaps ‘sufficient for an inductive inference to a causal conclusion’. If the former kind was logical or deductive sufficiency, we can call this inductive sufficiency. The claim that there are inductively sufficient conditions for causal inference is plausible. However, it is not news. It amounts to saying ‘causal inference is possible for us’, or ‘sometimes, we can make causal inferences’. This is true, and has been true since the Stone Age. It remains a mystery exactly how inductive inference is possible, given our inability either to justify or even to accurately describe it. However, it certainly appears that we have made successful causal inferences on many occasions.

We suspect that the assertion that ‘sometimes, we can make causal inferences’ is much weaker than what POA advocates would intend, and is also weaker than what readers would naturally read into it. We suspect that the intention is to indicate that the POA has contributed a set

of logically sufficient conditions, where before we did not have such a set. We do not accept this. We agree that the POA has contributed to our causal inference abilities, by contributing to the large and diverse set of inductively sufficient conditions for causal inference; but we maintain that it has not contributed logically sufficient conditions, and that the inductively sufficient conditions it has contributed join a large array of pre-existing conditions of this kind. Thus we see a risk that claims of this kind may give a false sense of progress.

The fact that there are inductively sufficient conditions for causal inference certainly has not been made true by the advent of the POA or any other recent developments. Probably what POA advocates would want to say in response is that POA represents a development or an advance in our abilities to make causal inference. We tend to agree and we made a similar claim in our original paper, arguing that the POA represented a set of extremely useful conceptual and methodological tools that may help causal inference in particular circumstances. However, this simply is not equivalent to the stronger (logical) sense of sufficiency, and we suspect that the stronger sense is often in the background of discussions of the POA and supplies some of the associated glamour and excitement, as well as a false sense of certainty.

A3. The identification/estimation distinction

One innovation of this debate has been the introduction, or at least the clarification, of a distinction between two kinds of causal inference. Causal identification occurs when we identify some exposure as among the causes of the outcome. Causal estimation occurs when we seek to quantify the contribution of some exposure to an outcome.

One of us (A.B.) trained with the late Peter Lipton, a philosopher of science famous for turning 'Inference to the Best Explanation' from a slogan into a detailed (partial) theory of inductive inference, and who emphasized the importance of contrasts in causal explanation and hence (via IBE) in causal inference.^{21,22} One of his mottos was: 'When faced with a contradiction, make a distinction.' Our view of the distinction between identification and estimation is that it represents a way for the POA to maintain that counterfactual dependence of certain kinds is necessary for certain kinds of causal inference, even though it is not necessary for causal inference in general. We are uncertain whether it is useful in any other way, besides carving out a domain in which the POA may seek to assert necessary conditions on a certain demarcated subset of causal inferences, namely those relating to the quantitative estimation of causal effects. Below (in section B2) we express

our view that even causal estimation is the result of an inferential process involving the integration of many pieces of evidence from many sources.

A4. The POA's semantic stance on estimation

Characteristic of the POA is an insistence that it is not meaningful, or at least not clear, to make quantitative attributions of causal responsibility without properly specifying the counterfactual contrast against which the contribution of the putative cause to the actual effect is quantified. This is a semantic stance (concerning the meaningfulness of certain claims) but it has epistemological consequences (consequences for what we can know). Meaningless claims cannot be true (or false), and thus cannot be known; nor can they be the conclusions of any inference (since an inference 'to a claim' is to the truth of that claim). Thus (iv) implies that putative causal inferences are mistaken if they attempt to estimate causal effects without adequately specifying counterfactual contrasts. We regard (iv) as a useful and important point, and a real contribution to epidemiological thinking. Perhaps it ought to be obvious, but if one is going to specify how much of a difference something makes, one needs to be clear about the baseline against which the difference is being measured.

However, we do note two points of caution about (iv). First, we suppose that 'adequately defined' is understood in a suitably liberal way, and not as implying the universal use of the set of methods closely associated with the POA/RPOA. Second, necessity and sufficiency must not be mixed up. We remain of the view expressed above (in A2) that there are not logically sufficient conditions for any inductive inference. Adequately defined contrasts are necessary to make sense of estimates of effects of potential causes. However, the fact that we have adequately defined our contrasts does not by itself guarantee that our causal estimate is correct, either in its quantity or in the underlying causal relationship it purports to measure. Background information and knowledge from other sources is indispensable, a point we return to in B2 below.

