Abstract

Many methods have been proposed to solve the age-period-cohort (APC) linear identification problem, but most are not theoretically informed and may lead to biased estimators of APC effects. One exception is the mechanism-based approach recently proposed and based on *Pearl's front door criterion*; it ensures consistent APC effect estimators in the presence of a complete set of intermediate variables between one of age, period, cohort and the outcome of interest, as long as the assumed parametric models for all the relevant causal pathways are correct. Through a simulation study mimicking APC data on cardiovascular mortality, we demonstrate possible pitfalls that users of the mechanism-based approach may encounter under realistic conditions, namely when 1) the set of available intermediate variables is incomplete; 2) intermediate variables are affected by two or more of the APC variables, but this feature is not acknowledged in the analysis 3) unaccounted confounding is present between intermediate variables and the outcome. Furthermore, we show how the mechanism-based approach can be extended beyond the originally proposed linear and probit regression models to incorporate all generalized linear models, as well as non-linearities in the predictors, using Monte Carlo simulation. Based on the observed biases that resulted from departures from underlying assumptions we formulated guidelines for the application of the mechanism-based approach (extended or not).

Introduction

Demographers, epidemiologists, sociologists and others have attempted to break down outcomes of interest into constituent effects caused by, or associated with, age, calendar time, and time of birth. This is also known as age-period-cohort (APC) analysis. Age effects refer to changes in the outcome as the age of individuals in the study population progresses. For example, as individuals age, cardiovascular function declines and hence older individuals tend to have worse cardiovascular health than younger individuals. Period effects refer to changes that occur in an outcome as calendar time progresses. They can represent sudden changes or temporary changes in an outcome, such as spikes in death rates due to war or famine, but may also represent gradual changes such as those produced by the accumulation of minor improvements in public health infrastructure over time influencing mortality rates in all age-groups. Finally, birth cohort effects represent differences between generations that aren't attributable to differences in age or calendar time. Conceptually, they commonly represent the effects of shared formative experiences of individuals in a birth cohort, either *in utero* or during other critical phases in the life course (Ben-Shlomo and Kuh 2002). The effect of these formative years would then remain largely constant in that cohort throughout the remaining life course, and are therefore independent of age and calendar time. For example, the cohort that was *in utero* during the Dutch Hunger Winter in 1944-1945 had worse health even later in life (Ekamper et al. 2015), compared to other cohorts. Furthermore, birth cohort has been found to be strongly tied to smoking behaviour in various Western countries (Preston et al. 2006; Verlato et al. 2006).

Unfortunately, decomposing an outcome (Y) into the separate effects of age (*A*), period (*P*) and cohort (*C*) using for example a linear regression model (e.g. $E(Y|A,P,C) = \eta + \alpha \cdot A + \beta \cdot P + \theta \cdot C$) imposes an identification problem. Since A = P - C there is linear dependency between the three variables, and consequently any linear model involving these three variables cannot

have a unique solution. To circumvent this, various techniques have been introduced that constrain the model specification (e.g. Clayton & Schifflers 1987; Held and Riebler 2012; Holford 2006; Yang et al. 2008), so that a solution can be found. However, the technical constraints that have been proposed are arbitrary and do not lead to meaningful measures of effect (Bell and Jones 2014; Luo 2013). Estimation of the parameters may be unbiased, but only under the constraints that have been imposed, and hence do not reflect the 'true effects' (see below) of age, period and cohort that we seek (Luo 2013). Luo and Fienberg both argue in favour of a 'paradigm shift' in a recent discussion in *Demography* (Fienberg 2013; Kuhn 1996; Luo 2013); APC analysis, they argue, needs to become more theoretically informed. Simply fitting a regression model to an outcome given age, period and cohort, without any forethought or theoretical reasoning, cannot result in meaningful effect estimates for these variables.

Few authors explicate what they mean by 'true' or 'meaningful' effect estimates. Viewed from one perspective, since the relationship A = P - C always holds, then a regression model with age, period and cohort as covariates truly has infinitely many solutions, and thus there is no problem to be solved. However, since those who write about this problem talk about there being one special solution out of those infinitely many solutions which is 'correct'/'true'/'valid'/'meaningful', it must be that they are (albeit implicitly) thinking of a hypothetical world, different from the actual world, in which age, period and cohort can be manipulated such that the identity A = P - C is broken.

More formally, one could take an explicitly causal perspective using potential outcomes with age, period and cohort as 'independent' exposures. Let Y(a,p,c) be the potential outcome which would occur if A were set to a, P to p and C to c, without necessarily abiding by the relationship a = p - c (Rubin 1974). Then the causal model

$$E\big(Y(a,p,c)\big) = \eta^* + \alpha^* \cdot a + \beta^* \cdot p + \theta^* \cdot c \qquad (1)$$

has one solution, and this is presumably the 'true' solution to which the various authors on this topic refer.

In fact, imagining a hypothetical world in which time can be manipulated is difficult enough; contemplating one in which three different 'aspects' of time, namely age, period and cohort, can be independently manipulated requires an even wilder imagination, and is therefore unlikely to be truly of interest. More realistically, we can view Eq. 1. as being a shorthand for

$$E\left(Y\big(c_a, c_p, c_c\big)\right) = \eta' + \alpha' \cdot c_a + \beta' \cdot c_p + \theta' \cdot c_c \qquad (2)$$

where c_a, c_p and c_c are the set of all immediate *consequences* of age, period and cohort, respectively. So if being born in a particular cohort meant that you would be born during a famine, it is this famine that we imagine we could manipulate (and, say, prevent), rather than the cohort of birth itself. But since we may not have all these consequences at our disposal, model (2) is replaced (as a shorthand) by model (1).

Given this reframing of the model of interest as a causal model, it makes sense to consider methods from causal inference (Pearl 2000) to analyse data from APC studies. This was done by Winship & Harding (Winship and Harding 2008) in 2008, and was called the mechanism-based approach. In particular, their approach uses *Pearl's front door criterion* to identify the APC causal parameters α*, β*, and θ* (Pearl 2000). In short, the mechanism-based approach uses intermediate variables on the path between one of the three APC variables and the outcome in order to estimate the effect of one of these three variables on the outcome indirectly, and to estimate the effect of the remaining two APC variables directly (with the method generalizable to modelling intermediate variables for two of the three APC variables). The approach naturally leads to drawing a directed acyclic graph (Glymour 2006; VanderWeele et al. 2008) depicting the

assumed relationships between A, P, C, the intermediate variables being considered, and the outcome. It thus motivates researchers to be explicit about their substantive assumptions.

The method requires that a complete set of intermediate variables can be found for at least one of the three APC variables. By a '*complete* set of intermediate variables' for A, for example, we mean a set of variables M_1, M_2, …, M_K that are affected by age and which themselves affect the outcome Y in such a way that *all* of the effect of A on Y is via this set of intermediate variables.

However, in a realistic setting, finding a complete set of intermediate variables even for just one of the three APC variables is unlikely. Also, the partial set of intermediate variables that may be available could be dependent on more than one APC variable. Furthermore there may be variables that affect both the intermediate variable(s) and the outcome. All these settings (if they cannot somehow be accounted for) threaten the mechanism-based approach with bias, and one of the aims of this paper is to demonstrate these potential sources of bias and their magnitude in these realistic scenarios.

