

Commentary: Incorporating concepts and methods from causal inference into life course epidemiology

Bianca L De Stavola and Rhian M Daniel

LSHTM Centre for Statistical Methodology and Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London, UK

*Corresponding author. Bianca L. De Stavola, LSHTM Centre for Statistical Methodology and Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London, UK. E-mail : bianca.destavola@lshtm.ac.uk

Accepted 15 March 2016

International Journal of Epidemiology, 2016, 1006–1010
doi: 10.1093/ije/dyw103



Introduction

The review by Ben-Shlomo *et al.*¹ highlights how life course epidemiology is evolving and adapting to accommodate increasing access to data on novel dimensions and over extended periods. This enriched framework raises ever greater methodological challenges, leaving statisticians like us daunted by the task of translating life course enquiries into suitable analyses of the data at hand.

Take for example Figure 4 of Ben-Shlomo *et al.*¹ This is very useful for gaining a ‘big picture’ understanding of a complex area such as ageing, and for establishing which processes may benefit from a more detailed investigation. However, the leap from such a diagram to a specific data analysis should not be (and is not typically) made without greater thought. We will argue in this commentary that some recent developments from the field of modern causal inference may be helpful in this regard. First, in order to state unambiguously the question (or questions) of interest, the potential outcomes framework, a cornerstone of modern causal inference thinking, is invaluable. Then, the conceptual framework should be refined to a causal directed acyclic graph (DAG) relevant to the question, and the causal DAG should be formally interrogated to see if the question can be addressed, and if so how. Indeed, depending on the question, the causal DAG and the data available, we may find that standard statistical methods traditionally used in epidemiology are sufficient; in other settings we may find that more novel techniques are needed.

We will discuss each of these points next, mentioning also the issues of missing data and measurement error, as well as highlighting concerns about the difference between

the processes which are the focus of investigations and their manifestations in observed data.

What is the question of interest?

For illustration, we take an example briefly discussed by Ben-Shlomo *et al.*,¹ namely the relationship between nutrition and type II diabetes. Whereas the broad aim of a project may be to understand the effect of nutrition across the life course on the risk of developing adult type II diabetes, a more specific question must be established before we can proceed. In this section, we will highlight the range of different questions that may be of interest, and how they can be unambiguously distinguished using the potential outcomes notation.

Potential outcomes

Suppose for simplicity that nutrition is reliably measured (e.g. via detailed food frequency questionnaires together with analysis of urinary samples) twice during the life course on a cohort of people: once in childhood and again in early adulthood. Let X_1 and X_2 denote relevant summaries of nutritional status in childhood and early adulthood, respectively, and let the binary variable Y denote the development of type II diabetes by age 70 years, say. Let $Y(x_1)$ be the potential outcome, i.e. the value that Y would take if we were hypothetically to intervene on X_1 and set it to the value x_1 . We can similarly define potential outcomes $Y(x_2)$ for hypothetical interventions on X_2 . We may also define $Y(x_1, x_2)$, the value that Y would take were we to intervene on both X_1 and X_2 and set them to x_1 and x_2 , respectively. We will now use this very simple example to

illustrate how subtly different questions of interest can be articulated using these potential outcomes.

Total and joint effects

The total causal effect (TCE) of X_1 on Y (including both its direct effect and its indirect effect via X_2) can be expressed as a comparison of the distribution of $Y(x_1)$ for different values of x_1 . Often the mean is compared, as in $TCE_1(x_1) = E\{Y(x_1)\} - E\{Y(x_1^*)\}$, where x_1^* is a reference (or baseline) value of X_1 ; for binary X_1 , $TCE_1 = E\{Y(1)\} - E\{Y(0)\}$. If Y were a time-to-event outcome, such as time to onset of type II diabetes, we could compare the survivor functions of the potential outcomes, for example.² The TCE of X_1 would be of public health interest, for example if primary school nutrition programmes were being considered. If a public health nutrition initiative targeted at adults would instead be considered, then the causal effect of X_2 on Y , TCE_2 would more likely be of interest, which would involve a comparison of the distribution, e.g. the mean, of $Y(x_2)$ for different values of x_2 . Alternatively, the likely impact of a general intervention such as increased taxation of unhealthy food would more naturally lead to a comparison of the distribution of $Y(x_1, x_2)$ for different values of x_1 and x_2 : this is the joint effect of (X_1, X_2) on Y .^{3,4}