A5. The RPOA's restriction to interventions

More problematic is the insistence that 'causal effects cannot be defined ... in the absence of well-defined interventions'.⁶ This is the key element of the RPOA, in contrast with the POA, and the one that we believe is most problematic. We doubt the usefulness of the term 'intervention', and, as we emphasize in our authors' reply²³ to a letter by VanderWeele *et al.*,²⁴ the term itself has not been defined by those who rely on it. In epidemiological circles,

‘intervention’ usually denotes a human act of some kind. Therefore if we have misunderstood, this at least partly reflects the vagueness with which the term ‘intervention’ has been used. The commentaries universally argue that ‘intervention’ ought not to be interpreted so as to indicate humanly feasible interventions. Even so clarified, we reject the claim that interventions are necessary or sufficient for the adequate specification of counterfactual contrasts. In our original article we said: ‘The deeper problem with the RPOA concerns its reliance on the notion of a well-specified intervention, whether humanly feasible or not.’⁹ Thus we made clear that our criticisms extended beyond a commitment to human feasibility. None of the respondents has addressed our remaining criticisms of the notion of intervention in general as it appears in the RPOA, focusing instead on our criticisms of humanly feasible interventions. We take this opportunity to seek to further clarify why we think that interventions (human or otherwise) should not have an essential or privileged place in the POA.

In our original paper we identified three difficulties with the role assigned to interventions by the RPOA: (a) merely specifying an intervention is not sufficient for adequately specifying a causal contrast; (b) one may not always be able to determine in advance whether an intervention is sufficiently well-specified; and (c) specifying an intervention is not necessary for adequately specifying a causal contrast.

The first point (a) is illustrated by the interventions in Hernán and Taubman’s paper.⁶ ‘One hour of strenuous exercise per day’ includes many different forms of exercise which may have very different effects on mortality, showing that specifying an intervention does not necessarily get you all the way to a well-specified contrast. The intervention itself needs to be adequately specified; by introducing the notion of intervention, the RPOA has just pushed the problem of defining ‘adequate specification’ back a step, and not solved it.

The second point (b) is essentially a conditional: if specifying interventions were necessary, we would be hamstrung in many cases, because we would be unable to say whether the study we were about to embark on satisfied this criterion or not without actually doing the study. In advance of a study, we may not know whether different forms of exercise affect mortality differently. More plausibly, we may not know (or may learn more about) which properties of the exercise matter (intensity, duration, dominant energy system, time of day, etc.).

The third point (c) is that specifying interventions is, in fact, not necessary for specifying counterfactual contrasts. Specifying an intervention to counterfactually alter a causal variable is not the same thing as specifying a different value for that variable. Specifying an intervention to

bring about a counterfactual value of the causal variable is not necessary for specifying a counterfactual value of the causal variable. This is clear even within the POA. As Bollen and Pearl note, ‘The essential ingredient of causation is responsiveness, namely the capacity to respond to variations in other variables, regardless of how those variations came about.’⁷

We stand by our three original objections to the RPOA, which have not been addressed in any of the commentaries. We suspect that this may partly be because of a tendency to mix up three distinct things: the causes of an exposure of interest; the effect of the exposure on the outcome of interest; and the effects on the outcome of interest of interventions on the exposure. For example, obesity: may have many causes, such as genetic factors, diet, exercise habits, sleep patterns, etc; may have effects on mortality; and may respond to different potential interventions which may have different effects on mortality, both through their differing effects on obesity and their different direct effects on mortality. None of this means that obesity is a ‘composite exposure’ as VanderWeele suggests.⁴ Rather it is a single cause of mortality which can be caused by multiple factors and can be reduced by multiple interventions. Of course, it could be argued that different types of obesity carry different mortality risks, and that we should study more specific subtypes of obesity, but this applies to most exposures. We are always estimating average population effects, for example the average effect on mortality of a BMI of 35 vs 25, or the average effect on mortality of smoking 10 pack-years vs 0 pack-years. This does not depend on specifying interventions.