Another challenge also arises; the mechanism-based approach has been developed for use in linear and probit regression models for Y, and linear and probit regression models for the intermediate variables M_1, M_2, …, M_K. While some analytical solutions (e.g. Winship and Mare 1983) could be adopted to extend this approach to deal with logistic regression models for outcome and/or mediators, they are complex to implement. Moreover, only approximate methods are available to deal with settings where the variables included in the outcome model interact or have some other non-linear effects, even when Y, M_1, M_2, …, M_K are all continuous and modelled using linear models (Preacher and Hayes 2008; VanderWeele 2015; Jiang and VanderWeele 2015).

In this paper, in order to illustrate possible pitfalls one may encounter using mechanism-based APC models, we assess their performance under realistic settings, namely when 1) only a partial set of mediators is available; 2) some of the intermediate variables are affected by two or more of the APC variables, but this feature is not acknowledged in the analysis; 3) unmeasured confounding affects the intermediate variables and the outcome. Furthermore, we extend the mechanism-based approach to settings with any fully-parametric model for the outcome and intermediate variables by approximating the estimation of the APC parameters by Monte Carlo simulation. R code demonstrating the mechanism-based approach, and its extension, is available as supplementary material.

Methods

The mechanism-based approach

The mechanism-based approach is based on exploiting the fact that age, period and cohort affect the outcome through intermediate variables (hereafter called 'mediators') (Winship and Harding 2008). The key idea is that while age, period and cohort are deterministically related, the intermediate variables along the paths from these to the outcome will be affected by other (APC-independent) causes, and hence can be used (if measured) to circumvent the identification problem. We now discuss this in more detail.

Consider for simplicity the setting depicted in Figure 1, which shows a causal directed acyclic graph (DAG). Causal DAGs are formal graphical representations of the assumed causal relationships between the variables under study (Glymour 2006). Here the number of mediators K is equal to 2, and the mediators M_1 and M_2 being considered lie on causal pathways from P to Y. Note that there is no arrow from P to Y in the DAG, representing the assumption that all of the effect of P on Y is via M_1 and M_2. Also note that there is no arrow from either of A or C to the mediators, nor shared common causes of the mediators and any other variables in the DAG. Finally, the two paths from P to Y are 'separate', in the sense that M_1 does not affect M_2 nor vice versa, and neither does any variable along either path affect a variable on the other path, nor share any common causes: M_1 and M_2 are assumed to be conditionally independent given P. These strong structural assumptions concerning the roles of M_1 and M_2 (some of which can be relaxed – more on this later) allow the identification of the APC effects in a two-stage procedure as we now describe. In the first step, separate models for each mediator on P are fitted. In the second step, a model for the outcome Y on A, C, M_1 and M_2 is fitted. If Figure 1 is correct and the outcome and the mediators are continuous variables and modelled using linear regression, or the outcome and mediators are binary and modelled using probit regression models, and none of these models in

truth includes product terms or other non-linearities, then the effect of A and C on Y (α^* and θ^*) are equal to their regression coefficients in the outcome (second-step) model, while the effect of P on Y (β^*) is equal to the sum of the effects along the two pathways involving M_1 and M_2. The effect along a pathway is equal to the product of two regression coefficients; for the P – M_1 – Y pathway, it is the product of the coefficient for P in the (first-step) regression of M_1 on P and the coefficient for M_1 in the (second-step) regression of Y (that also includes M_2, A and C as covariates). Similar calculations apply to the P – M_2 – Y pathway, and the effects along the two pathways are then summed to obtain the effect of P on Y. These calculations are an application of the path tracing rules that are widely used in structural equation modelling (Mulaik 2009; Wright 1934), with standard errors for the estimated effect of P estimated using the delta method (in simple settings) or more generally the bootstrap (MacKinnon 2008). See the supplemental material for an applied example of the path tracing rule.

Complications

In a real life setting, a number of situations may occur that make mechanism-based estimation of APC effects less straightforward. Firstly, if a complete set of mediators is not available for the selected APC variable(s), then the effect estimators of the three APC variables described above will be biased for α^*, β^* and θ^* because the required assumption that (at least) one of the three APC variables is fully mediated by a set of measured mediators would not be met. Secondly, a variable that we believe to be a mediator for one of the three APC variables may actually be a mediator for more than one (for example, daily caloric intake is probably affected by age as well as period). In this case, the regression coefficient for the APC variable that we did not believe to be mediated by any of the mediators for P, M_1, M_2, …M_K, will only capture the component of its effect that is not mediated. Thirdly, the relationship between mediators and outcome may be

confounded, i.e. there may be a variable (either measured or unmeasured) that has a causal effect on both the mediator(s) and the outcome. If this confounding is not controlled for in the outcome model, the effect of the mediator on the outcome will be estimated with bias, and consequently also the effect of the APC variable that is assumed to be mediated by it. Finally, the outcome and/or the mediators may not be of a type that can be modelled by linear or probit regression, or even if they can, when the models require product terms or other non-linearities. In this situation, the path tracing rule needed to derive the causal effect of the mediated APC variable cannot be used (Mulaik 2009).

Simulation: approach

We assess the mechanism-based approach through simulations. In our simulations, we attempt to recreate a realistic setting in which APC analyses are performed, namely the study of cardiovascular mortality. However, in order to demonstrate particular pitfalls we isolate sources of bias in the APC effect estimates, and therefore simplify this real world setting into three scenarios.

In each scenario, the study population and outcome are the same. The individuals in the study population are generated to be aged between 40 and 95 years during the calendar years 1990 to 2015, and hence the whole data consists of birth cohorts ranging from 1896 to 1975. Age and period for each record are generated according to uniform distributions but then categorised into 5-year groups (for A and P), and cohort dependent on these categories (C = P – A). The outcome is mortality due to cardiovascular disease (CVD), coded as 1 (death due to CVD) or 0 (alive or dead from other causes). It was generated in all scenarios according to either a probit or logistic regression model, with the probability of CVD death generated as a function of age, cohort and the period mediators. Age is set to account for 70% of the effect of the APC variables

on cardiovascular mortality, period (via its mediators) for 20%, and birth cohort for 10%. We believe these percentages approximately correspond to realistic effects of age, period and cohort in the period 1990 to 2015 in Western countries; the difference in incidence of CVD death between individuals aged 40 and those aged 95 is very large, whereas the difference in incidence of CVD death in these age groups between the year 1990 and 2015 is much smaller (Peeters et al. 2011). Due to the linear dependency phenomenon, it is unknown what part of these differences is in truth attributable to each dimension.

Simulation: mediators and confounders

We simulate settings where P is the variable which has measured mediating variables. Results however easily generalise to the alternative scenarios where A or C play this role (with due numerical differences given their unequal assumed strength of effects). Four mediators on the path from P to Y are included in the simulation study. They are body mass index (BMI), smoking, statin therapy, and 'unmeasured'. We choose the former three variables because they are commonly described variables that are believed to affect CVD mortality, but many other variables are also believed to affect CVD mortality (e.g. Blackmore et al. 2015; Capewell et al. 2000), which are represented by the 'unmeasured' variable. Together, these four variables account for the entire period effect on the outcome. The 'unmeasured' and BMI variables are continuous, whereas smoking and statin therapy are binary. We set each of the 'measured' mediators to account for ~20% of the period effect on cardiovascular mortality, while we set the 'unmeasured' mediator to account for ~40%. Table 1 shows the direction of the effects of period on the mediators, and of each mediator on the outcome; the effect of period on the 'unmeasured' mediators is linear, while its effect on the measured mediators is non-linear but monotonically increasing or decreasing. Initially, these variables are made to act as mediators only on the path

between P and CVD mortality, but in some scenarios they are also affected by A or C, in which case the total effects of these latter variables changes. In the scenario that includes confounding (see below), the confounder is presence or absence of a particular gene, randomly assigned to be present in 50% of individuals and set to have a positive effect on both the mediators and the outcome (Sabol et al. 1998; Smith et al. 2015).