Controlled direct effects and conceptual models in life course epidemiology

Another possible aim might be to gain a better biological understanding of the timing and strength of the mechanisms linking nutrition to type II diabetes prevalence. In this case we might compare the distribution of $Y(x_1, x_2)$ for different values of x_1 but for a fixed value of x_2 , known as a controlled direct effect, for example $CDE_1(x_2) = E\{Y(x_1, x_2)\} - E\{Y(x_1^*, x_2)\}$.⁵ Evidence of variation in $CDE_1(x_2)$ for different values of x_2 would indicate that the effect of nutrition in childhood on the risk of diabetes varies according to the level at which adult nutrition is set. This would support the so-called pathways model discussed in life course epidemiology^{6,7} according to which sensitive periods of exposure interact in their impact on risk (see below for further discussion of this). In the absence of such effect heterogeneity, we could compare the common CDE_1 , which represents the effect of X_1 that is not mediated by X_2 , with the total causal effect of X_2 , TCE_2 . A similarity between them would support the cumulative exposure model. This is because the similarity (in addition to the lack of interaction) implies that experiencing the exposure during each of these two periods (directly) influences the risk of type 2 diabetes by the same amount,

and hence it is the cumulative exposure, rather than the timing of it, that matters. If instead one or other effect were much smaller than the other, there would be support for the sensitive period model, and one or other effect being zero would support the critical period model.^{6,8}

Interaction versus effect modification

Returning again to the pathways model, there are two subtly different possible questions even here, which can clearly be articulated using potential outcomes, namely the difference between interaction and effect modification.⁹ For ease of explanation, suppose that X_1 and X_2 are both binary, and that we are interested in comparing the means of the potential outcomes. An interaction is said to be present if the two CDEs differ, i.e. if $E\{Y(1, 1)\} - E\{Y(0, 1)\} \neq E\{Y(1, 0)\} - E\{Y(0, 0)\}$. This is the same as saying that: $E\{Y(1, 1)\} - E\{Y(1, 0)\} \neq E\{Y(0, 1)\} - E\{Y(0, 0)\}$, i.e. the causal effect of changing adult nutrition on the risk of type II diabetes differs according to the level at which we set childhood nutrition. Conversely effect modification, as defined by VanderWeele,⁹ would be present if the causal effect of changing adult nutrition on the risk of type II diabetes differs between those who in reality have different childhood nutrition statuses, i.e. if $E\{Y(x_2 = 1) - Y(x_2 = 0) | X_1 = 1\} \neq E\{Y(x_2 = 1) - Y(x_2 = 0) | X_1 = 0\}$. (We explicitly write $Y(x_2 = 1)$ instead of $Y(1)$ here, to clarify that the hypothetical intervention being considered is on X_2 rather than X_1 .)

In this literature, interaction has a causal connotation with respect to both exposures, whereas effect modification is causal only with respect to (in this case) the later exposure X_2 . Which of these questions is of interest will depend on the broader aim of the investigation, and will have an impact on how the data are analysed.

Effect decomposition

Alternatively, we might be interested in effect decomposition, i.e. in asking what proportion of the effect of childhood nutrition on type II diabetes is mediated by early-adult nutrition. For these questions, so-called natural direct and indirect effects are relevant,^{10,11} and can again be unambiguously stated as a counterfactual comparison; the natural direct effect, for example, is a comparison of the distribution of $Y\{x_1, X_2(x_1^*)\}$ for different values of x_1 where $X_2(x_1^*)$ is the potential value of X_2 were we to set X_1 to x_1^* . More specifically, the natural effect of X_1 , expressed as a mean difference, is defined as $NDE_1 = E\{Y\{x_1, X_2(x_1^*)\}\} - E\{Y\{x_1^*, X_2(x_1^*)\}\}$. These effects and their estimation have received much attention in the recent causal inference literature.⁵ Since these effects involve

nested counterfactuals, they require very strong untestable assumptions for identification, assumptions that could not even be hypothetically verified in an experimental setting. For this reason, attention is currently being diverted to more policy-relevant effects known as interventional direct and indirect effects identifiable under weaker conditions.¹² We expect that these effects will soon be estimated in applications in life course epidemiology.

Multiple exposures/mediators

Suppose we had an additional exposure time point, such as nutrition during infancy; we would then be in a setting in which we could potentially be interested in the joint effects of more than two exposures,^{3,4} or we might be interested in effect decomposition with multiple mediators.^{13,14}

Comment

In any particular study, it is unlikely that all of the above would be relevant. Our main message, however, is that life course investigations are causal enquiries. Familiarity with the modern causal inference literature—and with the subtly different flavours of causal effects defined therein—has the potential to aid researchers in formulating and communicating the question(s) of interest.