Moreover, even if we are seeking to intervene on obesity to reduce mortality, the estimates we can obtain of the effects of obesity itself remain useful and provide upper limits to what can theoretically be achieved by public health interventions.^{2,5} Thus, estimating the overall causal effect of obesity is important in itself; furthermore, this is usually the only relevant effect that can be estimated in epidemiological studies—information on specific hypothetical interventions is usually not available. So we can estimate (or attempt to estimate) the causal effect(s) of obesity, and we may also wish to estimate (or attempt to estimate) the effects of specific interventions to reduce obesity. The former information is usually available in epidemiological studies, whereas the latter is usually not.

Hernán and Taubman mix up these issues in their paper.⁶ They point out that it is difficult to understand what ‘excess mortality attributable to obesity’ means unless one specifies a counterfactual contrast class in which the excess mortality is absent. However, their proposed contrast class is one in which an intervention is introduced on an obese population to reduce its obesity. In fact, the

counterfactual contrast of interest in estimating the effect of obesity is not between a population with too much obesity and one whose distribution of BMI has been manipulated towards greater average leanness by a particular intervention. The latter is the relevant counterfactual contrast for determining the effect of an intervention on obesity.

The relevant counterfactual contrast for determining the effect of obesity is better approximated by the general question ‘What if there never was obesity?’ (or more precisely, ‘What if this group of people had always had a BMI of 25 rather than 35?’), in the same way that the effect of smoking is better approximated by ‘What if people never smoked?’ (and not ‘What if they gave up?’—that helps us estimate the effect of giving up, not the effect of smoking itself).

In fact, for all estimations about factors that we believe to be causal, it is a truism that it might be unforeseeable how a real life intervention might play out. Some drug interventions for obesity have led to suicidal ideations and suicide. The original interventions on hypercholesterolaemia using diet had very little effect—if any—but epidemiologists continued to believe that hypercholesterolaemia was a cause, worthwhile to intervene upon; for that same reason, the pharmaceutical industry sought for means to influence it, which led to the advent of statins. Thus, what type of causal assessment is needed will depend on the state of the question: in the beginning it is the likely effect of removing a cause (by unspecified means); later, specific interventions might be considered.

Of course, all of the above is highly contextual. For some counterfactual contrasts one expects very little confounding, such as genetic differences in populations where the genetic differences can be regarded as being random, or when studying unexpected and unpredictable adverse effects of medications. Other counterfactual contrasts may always need intervention studies (or a search for intervention-like situations) because the confounding of the counterfactual contrast is believed to be intractable, such as confounding by indication in drug treatment, being a vegetarian or religious affiliation. Most counterfactual contrasts will be in between these extremes.²⁶ Finally, there is a distinction between cure and prevention: the cure of an acute myocardial infarction is (among other measures) stenting, which aims at the last part of the causal chain and has little to do with the original population causes; prevention of a second myocardial infarction, however, aims at causes ‘upstream’ in that chain.

Finally, the significance of mixing up the effect of an exposure with the effect of an intervention on the exposure is particularly apparent when one considers politically important exposures like race and gender. The effect of being a

woman on income is not well approximated by the effect of undergoing a gender reassignment process. Far from clarifying the study of such exposures, the notion of an intervention actually introduces an opportunity for further confusion.

B. ‘Soft’ issues

Having set out our stance on the substantial matters arising in the course of the debate, we turn now to various ‘soft’ issues which, as stated before, may actually be more important in the development of the science.

B1. Straw man or bendy man?

A repeated criticism from the various commentators is that we have created a ‘straw man’. We have supposedly set up an inaccurate caricature of the RPOA approach, and then criticized this caricature. We address this point reluctantly, because it is hard to do so without sounding argumentative, but we feel we must address it. We are left with the feeling that every time we attempt to address some of the more extreme views of the RPOA, we are told that remarks have been taken out of context or misunderstood (by us). One of the critical commentaries is structured largely around our ‘misconceptions’, as if the whole thing were an unfortunate misunderstanding.² In fact, the quotes that we have taken from Hernán, VanderWeele and others are not cherry picked or taken out of context, but have counterparts in many papers on these issues. We revisit some of them here just to show, once again, that our interlocutors have committed themselves to the RPOA and moreover have often assumed ‘intervention’ as implying human feasibility, despite their different point of view in the commentaries. Rather than a straw man, we are dealing with a ‘bendy man’.