Simulation: scenarios and variants

In all simulations, we generate data for 100,000 individuals, each measured once. We simulate three scenarios, with each scenario simulated 1000 times. In the data generating process for the first scenario ('simple'), A and C have a direct effect on Y while only the effect of P is mediated. In the second scenario ('more causes'), there is an additional (negative) effect of A on the mediator BMI, which amounts to roughly 30% of the total age effect, and an additional (positive) effect of C on smoking, which amounts to roughly 30% of the total cohort effect. Finally, in the third scenario ('confounding'), genotype confounds the relationship between BMI and Y and between smoking and Y (Figure 2); genotype has a positive effect on both BMI and smoking, and a positive effect on CVD mortality. Genotype accounts for roughly 33% of the association between age and CVD death and 35% of association between cohort and CVD death.

Each scenario has two variants. In the first variant, we generate Y and the binary mediators using probit regression models, while in the second variant logistic regression models are used instead. In both variants, linear regression models are used to generate continuous variables. Since probit and logistic regression models transform parameters into probabilities in a different way, the probabilities of cardiovascular mortality somewhat differ between the two variants. We varied the value for the intercept in each variant so that the age-specific probabilities of cardiovascular mortality were similar to those found in high income countries. However, due

to these differences in transformation, the extent of the bias found with the variants is not directly comparable.

Finally, to demonstrate how the unequally distributed strength of age, period and cohort affect the bias, we perform two sets of simulations in which we vary these strengths. In the first set, we vary the size of the period effect from 0 to 100% in 20% increments, while correspondingly reducing the size of the age effect, and keeping the size of the cohort dimension constant at 10% (excluding the last increment, where cohort is necessarily set to 0%). The second set is identical, but then the cohort effect size is varied and period kept constant at 20%. In both sets, bias is generated by removing the first mediator ('unmeasured') from the estimation model in the 'simple' scenario. These simulations were done with probit, logit and linear regression variants. The logit and linear regression variants, plus a third set of simulations in which the age effect size is varied, are described in the supplemental material.

Simulation: estimation

Estimation of the APC effects according to the mechanism-based approach consists of fitting separate regression models for CVD death as function of A, C and the P-mediators and of each of the mediators as functions of P. In all estimation models A, P and C are treated as categorical variables through dummy coding (10 dummies for age, 4 dummies for period and 14 dummies for cohort because cohort categories were forced to overlap in order to maintain the linear identity, C = P – A). By generating cohort in this way, we maintain the linear dependency between age, period and cohort and therefore follow the equal interval width definition (Luo and Hodges 2016).

Following how we generated the variables in each variant, in the probit variant we use probit regression models for CVD death and logistic regression models for the logistic variant.

The effects of P (and A and C when appropriate) on continuous mediators are estimated using linear regression and for binary mediators using logistic or probit regression models. In all three scenarios, we first perform our estimation under entirely correct assumptions; i.e. in the second scenario we also model the additional paths from A and C to their mediators (as described in the supplementary material), and in the third scenario we also control for the confounder. Moreover in all scenarios the parametric forms used to fit these models is the same as that used to generate the data. Then, to investigate the effect of incorrect assumptions, in the second scenario we omit to model the path from A and C to the mediators, and in the third scenario we do not control for confounding. Additionally, to explore the effect of including an incomplete set of mediators, in all three scenarios we first remove the 'unmeasured' mediators from the estimation model, then 'BMI', followed by 'smoking', and finally we remove all period mediators (i.e. fit an age-cohort model). For completeness we also report the results of adopting a more traditional APC-approach in the supplemental material.

Results obtained from each scenario, variant and model specification are summarised as means of each parameter's estimates over the 1000 simulations, which we compare to the estimated values obtained from the correctly specified models to estimate the bias. Due to the very large sample size, these estimates obtained from the correct models can be interpreted as the true values.

Extending the mechanism-based approach through Monte Carlo integration

When a model with a general non-linear link function is used for a mediator or for the outcome (or both), e.g. Poisson or logistic regression, or if the models include product terms or other non-linearities, then the path tracing method cannot be used (Mulaik 2009). A different approach is then required.

The basic intuition of our approach is as follows. First, similar to the traditional approach, we estimate individual relations between age, period, cohort and mediators, and then between mediators and outcome (step 1 and step 2, respectively). The difference now being that the statistical models used for these steps are allowed to have non-linear functional forms. However, this enhanced flexibility comes at a price; the traditional multiplication of coefficients along pathways is no longer possible. Therefore, Monte Carlo integration is used instead (Robert and Casella 2004); the coefficients from step 1 and 2 are used in step 3 and 4 to generate a new dataset that does not suffer from APC linear dependency (*Pearl's front door* criterion makes this possible) but that reflect the original data structure (more details below). In step 5, an APC model is then fitted to this newly created dataset to provide estimates of the effects of age, period and cohort. Our approach is, in many ways, analogous to the Monte Carlo estimation of the parametric G-formula (Keil et al. 2014, Hernán and Robins 2013). A similar approach has also been suggested in mediation analysis (VanderWeele 2015).

We treat age, period and cohort as continuous variables to simplify the presentation below, but in our simulations we model age, period and cohort as categorical variables. Furthermore, we describe here only the case where there is one mediator for period, and where the mediator has only one cause. See the supplementary material for a description of a more general setting. We proceed as follows:

Step 1) Mediator estimation: fit a model for the mediator. If it is continuous, we can use linear regression, e.g.

$$M = \gamma_0 + \gamma_1 \cdot P + \nu$$

Where we assume $\nu \sim N(0, \sigma_M^2)$ Note that the assumption on the distribution of the error terms in the linear regression model is non-trivial if there are non-linearities involving M in the model for Y. If instead a mediator is binary, we can use logistic regression, e.g.

$$logit\{E(M|P)\} = \gamma_0 + \gamma_1 \cdot P$$

Let $(\hat{\gamma}_0, \hat{\gamma}_1)$ be the estimates of (γ_0, γ_1) from the appropriate model. If the mediator is continuous, also save the estimate of the error variance $\hat{\sigma}_M^2$. These estimates will be used in step 3.

Step 2) Outcome estimation: fit a model for the outcome. If the outcome is continuous, fit a model using linear regression, e.g.

$$Y = \delta_0 + \delta_1 \cdot A + \delta_2 \cdot C + \delta_3 \cdot M + \xi$$

where $\xi \sim N(0, \sigma_Y^2)$. Whereas if the outcome is binary, use logistic regression, e.g.

$$logit\{E(Y|A, M, C)\} = \delta_0 + \delta_1 \cdot A + \delta_2 \cdot C + \delta_3 \cdot M$$

Let $(\hat{\delta}_0, \hat{\delta}_1, \hat{\delta}_2, \hat{\delta}_3)$ be the estimates of $(\delta_0, \delta_1, \delta_2, \delta_3)$ from the appropriate model. If the outcome is continuous, also save the estimate of the error variance $\hat{\sigma}_Y^2$. These estimates will be used in step 4.