What is the appropriate causal DAG and what can it tell us?

Once the question of interest has been stated, establishing whether it can be answered under plausible assumptions using the data at hand, and if so how, can be aided by drawing a causal diagram [more precisely, a causal directed acyclic graph (DAG)].¹⁵ Such a causal DAG should reflect a priori subject-matter knowledge regarding the likely causal structure of the variables being studied. Unlike conceptual frameworks, causal DAGs are well-defined mathematical objects that can be interrogated using procedures such as *d*-separation;¹⁶ as such, care must be taken to draw them correctly, otherwise the resulting conclusions will be

unreliable. Further, although a unique causal DAG can never be determined from the data and hence subject-matter knowledge is crucial, some candidate causal DAGs may be incompatible with the data, and thus compatibility should be investigated. There exists some recent technical literature on the various possible causal interpretations of DAGs,^{17,18} but one shared feature is that to be causal, any common (measured or unmeasured) cause of two or more variables in the diagram must itself be in the diagram.

See, for example, our Figure 1 for the nutrition-diabetes example. The most naive causal DAG would assume no common causes of any pair of X_1, X_2, Y and, under this assumption, finding a causally interpretable statistical analysis of the data would be straight-forward. More realistically, however, the nutrition-diabetes relationship will be confounded by a number of factors that we denote by C [e.g. socioeconomic position, physical activity and body mass index (BMI)]. In most life course settings, many of these confounders themselves will change over time (hence C_1 and C_2 in the figure), and to make progress, reliable repeated measures on these confounders would be needed. In particular, note that we have allowed the later confounders C_2 to be affected by the earlier exposure X_1 (as would be expected, particularly with say physical activity and BMI). This dependence of C_2 on X_1 introduces a potentially problematic feature common in life course studies, namely time-dependent (or intermediate) confounding.

If Figure 1(b) were correct, standard regression methods could be used to make inferences about the total effect of X_2 (by adjusting for X_1, C_1, C_2), and also about effect modification by X_1 of the total effect of X_2 on Y .¹⁹ In the unlikely event that we could additionally assume no unmeasured common causes U of X_1 and X_2 , then inference about the total effect of X_1 could also be made using standard methods upon adjustment for C_1 . If, however, we wished to make inference about the joint effects of X_1, X_2 , or the controlled direct effect of X_1 on Y fixing X_2 to a particular value, or to learn about interaction between X_1 and X_2 in their effect on Y , then so-called g-methods that deal with time-dependent confounding would be needed.^{3,4}

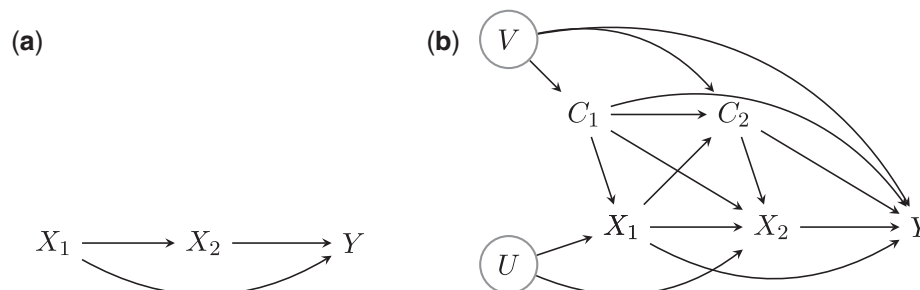


Figure 1. (a) A naive causal DAG. (b) A more realistic causal DAG, with the unmeasured variables circled.

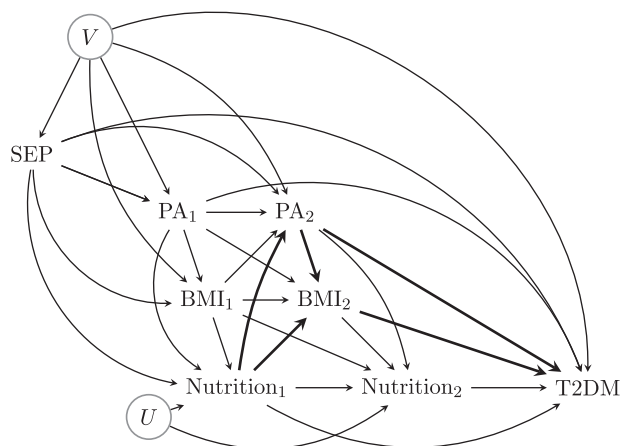


Figure 2. A causal DAG for the nutrition-diabetes example.