VanderWeele has endorsed the RPOA, giving a central role to interventions, for example:

There has thus been considerable debate as to what, if anything, is meant by the effects of race. The formal causal inference literature has generally conceived of causal effects as a comparison between counterfactuals or potential outcomes. Often in the causal inference literature, the position is taken that it is meaningful to speak of a contrast of counterfactual outcomes only to the extent that we can specify an intervention.¹⁵

This is an endorsement of (v) above, that is, of the RPOA.

VanderWeele has also endorsed an interpretation of ‘intervention’ that is restricted to humanly feasible interventions, in this passage (which follows directly from the previous):

Sometimes this position is associated with the slogan ‘no causation without manipulation.’ A literature has

begun to develop considering this issue of ill-defined ‘treatment’ or nonmanipulable exposures in more detail. However, race is not something we can intervene on, and the associated counterfactual queries generally strike researchers as meaningless. The question of what would a black person’s health outcome have been had they been white seems like a strange one to pose. It is sometimes cautioned that one should not discuss the effects of race except in very special circumstances when such effects do correspond to a manipulable variable such as in the examples above of job application studies.¹⁵

Assuming that we are human, the phrase ‘race is not something we can intervene on’ indicates that ‘intervene’ is being used to imply humanly feasibility. The passage (and the paper as a whole) makes no mention of the idea that an intervention might be humanly impossible. Moreover, it is not possible to make sense of this passage if ‘intervention’ is not supposed to imply human feasibility. This is because the human impossibility of intervening on race is used as the motivation for the paper’s project of finding mediating variables for the effects of race. Drop that implication, and the entire paper loses its stated motivation. (Others might be supplied *post hoc*, but the one that was actually supplied falls away.) The analytical solutions of the paper are imaginative; what we are concerned with is the rationale for proposing them. Elsewhere, humanly feasible interventions are identified as good things to investigate without any particular reason given:

Essentially, we give a plausible causal interpretation of the race coefficient by considering how much a racial inequality could be eliminated by intervening on a different variable, namely socioeconomic status, which may be more manipulable than race.¹⁵

We can think of two reasons for which one might prefer a more manipulable variable. One is that, for practical purposes, one may want to manipulate it. This is benign. But the other possible reason is that one can only meaningfully assign a causal effect to a (humanly) manipulable variable. This is not benign, in our view. The trouble is that the two reasons are not distinguished; indeed, no reason is given for seeking manipulable variables.

This leaves us with unanswered questions. Why are these authors looking for manipulable factors? Who is the intended manipulator—governments, individuals, corporations . . . ? Is this just a pragmatic preference? Is there anything wrong conceptually with just attributing effects to race? Why do we need to tie ourselves into knots trying to decide whether a particular counterfactual contrast can be conceived in terms of interventions? This line of thought leads to strange differentiations. For example, sex can be

considered as a cause (because it is randomized at conception), and genetics may or may not be considered a cause, whereas race/ethnicity cannot be (or at least its causal effect cannot be estimated) although the effects of the mediators of this non-cause can be estimated.¹⁵ These differentiations are bound to seem strange unless more is said about the practical reason for focusing on manipulable variables, if there is one; and from the theoretical point of view of achieving causal understanding, the differentiations appear arbitrary.

It is much more straightforward to assume that causes include (but are not restricted to) any well-defined counterfactual contrast, and that we can attempt to estimate the causal effect of any such contrast (while attempting to control for confounding etc.). It is then straightforward to consider the subgroup of causes for which interventions can be defined, if one so wishes. VanderWeele concedes this point in his commentary, albeit with different terminology from ours, writing: ‘The potential outcomes approach . . . is concerned with a subset of causal questions that can be defined as a contrast of hypothetical interventions.’⁴ This is the first time, to our knowledge, that this point has been explicitly conceded in the various discussions of this topic.