Step 3) Mediator simulation: for each of a range of period values $\tilde{p}$, simulate the mediator $\widetilde{M}(\tilde{p})$. The values $\tilde{p}$ could be randomly generated for example using a discrete uniform distribution, but its range should be equal to the range empirically observed in the data that were used for estimation. For example, if we have data ranging from 1990 to 2015, then that would be the range of values for $\tilde{p}$ that we use. If instead we categorised this into 5-year periods (1990-1994, 1995-1999, etc.), then we generate values of $\tilde{p}$ corresponding to these categories. If a mediator is continuous, use the estimates of the linear regression model in step 1 to simulate

$$\widetilde{M}(\tilde{p}) = \hat{\gamma}_0 + \hat{\gamma}_1 \cdot \tilde{p} + \hat{v}$$

Where $\hat{v}$ is randomly drawn from $N(0, \hat{\sigma}_M^2)$. If instead the mediator is binary, use the estimates from the logistic regression model in step 1 to simulate $\widetilde{M}(\tilde{p})$ from a Bernoulli distribution with mean

$$\hat{\mu}_{m,\tilde{p}} = \frac{\exp(\hat{\gamma}_0 + \hat{\gamma}_1 \cdot \tilde{p})}{1 + \exp(\hat{\gamma}_0 + \hat{\gamma}_1 \cdot \tilde{p})}$$

The number of $\widetilde{M}(\tilde{p})$ values to be simulated need not be equal to the number of observations in the data, as long as the entire empirical range is covered, but the more values we simulate, the less our final estimates will be affected by Monte Carlo error. The values of $\widetilde{M}(\tilde{p})$ will be used in step 4.

Step 4) Outcome simulation: for each of a range of age, period and cohort values $(\tilde{a}, \tilde{p}, \tilde{c})$, simulate the potential outcome $\tilde{Y}(\tilde{a}, \tilde{p}, \tilde{c})$. Since $\tilde{p}$ is already generated in step 3, they can be re-used instead of re-generated. As previously, the range of these values should be equal to the range empirically observed in age, period and cohort in the data respectively. However, we choose $\tilde{a}$, $\tilde{p}$ and $\tilde{c}$ independently, i.e. the identity $\tilde{a} = \tilde{p} - \tilde{c}$ should not hold. If the outcome is continuous, then using the linear regression estimates of step 2, and the simulated mediator values of step 3, simulate

$$\tilde{Y}(\tilde{a}, \tilde{p}, \tilde{c}) = \hat{\delta}_0 + \hat{\delta}_1 \cdot \tilde{a} + \hat{\delta}_2 \cdot \tilde{c} + \hat{\delta}_3 \cdot \widetilde{M}(\tilde{p}) + \hat{\xi}$$

Where $\hat{\xi}$ is randomly drawn from $N(0, \hat{\sigma}_Y^2)$ and $\widetilde{M}(\tilde{p})$ is as generated in step 3. If instead the outcome is binary, use the estimates from the logistic regression model in step 2, and the simulated mediator values of step 3, to simulate $\tilde{Y}(\tilde{a}, \tilde{p}, \tilde{c})$ from a Bernoulli distribution with mean

$$\hat{\mu}_{y,\tilde{a},\tilde{p},\tilde{c}} = \frac{\exp(\hat{\delta}_0 + \hat{\delta}_1 \cdot \tilde{a} + \hat{\delta}_2 \cdot \widetilde{M}(\tilde{p}) + \hat{\delta}_3 \cdot \tilde{c})}{1 + \exp(\hat{\delta}_0 + \hat{\delta}_1 \cdot \tilde{a} + \hat{\delta}_2 \cdot \widetilde{M}(\tilde{p}) + \hat{\delta}_3 \cdot \tilde{c})}$$

Step 5) Estimate age, period and cohort effects using the simulated values $\tilde{Y}(\tilde{a}, \tilde{p}, \tilde{c})$. If $\tilde{Y}\ (\tilde{a}, \tilde{p}, \tilde{c})$ was generated as a continuous variable, use linear regression with $\tilde{a}$, $\tilde{p}$ and $\tilde{c}$ as covariates. If instead $\tilde{Y}\ (\tilde{a}, \tilde{p}, \tilde{c})$ was generated as a binary variable, use logistic regression with $\tilde{a}$, $\tilde{p}$ and $\tilde{c}$ as covariates. These models will be identifiable because $\tilde{a}$, $\tilde{p}$ and $\tilde{c}$ should have been

chosen independently. The estimated parameters can be interpreted as the age, period and cohort effects from the causal model (Eq. 1).

Step 6) Use the non-parametric bootstrap to estimate the standard errors for the parameter estimates (Efron and Tibshirani 1993). This step consists of re-sampling with replacement from the original data, a dataset of equal size, and then repeating steps 1 through 5, and saving the parameter estimates of the effects of $\tilde{a}$, $\tilde{p}$ and $\tilde{c}$ at the end of step 5. The sample standard deviations of the distributions of the $\tilde{a}$, $\tilde{p}$ and $\tilde{c}$ effect estimates can be used as estimates of the standard errors, or for example the empirical 2.5% and 97.5% quantiles of these distributions can be used to derive 95% confidence intervals directly (with improvements such as 'bias-corrected and accelerated' intervals to be recommended).

The supplemental material includes R-code demonstrating the application of this technique. In linear settings, this technique results in findings identical to the traditional path tracing method, as long as a sufficiently large number of Monte Carlo simulations are carried out.

Linear dependency and expected bias

While age, period and cohort could be pairwise independent, the three variables together have a deterministic relationship: A = P – C. This linear dependency will determine the (direction of the) bias when the mechanism-based approach is used and not all the mediators are included in the estimation model. We demonstrate this here.

In a linear setting, the causal model described in Eq. 1. could be expressed as

$$Y(a,p,c) = \eta^* + \alpha^* \cdot a + \beta^* \cdot p + \theta^* \cdot c + \varepsilon^*$$

Where α*, β* and θ* represent the effects of age, period and cohort, η^* the intercept and ε^* the mean-zero error terms. This model corresponds to an (unidentified) associational model:

$$Y = \eta + \alpha \cdot A + \beta \cdot P + \theta \cdot C + \varepsilon \qquad (3)$$

Since P = A + C, we can rewrite Eq. 3. as

$$Y = \eta + \alpha \cdot A + \beta \cdot (A + C) + \theta \cdot C + \varepsilon \quad \Leftrightarrow \quad Y = \eta + (\alpha + \beta) \cdot A + (\theta + \beta) \cdot C + \varepsilon \qquad (4)$$

Eq. 4. shows that if we fit an age-cohort model (omitting period) and interpret the coefficients of A and C as representing the age and cohort effect respectively, the period effect is attributed to age and cohort in equal parts, thereby biasing the presumed age and cohort effects in the direction of the period effect. In the same way, since A = P - C, and C = P - A, we have

$$Y = \eta + (\beta + \alpha) \cdot P + (\theta - \alpha) \cdot C + \varepsilon \qquad (5)$$

$$Y = \eta + (\alpha - \theta) \cdot A + (\beta + \theta) \cdot P + \varepsilon \qquad (6)$$

when fitting period-cohort and age-period models, respectively. As Eq. 5. and Eq. 6. show, in the period-cohort model, the period parameter is biased in the direction of the age effect, while the cohort parameter is biased in the opposite direction. Similarly, in the age-period model, the period parameter is biased in the direction of the cohort effect, while the age parameter is biased

in the opposite direction. Of course, in all three of these models, the effect of the omitted variable would also be biased (unless its effect is null), as its effect estimate is effectively set to 0.