Effect decomposition using natural direct and indirect effects would not be possible in the presence of U , and would anyway require additional strong parametric assumptions (due to the presence of intermediate confounding).^{20,21} However, progress could be possible using interventional mediation effects¹² using the observed (rather than counterfactual) distribution of X_2 given X_1 and C_1 ; see a related discussion by VanderWeele and Robinson.²²

Figure 2 is an expanded version of Figure 1(b) specific to the nutrition-diabetes example. It highlights how controlling for adult physical activity and BMI (C_2) would remove part of the effect of childhood nutrition, the part that they mediate (thicker arrows in the DAG), whereas not controlling for adult physical activity and BMI would confound the effect of adult nutrition. Hence a traditional regression approach would not achieve the estimation of the joint effects of childhood and adult nutrition (including their interaction), nor of the controlled direct effect of childhood nutrition.

Missing data and measurement error

As well as confounding, other challenges facing those attempting to estimate causal effects in life course studies are: the likely depletion of participation during the course of follow-up, with the resulting missing data; and measurement error in particular when it affects a mediator in the considered analysis. The first challenge is a particularly relevant problem in studies of ageing, as also discussed by Ben-Shlomo *et al.*,¹ where data missing due to death or other competing events present an especially difficult problem.²³ Measurement error is ubiquitous in observational epidemiology and is arguably of even greater concern when it affects a mediator, with induced biases in opposite directions for direct and indirect effects.^{24,25} Inevitably when faced with incomplete

and/or mismeasured data, as with confounding, assumptions have to be made and then sensitivity to these assumptions assessed. Causal diagrams present an opportunity to represent these assumptions, for example by including missingness indicators (see for example Daniel *et al.*²⁶ and the extensions by Mohan, Pearl and others)²⁷ or measurement error mechanisms in the DAG.²⁸ The implications of different assumptions can then be formally assessed, and appropriate analyses, if they exist, identified.

Processes versus snapshots

We end our commentary by drawing attention to a recent important cautionary note on the use of causal DAGs, by Aalen *et al.*,²⁹ particularly relevant to studies in life course epidemiology. Conceptual frameworks, e.g. Figure 4 in reference 1, rightly concern processes, usually latent, such as the continuous-time evolution of an individual's reproductive function or nutritional trajectory over the life course. However, when we draw causal DAGs that are to be useful for informing data analysis, we naturally focus on the 'snapshot' observations of this latent process that are available in our data. As long as the translation from process to snapshot is done knowingly and carefully, this is wise; however, if we treat the snapshots in our causal DAGs as if they actually represented the whole process, then mistakes can be made and conclusions adversely affected. For example, one reasonable hypothesis may be that the effect of nutrition on type II diabetes risk is entirely mediated by BMI. However, the mediation in this instance would be through the entire BMI process. If BMI is measured only on a relatively small number of occasions, then we would expect only a proportion of the effect of nutrition on type II diabetes risk to be mediated by the observed measures of BMI.

Summary

In this commentary we have highlighted concepts and methods from the field of causal inference, many of them recent contributions, which we believe to be relevant in life course studies. We have discussed how the language of potential outcomes can help to articulate the precise question(s) being asked and how causal DAGs—distinct from conceptual frameworks—should be carefully drawn and interrogated and missing data and measurement error mechanisms included.

Funding

R.D. acknowledges support from a Sir Henry Dale Fellowship jointly funded by the Wellcome Trust and the Royal Society (Grant

Number 107617/Z/15/Z); and B.D.S. acknowledges support from the Economic and Social Research Council (grants ES/I025561/1, ES/I025561/2 and ES/I025561/3).

Conflict of interest: None declared.