We have also encountered the sentiment that there are no substantial disagreements between our views and those of R/POA adherents, once things are clarified: that it is all a ‘storm in a teacup’. We believe that it is now even clearer what the disagreements are. In particular, we reject the claim that it is necessary to specify an intervention (human or otherwise) when seeking to estimate a causal effect. Hernán, on the other hand, sees the specifying of interventions as obligatory for calculating causal effects: ‘Causal effects cannot be defined, much less computed, in the absence of well defined interventions.’⁶ To us, the word ‘defined’ appears to rule out all causal talk, not just the estimation (or computation) of quantified causal effects. But even if this sentence is understood as limited to quantitative estimate of causal effect (a limitation to our knowledge only explicitly introduced in the commentaries on our original paper), we do not accept this claim. We may accept an adequate specification of counterfactual contrasts as necessary for meaningful quantitative estimates, but we do not regard the specification of interventions as necessary for specifying counterfactual contrasts.

In places, the specification of an intervention appears to be the same thing as the specification of a counterfactual contrast: ‘A proper definition of causal effect requires well-defined counterfactual outcomes, that is, a widely shared consensus about the relevant interventions.’⁸ Here the two stances are conflated—POA and RPOA. We suspect that this conflation of the two stances is one source of the

feeling we have sometimes encountered that we do not really disagree. We hope that the difference is now clearer.

At other times, there is an appeal to usefulness. For example:

The crucial question is then this: What is the point of estimating a causal effect that is not well defined? The resulting relative risk estimate will not be helpful to either scientists, who will be unable to relate it to a mechanism, or policy makers, who will be unable to translate it into effective interventions.⁸

We see this appeal to usefulness as a different point, one that we addressed in our original paper.⁹ It has not been commented on subsequently, so we assume that what we say about it has been accepted—that causal knowledge can be useful regardless of whether it concerns interventions. Nonetheless we have reiterated in the present paper that causal knowledge without specified interventions can be useful: we provide several examples in our original paper,⁹ and we reiterate this position in A5 above.

So we disagree that we have set up a straw man, and we also disagree that we agree with our interlocutors on all main points, as these are expressed in previous writings. To our knowledge, the commentary by VanderWeele⁴ and the letter by VanderWeele *et al.*²⁴ are the first occasions on which it has been explicitly conceded that the R/POA approach does not address all causes of disease, and does not constitute a general and complete theory of causal inference. We agree with these comments and are grateful for the clarification.

B2. Science as a process

There may be differences of perspective on causal inference and on science itself underlying the specific differences described above. We believe that causality is not a statistical property, but a theoretical entity. Saying that a certain relationship is causal amounts to asserting a scientific theory, even if it is a simple, localized theory. As for any scientific theory, we believe that the scientific discovery of causality is not a single-study phenomenon, and that epidemiological research is a process whereby a variety of types of evidence are gathered.

We appreciate that all the main contributors to this debate have now explicitly said that they accept that a variety of types of evidence from multiple studies is important. But the points about sufficiency of circumstances for causal inference (raised above in A2) sit uneasily with this avowal. Moreover, the actual methodological focus of the R/POA is on inference within a single study, or sets of studies.

Methods for reconciling different kinds of evidence are rarely, if ever, mentioned.

The process of discovering a cause of disease unfolds differently on each occasion, but there are some commonalities. One pivotal way in which epidemiologists have been able to generate ideas has been through their population focus.^{27,28} Because populations are particular, time-bound, local entities and not abstract theoretical entities, this means that epidemiological discoveries often have a particularly local, unrepeatable and sometimes even coincidental character. It is the interaction between studies and ideas at the population, individual and molecular levels which so often produces an ongoing cascade of hypothesis generation and testing. Ultimately, this may result in the formulation of highly specific hypotheses that are suitable for testing in a trial (e.g. hepatitis B virus is a cause of liver cancer), or for the estimation of a precisely quantified effect. But this usually occurs at the end of a long process with many false starts and blind alleys.