The same logic applies in the situation where we use mediators to estimate the effects of age, period or cohort, while estimating the other two effects directly. Consider the earlier described example where we used two mediators on the period path, while the effects of age and cohort are estimated directly (Figure 1). However, this time, we omit the variable M_2 from our estimation model, perhaps because it was not measured (Figure 3, left). When fitting a model for the outcome (Y) using A, C and the measured period mediator (M_1) as explanatory variables, there will be an association between A and the period mediators and between C and the period mediators due to the aforementioned linear dependency (Figure 3, right). Since the model for Y includes M_1 (i.e. we condition on M_1) the paths from A or C to Y via M_1 are blocked. However, since we do not condition on the unmeasured mediator (M_2), the pathways from A and C to Y via M_2 are not blocked, and hence their regression parameters will be biased by the contribution of the additional paths from respectively A and C via M_2 (i.e. $\delta_4 \cdot \gamma_3$ for both).

Results

Scenario 1: 'Simple'

Removing mediators from the estimation of the period path introduces bias (Figure 4 for probit variant, see supplementary material eFigure 3 for logistic variant). Initially, the magnitude of the bias is relatively weak, but becomes stronger as more mediators are removed. The relative magnitude appears strongest for birth cohort, which had weak negative parameter estimates when the model was correctly specified, but these estimates become strongly negative when mediators are removed from the estimation model. The directions of the bias are as expected (see section 'linear dependency and expected bias'). On average, removing the 'unmeasured' period mediators has a negative effect on the estimated age and cohort effects, while it introduces a positive bias on the period estimates; the estimated parameters for the age and cohort effects become less positive, while those for the period effect become less negative. The same occurs when smoking and statins are additionally removed. This is as expected, because the paths via the 'unmeasured' period mediators, smoking, and statins all have a negative effect on the outcome (Table 1). Removing BMI leads to the opposite, i.e. it introduces a positive bias on the estimated effects of age and cohort, and a negative one on that of period. Again, this is as expected because the path from P to Y through BMI is positive (Table 1). Finally, removing all mediators (i.e. fitting an AC model) results in the largest amount of bias (Figure 4).

Scenario 2: 'More causes'

In the scenario where age, in addition to period, is a cause of BMI, and cohort, in addition to period, is a cause of smoking, we find that not including these relationships in the estimation model also results in bias. As expected, this bias is largest for age, where the age effect is overestimated when the negative effect of age on BMI is not included in the estimation model

(Figure 5). A similar, but much weaker bias is found for cohort when the effect of C on smoking is not included (Figure 5 for the probit variant; supplementary material eFigure 4 for logistic variant). There is no bias in the estimation of the period effect (Figure 5).

These results follow our expectation. The effect of age on the outcome via BMI is negative. By not including age as a cause of BMI in the estimation model, this negative effect is subtracted from the total age effect, thereby resulting in an overestimation of the age effect. The effect of not modelling birth cohort as a cause of smoking follows the same logic, but the bias is weaker because the effect size from birth cohort to the outcome via smoking is also smaller. In this scenario, the period effect estimates are not biased because period is correctly modelled as a cause of the four mediators. Removing mediators from the estimation model results in biases in the same direction as found in the 'simple' scenario (see supplementary material eFigures 1 and 4), and the same explanations for these directions apply.

Scenario 3: 'Confounding'

In the scenario where there is a variable (genotype) confounding the relationships between BMI and CVD death, and between smoking and CVD death (both mediators on the P-Y path), we find that failing to control for this confounder results in bias in all three age, period and cohort effect estimates. The age parameter estimates become somewhat negatively biased when we do not control for genotype in our outcome model (Figure 6 for the probit variant; supplementary material Figure 5 for the logistic variant). The same occurs in the cohort effect, while the period effect suffers from a small positive bias (Figure 6).

The bias observed here is caused by collider stratification (Cole et al. 2010; Elwert and Winship 2014). Collider stratification bias occurs when two variables both have causal effects on a third variable (the 'collider'), and the collider is conditioned upon. In the scenario considered

here, genotype has a positive causal effect on BMI, on smoking, and on CVD mortality. Therefore, BMI and smoking are colliders in the paths between P and genotype. By including BMI and smoking (together with the other two mediators) in our outcome model, we create an additional spurious association between A and Y via A – P – BMI – genotype – Y, and one between C and Y via the path C – P – smoking – genotype – Y. Both of these spurious associations are negative because those induced by collider stratification, P-BMI-genotype and P-smoking-genotype, are both negative while A – P, C – P and genotype – CVD are all positive.

Varying effect sizes

Increasing the size of the period effect, and keeping the size of the cohort effect constant, resulted in increased bias when the 'unmeasured' mediator was removed from the estimation model (Figure 7 for the probit variant, eFigures 6 and 7 in the supplemental material for the logistic and continuous variants). Keeping the period effect constant and increasing the size of the cohort effect, resulted in roughly equal amounts of bias in the cohort effect estimate (Figure 7 and eFigures 6 and 7 in the supplemental material); the bias was equal in the continuous variant, as that estimation is done without link function transformations.

Discussion

We assessed the performance of mechanism-based APC models in realistic settings, and extended the method to incorporate all generalized linear models, as well as non-linearities in the predictors, using Monte Carlo simulation. We found that, in simple scenarios where there is a single set of mediators for P which is not affected by unmeasured confounding, the mechanism-based approach (extended or not) performed reasonably well for the estimation of all three age, period and cohort effects, especially if the mediators that were not included in the estimation had opposite signs (e.g. 'unmeasured' and 'BMI') so that their unmeasured contributions (partially) cancelled out. The bias we observed was in line with our expectations derived from the APC linear dependency. In the scenarios with additional complications, i.e. either when mediators were caused by more than one of three APC components, or there was unmeasured confounding of the mediator-outcome relationships, we found additional bias. This, again, was in line with our expectations. Findings were similar for the probit and logistic variants, though we did not directly compare these methods due to their differences in transforming effects into probabilities.

(Un)testable assumptions

APC models solve the linear dependency problem by imposing modelling constraints by fiat (Fienberg 2013). The mechanism-based approach does not differ from this, as it is (commonly, see next paragraph) not possible to identify from the data whether a variable is a mediator for age, period, cohort, a combination of these two, or all three variables. This is, of course, the same problem that occurs in any conventional APC analysis when attempting to decompose some outcome into age, period and cohort effects. Therefore, the untestable assumptions that are made in conventional APC analysis move to the mediator stage of the modelling procedure.

It is possible to statistically test the addition of mediators to the APC model (Winship and Harding 2008). We have omitted this test from our assessment because the test is conditional on having (at least) a full set of mediators for one of the three APC variables. We consider it likely that in the majority of real-life applications it will not be possible to find a full set of mediators for even one of the APC variables, and hence have chosen to focus our assessment on possible biases that may be encountered, such as due to missing mediators.

Simulating bias

In our simulations, we considered 4 mediators on the period path, and kept the 'more causes' scenario separate from the 'confounding' scenario. This was done in order to illustrate, separately, possible pitfalls which may be encountered when the mechanism-based approach is used. In a real application of the method, many more mediators may exist, which may also have more than one cause, and their relation with the outcome may be confounded. However, a more complicated scenario need not necessarily result in more biased estimates because biases of opposed sign may cancel each other out, such as when BMI was removed from the estimation model in our simulations (since BMI was the only positive mediator). Most importantly, when assessing potential bias, it is the size of the omitted pathways that matters, relative to the size of the included pathways. This was demonstrated in the analysis where we varied the size of the period effect and omitted a period mediator from the estimation model; a larger period effect resulted in a larger bias. Whereas keeping the size of the period effect constant and increasing the size of the cohort effect resulted in roughly the same magnitude of bias and smaller relative bias. Analogous simulations where instead age or cohort mediators had been removed would have yielded the same conclusions; only the sign of the bias would differ as shown in the section on linear dependency and expected bias.