References

- Ben-Shlomo Y, Cooper R, Kuh D. The last two decades of life course epidemiology, and its relevance for research on ageing. *Int J Epidemiol* 2016;**45**:973–88.
- Hernán MA. The hazards of hazard ratios. *Epidemiology* 2010; **21**:13–15.
- Robins JM, Hernán MA. Estimation of the causal effects of time-varying exposures. In: Fitzmaurice G, Davidian M, Verbeke G, Molenberghs G (eds). *Longitudinal Data Analysis*. Boca Raton, FL: Chapman and Hall/CRC Press, 2009.
- Daniel RM, De Stavola BL, Cousens SN, Kenward MG, Sterne JAC. Methods for dealing with time-dependent confounding. *Stat Med* 2013; **32**:1584–618.
- VanderWeele TJ. *Explanation in Causal Inference: Methods for Mediation and Interaction*. Oxford, UK: Oxford University Press, 2015.
- Mishra G, Nitsch D, Black S, De Stavola BL, Hardy R. A structured approach to modelling the effects of binary exposure variables over the life course. *Int J Epidemiol* 2009;**38**: 528–37.
- Nandi A, Glymour MM, Kawachi I, VanderWeele TJ. Using marginal structural models to estimate the direct effect of adverse childhood social conditions on onset of heart disease, diabetes, and stroke. *Epidemiology* 2012;**23**:223–32.
- De Stavola BL, Daniel RM. Marginal structural models: The way forward for life-course epidemiology? *Epidemiology* 2012;**23**:233–37.
- VanderWeele TJ. On the distinction between interaction and effect modification. *Epidemiology* 2009;**20**:863–71.
- Robins KM, Greenland S. Identifiability and exchangeability for direct and indirect effects. *Epidemiology* 1992;**3**:143–55.
- Pearl J. Direct and indirect effects. *Proceedings of the Seventeenth Conference On Uncertainty in Artificial Intelligence, 2-5 August 2001*. Seattle, WA: UAI, 2001.
- VanderWeele TJ, Vansteelandt S, Robins JM. Effect decomposition in the presence of an exposure-induced mediator-outcome confounder. *Epidemiology* 2014;**25**:300–06.
- VanderWeele TJ, Vansteelandt S. Mediation analysis with multiple mediators. *Epidemiol Methods* 2014;**2**:95–115.
- Daniel RM, De Stavola BL, Cousens SN, Vansteelandt S. Causal mediation analysis with multiple mediators. *Biometrics* 2015;**71**:1–14.
- Pearl J. Causal diagrams for empirical research. *Biometrika* 1995;**82**:669–710.
- Pearl J. *Causality*. 2nd edn. New York, NY: Cambridge University Press, 2009.
- Robins JM, Richardson TS. Alternative graphical causal models and the identification of direct effects. In: Shrouf P, Keyes K, Ornstein K (eds). *Causality and Psychopathology: Finding the Determinants of Disorders and their Cures* Washington, DC: APPA, 2011.
- Richardson TS, Robins JM. Single World Intervention Graphs (SWIGs): a unification of the counterfactual and graphical approaches to causality. Center for Statistics and the Social Sciences, University of Washington, 2013.
- Keogh RH, Daniel RM, VanderWeele TJ, Vansteelandt S. Analysis of longitudinal studies: Adjusting for time-dependent confounding using conventional methods. *Am J Epidemiol* in press.
- De Stavola BL, Daniel RM, Ploubidis G, Micali N. Mediation analysis with intermediate confounding: structural equation modeling viewed through the causal inference lens. *Am J Epidemiol* 2015;**181**:64–80.
- Tchetgen Tchetgen EJ, VanderWeele TJ. On identification of natural direct effects when a confounder of the mediator is directly affected by exposure. *Epidemiology* 2014;**25**:282–91.
- VanderWeele TJ, Robinson WR. On the causal interpretation of race in regressions adjusting for confounding and mediating variables. *Epidemiology* 2014; **25**:473–84.
- Rubin DB. Causal inference through potential outcomes and principal stratification: application to studies with censoring due to death. *Stat Sci* 2006;**21**:299–309.
- le Cessie S, Debeij J, Rosendaal FR, Cannegieter SC, Vandenbroucke J. Quantification of bias in direct effects estimates due to different types of measurement error in the mediator. *Epidemiology* 2012;**23**:551–60.
- VanderWeele TJ, Valeri L, Ogburn EL. The role of measurement error and misclassification in mediation analysis. *Epidemiology* 2012;**23**:561–64.
- Daniel RM, Kenward MG, Cousens SN, De Stavola BL. Using causal diagrams to guide analysis in missing data problems. *Stat Methods Med Res* 2012;**21**:243–56.
- Mohan K, Pearl J, Tian J. Graphical models for inference with missing data. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ (eds). *Advances in Neural Information Processing Systems*, Norwich, UK: Curran, 2013.
- VanderWeele TJ, Hernán MA. Results on differential and dependent measurement error of the exposure and the outcome using signed directed acyclic graphs. *Am J Epidemiol* 2012;**175**:1303–10.
- Aalen O, Røysland K, Gran J, Kouyos R, Lange T. Can we believe the DAGs? A comment on the relationship between causal DAGs and mechanisms. *Stat Methods Med Res* 2014, March 28. PMID: 24463886. [Epub ahead of print.]