Furthermore, we would note that generalizability should also be considered as a matter of scientific rather than statistical inference,^{16,17,29} and that estimates from a particular population may not apply to other populations because of different distributions of effect modifiers (and there will almost always be effect modification of the relative risk of risk difference or both,³⁰ or the presence or absence of co-factors. Furthermore, the ‘effect estimates’ generated in epidemiological studies do not even provide estimates of the probability of causation in the populations under study,³¹ let alone in other populations. So the most one can say is that if an exposure causes an increased risk then one might expect it to also cause increased risks in some, but not all, other populations. Thus even the estimation of causal effects is not an exact quantitative science.

Thus we suggest that it is good practice to refrain from calling any individual study’s estimate ‘causal’ even if it is a randomized trial. It is the totality of the evidence that leads to the verdict of causality. Causality is a scientific conclusion, a theoretical claim, and as such transcends any individual study. Estimations either of counterfactual contrasts or of interventions are interesting and important, but are often local effects in a particular time, place and population. And even these are not pure empirical findings, but are heavily theory-laden. They are not read or calculated from data, but inferred from it, and the inference depends upon a huge network of background hypotheses and scientific knowledge—even in the case where an estimate arises from a single study. Thus, causality is not a statistical concept whose presence or absence can be determined by statistical analysis of a set of data. It is a theoretical concept, even when invoked in quantitative estimates for particular

populations. As with any scientific theoretical finding, we infer causal conclusions (including estimations of causal effect) as the result of an inductive inference, considering all the available evidence.

Of course, a single study can be decisive in some circumstances but not because of its innate power to ‘prove’ causation: the context plays a crucial role in determining the significance of evidence. As we have emphasized previously,⁹ a piece of evidence might be very weak taken on its own, yet still be the keystone in a larger argument for causality. While our interlocutors have acknowledged some of these points, we nonetheless maintain that there is a fundamental difference in emphasis between this approach to causal inference and that of the R/POA.

Conclusion

So where does this leave us? Is this just a storm in a teacup? In our view it is not, because it reflects two very different visions of what epidemiology is, how it should be taught and how it should develop. We believe that epidemiology is a science with a particular subject matter to which we can apply a variety of methods. In the R/POA, we perceive a view of epidemiology as a set of methods which are then applied to those scientific questions which can be answered with these particular methods.

Epidemiology identifies potential causes of disease in populations. It is part of a larger undertaking, together with other medical and social sciences ranging from biochemical laboratory sciences to anthropology, that (unlike many RCTs) attempts to understand and explain the observed associations rather than simply estimate them. Explaining associations means obtaining understanding as to why the associations occur. This can and often does lead to understanding why they may not occur (or may have different strengths) in other populations. We see quantitative causal estimation as an important part of epidemiology, but it is just one aspect. We understand the appeal of the elegant mathematical tools and methods of the R/POA, but we do not believe that this is what epidemiology has been restricted to, nor what it is, nor what it ought to be.

Our aim has been to ‘allow’ all of the different causal questions that epidemiologists may wish to ask, and to consider the issues involved in answering them. Scientific enquiry, including physics, biology, and epidemiology, starts with a particular aspect of reality (matter, life, the distribution and determinants of disease in human populations) and questions about this reality (e.g. ‘What are the causes of asthma?’ or ‘Why is lung cancer increasing?’). In seeking to answer these questions, each science then simultaneously develops a methodology that seems appropriate,

by a kind of ‘reflective equilibrium’ process:³¹ methods that do not yield tenable results are rejected; and those that seem fruitful are retained and tested against each other.^{33,34} Every science must restrict the questions that its participants can ask and the methods used to answer them: otherwise, all causal inquiries would end up studying the Big Bang. However, over-restriction results in exclusion of areas of potential knowledge.

We conclude that an inclusive approach to scientific enquiry in epidemiology, which we call pragmatic pluralism, is not only appropriate but necessary. We agree with Bollen and Pearl⁷ that the essential ingredient of causation is responsiveness to variations in values of variables—however this is caused. In principle, all of types of causes of health states are susceptible to causal investigation and causal verdicts by epidemiologists, not just causes and situations amenable to certain methods. Furthermore, the issues in establishing causality for a particular factor are often highly hypothesis- and context-specific. Perhaps the only general conclusion that can be drawn is that any attempt at causal assessment involves integrating evidence of a variety of types from a variety of sources; no single study can establish causality for any of these types of causes, just as there is no single study design (not even an RCT) which can establish causality in itself.