Confounding and colliding

Because the mechanism-based approach uses a mediation approach to APC analysis, it is also subject to biases that are not present in traditional APC analysis. These are confounding of the mediator-outcome relationship, and collider stratification bias (Elwert and Winship, 2014, Greenland 2003).

Arguably, traditional APC analysis is not affected by confounding because age, period and cohort are time dimensions and therefore not causally affected by other variables. In our confounding scenario, the effect of BMI (and smoking) on CVD mortality was confounded by Genotype, which thereby also affected the estimation of the period effect. Ideally, such a confounder would be controlled. However, if the confounder is unmeasured, and therefore cannot be controlled, a difficult choice has to be made. Removing the mediator induces bias via an omitted pathway, whereas including the mediator induces confounding bias. Confounding of the mediator-outcome relationship (and hence collider stratification bias) resulted in more bias than omitting the relevant pathway in our simulations, but this is dependent on the relative strengths of mediation and confounding.

Collider stratification bias occurs when two variables both have causal effects on a third variable (the 'collider'), and the collider is conditioned upon (Elwert and Winship, 2014, Greenland 2003). This is also known as endogenous selection. Doing so creates an artificial association between the two causes of the collider which is of sign opposite to the product of the signs of the effects into the collider. In traditional APC analysis, collider stratification does not occur because APC effects on an outcome are estimated without adjusting for confounders or mediators. In the "Confounding" simulations, BMI (and smoking) was affected by Period and by Genotype. Due to the linear dependency of period with age and cohort, collider stratification bias

also affected the age and cohort estimates. Also in this scenario, conditioning on Genotype would have prevented this bias. If conditioning on both causes of the collider is not possible, a choice has to be made between removing and keeping the mediator in the analysis. Greenland's (2003) second 'rule of thumb' regarding endogenous selection states that including mediators which are also colliders induces larger bias than the size of the mediator's dependence on its causes (Elwert and Winship, 2014). This suggests removing such mediators from the analysis, despite inducing bias via omitted pathways.

Directed Acyclic Graphs

We encourage investigators to draw causal DAGs of the relations between APC variables, mediators and outcome. To represent the deterministic relationship between age, period and cohort in the causal DAG, we used bold arrows, which is in line with conventions in causal inference (Spiegelhalter et al. 2002; Spirtes et al. 2001) but differs from the representation used by Winship and Harding (2008). We used these arrows because the relationship between age, period and cohort is fundamentally different from that of other relationships in the DAG (which are stochastic and causal, rather than deterministic).

By drawing the relations between APC and mediators, we are clear about the assumptions underlying our analyses. This sets the mechanism-based approach apart from other APC approaches, where transparency about constraints can be lacking (Luo 2013). Drawing causal DAGs, and re-drawing once one of the three age, period or cohort variables is removed (as was done in Figure 3) helps explicate otherwise hidden assumptions about the relationships between age, period, cohort and their mediators, and can help identify possible biases such as collider stratification bias.

Guidelines for application of the mechanism-based approach

Based on our assessment, when estimating age, period and cohort effects by applying the mechanism-based approach (extended or not), we suggest considering the following questions:

a) Which APC variable is believed to have the weakest effect on the outcome in question?
b) Which APC variable has putative mediators available in the data?
c) How well-measured are these mediators?
d) How exhaustive are these mediators of all the pathways linking the respective APC variable and the outcome?
e) Are there common causes of these mediators (e.g. is more than one APC variable affecting the mediators)?
f) Are mediator-outcome relationships potentially confounded, and how severely?

To minimize bias, select the APC variable for which the answers to questions (a) through (f) are most favourable. For this variable, model the mediated pathways while the other two effects can be modelled directly. The degree of bias that remains in the analysis is dependent on the answers to the above questions for the chosen APC variable, and dependent on 'ordinary' modelling concerns such as correct modelling of the functional form of the relationship between APC variables, mediators and outcome.

Answering questions (a) through (f) will likely require making assertions which cannot be tested in the data, and is thereby similar to setting constraints by fiat like in various other APC approaches. However, the difference is that the answers to these questions, particularly questions (a), (d), (e), and (f), can be motivated based on theoretical reasoning. As described in our introduction, various authors have argued that age-period-cohort analysis needs to become more theoretically informed. A large suite of APC methods exist that solve the linear identification problem non-transparently and without substantive theoretical justification. A strength of the

mechanism-based approach is that it motivates researchers to be explicit about their (substantive theoretical) assumptions.

Conclusion

We demonstrated the performance of the mechanism-based approach to APC modelling in non-ideal circumstances. Biases occurred when the assumed causal relations did not coincide with the truth, such as when paths from causes to mediators were omitted from the estimation model, or if there was unmeasured confounding of the mediator-outcome relationship. The direction of the bias followed our expectations based on APC linear dependency. Size of bias is dependent on the size of the effects involving confounders, omitted intermediate variables or pathways. Our extension of the mechanism-based approach increases its utility, by allowing it to be easily useable for models with non-linear link functions and parameterizations of any complexity. Our brief guidelines, aided by causal DAGs, offer a useful tool for researchers who wish to implement this approach.

References

Bell, A., Jones, K. (2014). Another 'futile quest'? A simulation study of Yang and Land's Hierarchical Age-Period-Cohort model. *Demographic Research*, 30, 333-360.

Ben-Shlomo, Y., Kuh, D. (2002). A life course approach to chronic disease epidemiology: conceptual models, empirical challenges and interdisciplinary perspectives. *International Journal of Epidemiology*, 31 (2), 285-293.

Blackmore, H.L., Ozanne, S.E. (2015). Programming of cardiovascular disease across the life-course. *Journal of Molecular and Cellular Cardiology*, 83, 122-130.

Capewell, S., Beaglehole, R., Seddon, M., McMurray, J. (2000). Explanation for the decline in coronary heart disease mortality rates in Auckland, New Zealand, between 1982 and 1993. *Circulation*, 102 (13), 1511-1516.

Clayton, D, Schifflers, E. (1987). Models for temporal variation in cancer rates. II: age-period-cohort models. *Statistics in Medicine*, 6, pp. 469–481.

Cole, S.R., Platt, R.W., Schisterman, E.F., Chu, H., Westreich, D., Richardson, D., Poole, C. (2010). Illustrating bias due to conditioning on a collider. *International Journal of Epidemiology*, 39 (2), 417-420.

Efron B., Tibshirani R.J. (1994). An introduction to the Bootstrap. CRC Press: Boca Raton.

Ekamper, P., van Poppel F., Stein, A.D., Lumey, L.H. (2014). Independent and additive association of prenatal famine exposure and intermediary life conditions with adult mortality between age 18-63 years. *Social Science and Medicine,* 119, 232-239.

Elwert, F., Winship C. (2014). Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable. *Annual Review of Sociology,* 40, 31-53.

Fienberg, SE. (2013). Cohort Analysis' Unholy Quest: A Discussion. *Demography*, 2013, 50 (6), 1981-1984.

Glymour, MM. (2006). Chapter 16: Using causal diagrams to understand common problems in social epidemiology. In J.M. Oaks and J.S. Kaufman (eds.), *Methods in Social Epidemiology* (pp. 387-422). San Francisco, CAL: Jossey-Bass.

Greenland S. (2003). Quantifying biases in causal models: classical confounding versus collider-stratification bias. *Epidemiology* 14,300–306.

Hernán M.A., Robins J.M. Causal Inference. 2013 Boca Raton, Fla Chapman & Hall/CRC

Holford, T.R. (2006). Approaches to fitting age-period-cohort models with unequal intervals. *Statistics in Medicine,* 25 (6), 977-993.

Held, L., Riebler, A. (2012). A conditional approach for inference in multivariate age-period-cohort models. *Statistical Methods in Medical Research,* 21,311-329.

Jiang, Z., VanderWeele T.J. (2015). Jiang and VanderWeele Respond to "Bounding Natural Direct and Indirect Effects". *American Journal of Epidemiology,* 182 (2), 115-117.

Keil, A.P, Edwards, J.K., Richardson, D.B., Naimi, A.I., Cole, S.R. (2014). The parametric g-formula for time-to-event data: intuition and a worked example. *Epidemiology*, 25(6),889-97.

Kuhn, T. S. (1996). The structure of scientific revolutions (3rd ed.). Chicago, IL: The University of Chicago Press.

Luo, L. (2013). Assessing Validity and Application Scope of the Intrinsic Estimator Approach to the Age-Period-Cohort Problem. *Demography*, 50 (6), 1945-1967.

Luo, L., Hodges, J.S. (2016), Block Constraints in Age-Period-Cohort models with Unequal-Width Intervals. *Sociological Methods & Research*, 45(4), 700-726.

MacKinnon, D.P., Lockwood, C.M., Williams, J. (2004), Confidence Limits for the Indirect Effect: Distribution of the Product and Resampling Methods. *Multivariate Behavioral Research*, 39(1), 99-128

Mulaik, S. (2009). Structural equation models. In: *Linear Causal Modeling with Structural Equations* (pp. 119–138). Boca Raton, FL: CRC Press.

Pearl (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Peeters, A., Nusselder, W.J., Stevenson, C., Boyko, E.J., Moon, L., Tonkin, A. (2011). Age-specific trends in cardiovascular mortality rates in the Netherlands between 1980 and 2009. European Journal of Epidemiology 26 (5): 369–373.

Preacher, K. J., & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, *40*, 879-891.

Preston, S., Wang, H. (2006). Sex mortality differences in the United States: The role of cohort smoking patterns. *Demography*, 43, 631–646.

Robert, C, Casella, G (2004). Monte Carlo Integration. In Robert, C. (ed.), *Monte Carlo Statistical Methods* (pp. 79-122). Springer: New York.

Rubin, D.B. (1974). Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology*, 66 (5), 688–701.

Sabol, S.Z., Nelson, M.L., Fisher, C., Gunzerath, L., Brody, C.L., Hu, S., Sirota, L.A., Marcus, S.E., Greenberg, B.D., Lucas, F.R., Benjamin, J., Murphy, D.L., Hamer, D.H. (1999). A genetic association for cigarette smoking behavior. *Health Psychology*, 18 (1), 7-13.

Spiegelhalter, D.J., Thomas A, Best N.G., Lunn D. (2002). Winbugs User Manual. Version 1.4. Cambridge. England: MRC Biostatistics Unit.

Spirtes P., Glymour C., Scheines R. (2001). Causation, Prediction and Search. The MIT Press: Massachusetts: Cambridge.

Smith, J.G., Newton-Cheh, C. (2015). Genome-wide association studies of late-onset cardiovascular disease. *Journal of Molecular and Cellular Cardiology*, 83, 131-14.

VanderWeele, T.J. (2015). *Explanation in Causal Inference: Methods for Mediation and Interaction*, Oxford University Press.

VanderWeele, T.J., Hernán, M.A., Robins J.M. (2008). Causal directed acyclic graphs and the direction of unmeasured confounding bias. *Epidemiology*. 19 (5), 720-8.

Verlato, G., Melotti, R., Corsico, A.G., Bugiani, M., Carrozzi, L., Marinoni, A., Dallari, R., Pirina, P., Struzzo P, Olivieri, M., de Marco, R.; ISAYA Study Group. (2006). Time trends in smoking habits among Italian young adults. *Respiratory Medicine*, 100 (12), 2197-2206.

Winship, C, Harding, D.J. (2008). A Mechanism-Based Approach to the Identification of Age–Period–Cohort Models. *Sociological Methods & Research*, 36 (3), 362-401.

Winship, C., Mare, R.D. (1983). Structural Equations and Path Analysis for Discrete Data. *The American Journal of Sociology*, 89(1), 54-110.

Wright, S. (1934). The Method of Path Coefficients. *The Annals of Mathematical Statistics*, 5(3), 161-215.

Yang, Y., Schulhofer-Wohl, S., Fu, W. J., Land, K. C. (2008). The intrinsic estimator for age-period-cohort analysis: What it is and how to use it. *American Journal of Sociology*, 113, 1697–1736.

	Direction of effect		
Mediator	Period on mediator	Mediator on outcome	Period on outcome
BMI	Positive	Positive	Positive
Smoking	Negative	Positive	Negative
Statin	Positive	Negative	Negative
Unmeasured	Negative	Positive	Negative

Table 1. Direction of effect of period on the mediators, mediators on cardiovascular mortality, and the direction along the entire path of period on cardiovascular mortality.

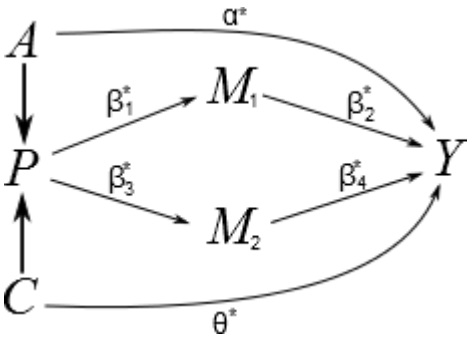

Figure 1. Causal Directed Acyclic Graph, showing the age effect (α*), cohort effect (θ*), and the period effect (the β*'s). The bold arrows represent deterministic relationships, while the non-bold arrows represent stochastic relationships.