Thus we hold that epidemiology, like other sciences, starts with its subject matter (distribution and determinants of the health of populations), and that introductory epidemiology courses should likewise start with this subject matter. As other sciences do, epidemiology should use a variety of study designs and types of data to attempt to identify causes.²² Although some approaches may be more valid in specific circumstances, such as randomization in the presence of strong confounding by indication, there are no definitive ‘rules’ as to which types of evidence are admissible and which are not. In our original paper,⁹ we give the example of the crucial role of time trend evidence in the debates about smoking and lung cancer. Time trend data are not intrinsically strong data for causality, yet in the context of the debate they proved crucial because they effectively ruled out one credible competitor to the smoking hypothesis (the constitutional hypothesis).

Moreover, putting all of the evidence together ultimately remains an exercise in judgement. This judgement is not formalizable, for reasons given in A2 above. Such judgements may be aided by articulations such as those of Bradford Hill³⁵ and the US Surgeon General,³⁶ but they cannot be developed into fully formal tools because there are no logically sufficient conditions for causal inference, as discussed in A2 above. There is nothing unusual about this: it applies to all sciences, since all empirical scientific enquiry depends on inductive inference.^{19,22,37}

In this context, the methods developed and proposed by the various commentators on our paper are invaluable as data analysis tools and as articulations which assist in the structuring and expression of the act of judgement that is at the heart of a causal inference. But these methods, and the principles that are often espoused along with them, do not define epidemiology. Rather, they are a set of useful tools which can be used alongside other approaches.

We believe that the activity of synthesizing evidence from multiple sources to arrive at an informed judgment is a much broader one than the estimation of the magnitude of causal effects from a single study or group of similar studies. Most of our commentators have now explicitly acknowledged this. But we reiterate our view that the effect estimation of potential interventions has been almost exclusively emphasized by the RPOA, at the cost of the wider set of causal inference activities. Calling a collection of new ideas, however worthwhile, ‘causal inference’ is misleading if only a small part of causal inference is treated. Our fear is that the messy, informal but very important expertise of doing epidemiology might be lost or diminished, given the intellectual appeal of the RPOA’s formalism and the lack of accompanying caveats about its limitations. We hope that this exchange of views has helped to clarify the strengths of these new methods as major contributions to causal inference, although emphasizing that they constitute just part, not the totality, of this field.

Funding

The Centre for Global NCDs is supported by the Wellcome Trust Institutional Strategic Support Fund, 097834/Z/11/B. Part of this research was supported by the National Research Foundation of South Africa (Incentive Funding for Rated Researchers).

Acknowledgements

The authors are grateful to the authors of the commentaries for their careful thoughts.

Conflict of interest: None declared.

References

- Blakely T. DAGs and the restricted potential outcomes approach are tools, not theories of causation. *Int J Epidemiol* 2016, in press.
- Daniel R, De Stavola B, Vansteelandt S. The formal approach to quantitative causal inference in epidemiology: misguided or misrepresented? *Int J Epidemiol* 2016, in press.
- Robins JM, Weissman MB. Counterfactual causation and street-lamps: what is to be done? *Int J Epidemiol* 2016, in press.
- VanderWeele TJ. On causes, causal inference, and potential outcomes. *Int J Epidemiol* 2016, in press.
- Weed DL. Causal inference in epidemiology: potential outcomes, pluralism and peer review. *Int J Epidemiol* 2016, in press.
- Hernán MA, Taubman SL. Does obesity shorten life? The importance of well-defined interventions to answer causal questions. *Int J Obes* 2008;32:S8–S14.
- Bollen K, Pearl J. Eight myths about causality and structural equation models. In: Morgan S (ed). *Handbook of Causal Analysis for Social Research*. New York, NY: Springer, 2013.
- Hernán MA. Invited commentary: hypothetical interventions to define causal effects—afterthought or prerequisite? *Am J Epidemiol* 2005;162:618–20.
- Vandenbroucke JP, Broadbent A, Pearce N. Causality and causal inference in epidemiology: the need for a pluralistic approach. *Int J Epidemiol* 2016; in press.
- Lewis D. Causation as influence. In: Collins J, Hall N, Paul LA (eds). *Causation and Counterfactuals*. Cambridge, MA: MIT Press, 2004.
- Lewis D. Counterfactuals and comparative possibility. *J Philos Log* 1973;2:418–46.
- Lewis D. Counterfactual dependence and time’s arrow. *Noûs* 1979;13:455–76.
- Stalnaker R. A defense of conditional excluded middle. In: Harper WL, Stalnaker R, Pearce G (eds). *Ifs*. Dordrecht, The Netherlands: D. Reidel, 1981.
- Pearl J. *Causality: Models, Reasoning and Inference*. 2nd edn. Cambridge, UK: Cambridge University Press, 2009.
- VanderWeele TJ, Robinson WR. On the causal interpretation of race in regressions adjusting for confounding and mediating variables. *Epidemiology* 2014;25:473–84.
- Cartwright N. *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. New York, NY: Cambridge University Press, 2007.
- Cartwright N, Hardie J. *Evidence Based Policy: A Practical Guide to Doing It Better*. New York, NY: Oxford University Press, 2012.
- Resnik D. *The Price of Truth: How Money Affects the Norms of Science*. New York, NY: Oxford University Press, 2007.
- Ladyman J. *Understanding Philosophy of Science*. London: Routledge, 2002.
- Douglas H. *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press, 2009.
- Lipton P. Contrastive explanation. In: Knowles D (ed). *Explanation and its Limits*. Cambridge, UK: Cambridge University Press, 1990.
- Lipton P. *Inference to the Best Explanation*. 2nd edn London and New York: Routledge, 2004.
- Broadbent A, Vandenbroucke J, Pearce N. Authors’ reply to: VanderWeele *et al.*, Chiolero, and Schooling *et al.* *Int J Epidemiol* 2016; doi: 10.1093/ije/dyw163.
- VanderWeele TJ, Hernán MA, Tchetgen Tchetgen EJ, Robins JM. Re: Causality and causal inference in epidemiology: the need for a pluralistic approach. *Int J Epidemiol* 2016;doi: 10.1093/ije/dyv341.
- Greenland S. Epidemiologic measures and policy formulation: lessons from potential outcomes. *Emerg Themes Epidemiol* 2005;2:1–7.
- Vandenbroucke JP. Observational research, randomised trials, and two views of medical science. *PLoS Med* 2008;5:0339–43.
- Pearce N. Tradition epidemiology, modern epidemiology, and public health. *Am J Public Health* 1996;86:678–83.

28. Pearce N. Epidemiology as a population science. *Int J Epidemiol* 1999;28:S1015–18.
29. Cartwright N. Predicting what will happen when we act. What counts for warrant? *Prev Med* 2011;53:221–24.
30. Pearce N, Greenland S. Confounding and interaction. In: *Handbook of Epidemiology*. Heidelberg, Germany: Springer, 2014.
31. Greenland S. Relation of probability of causation to relative risk and doubling dose: a methodologic error that has become a social problem. *Am J Public Health* 1999;89:1166–69.
32. Goodman N. *Fact, Fiction and Forecast*. 4th edn. Cambridge, MA: Harvard University Press, 1983.
33. Quine WV. Two dogmas of empiricism. In: *From a Logical Point of View*. Cambridge, MA: Harvard University Press, 1953.
34. Quine WV. Epistemology naturalized. In: *Ontological Relativity and Other Essays*. New York and London: Columbia University Press, 1969.
35. Bradford Hill A. The environment and disease: association or causation? *Proc R Soc Med* 1965;58:259–300.
36. Advisory Committee to the Surgeon General of the Public Health Service. *Smoking and Health*. Washington, DC: Department of Health, Education and Welfare. 1964.
37. Broadbent A. *Philosophy for Graduate Students: Metaphysics and Epistemology*. London and New York: Routledge, 2016.