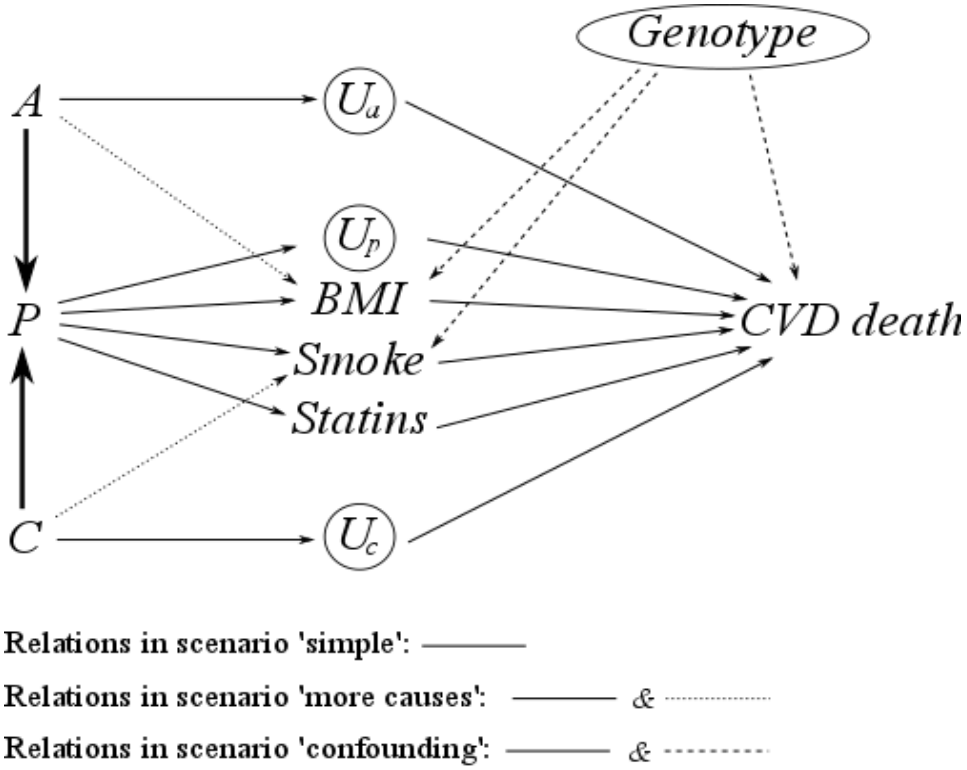

Figure 2. Causal Directed Acyclic Graph of the 3 scenarios investigated by simulation. Bold arrows between age (A), period (P) and cohort (C) represent a deterministic relationship, whereas the remaining arrows represent stochastic causal relationships. Circled variables represent variables that are omitted from the estimation model (in some simulation setups).

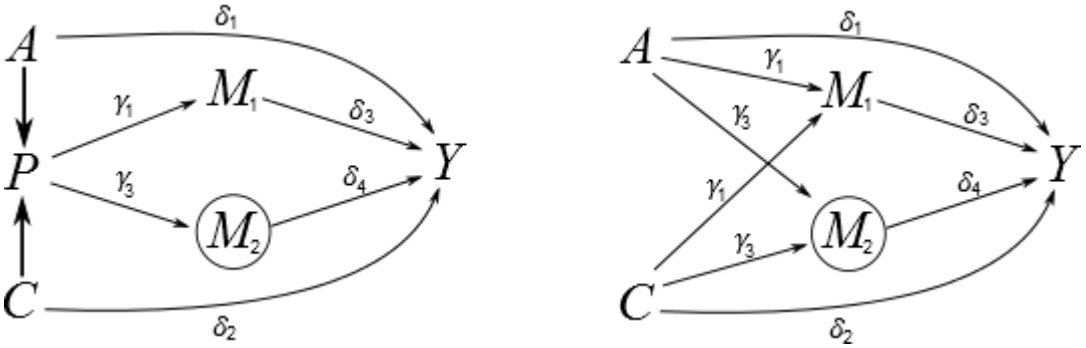

Figure 3. Causal Directed Acyclic Graph when the age effect (δ_1) and cohort effect (δ_2) are estimated directly, and the period effect is estimated using mediators as per *Pearl's front door criterion,* while one period mediator is unmeasured. Left: effect estimates if M_2 is measured and P is included in the estimation process. Right: relationships that form due to linear dependency when period is excluded from estimation.

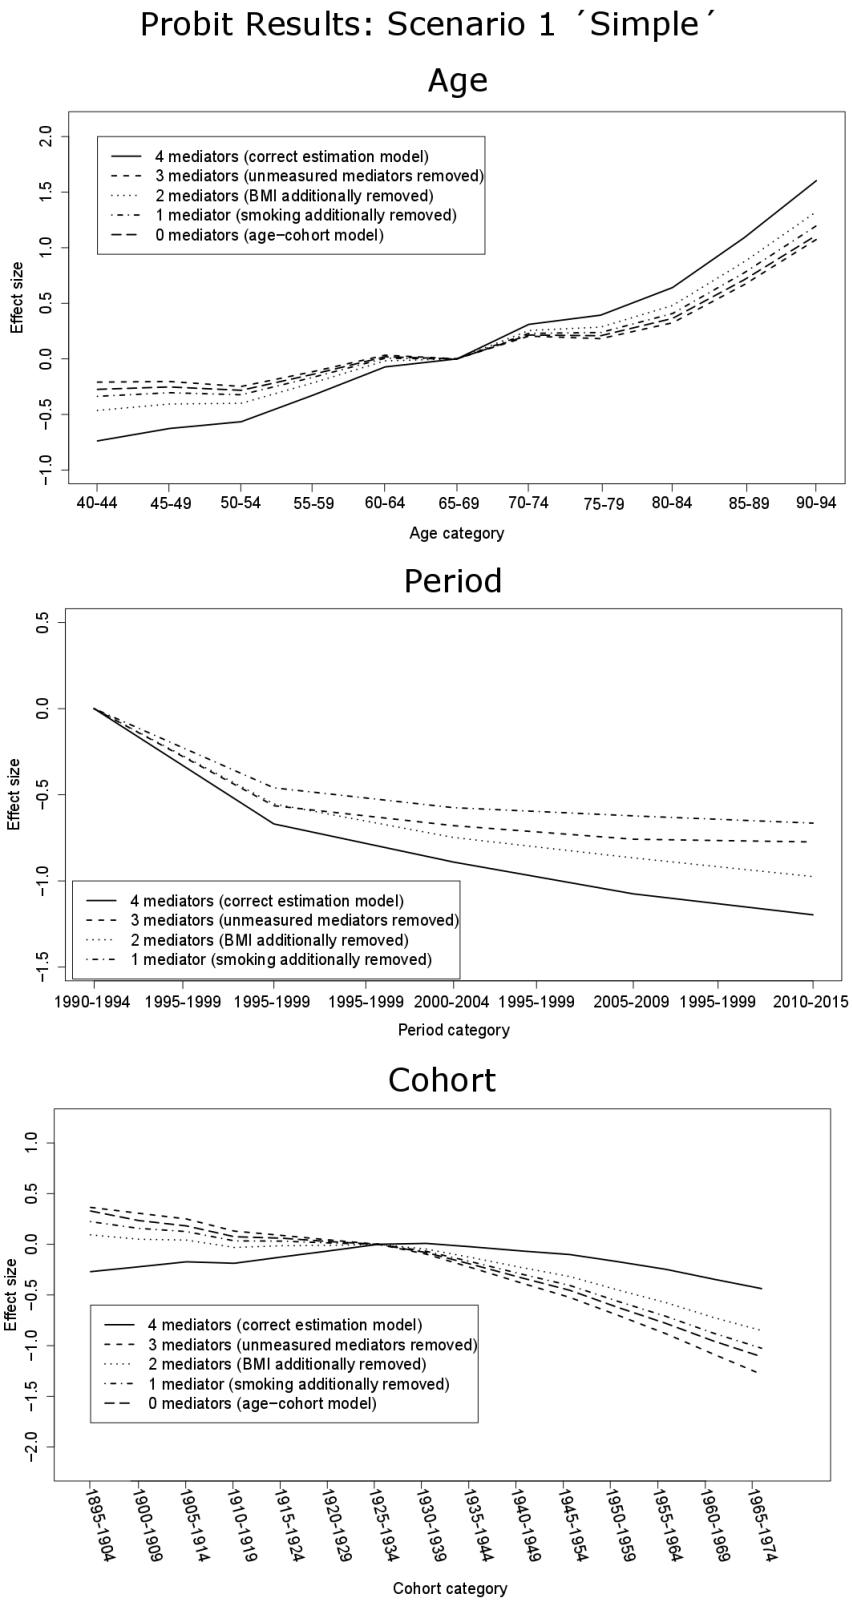

Figure 4. Average estimated parameters for the APC effects in scenario 1 ("Simple") using the mechanism-based approach: summary of 1000 simulations.

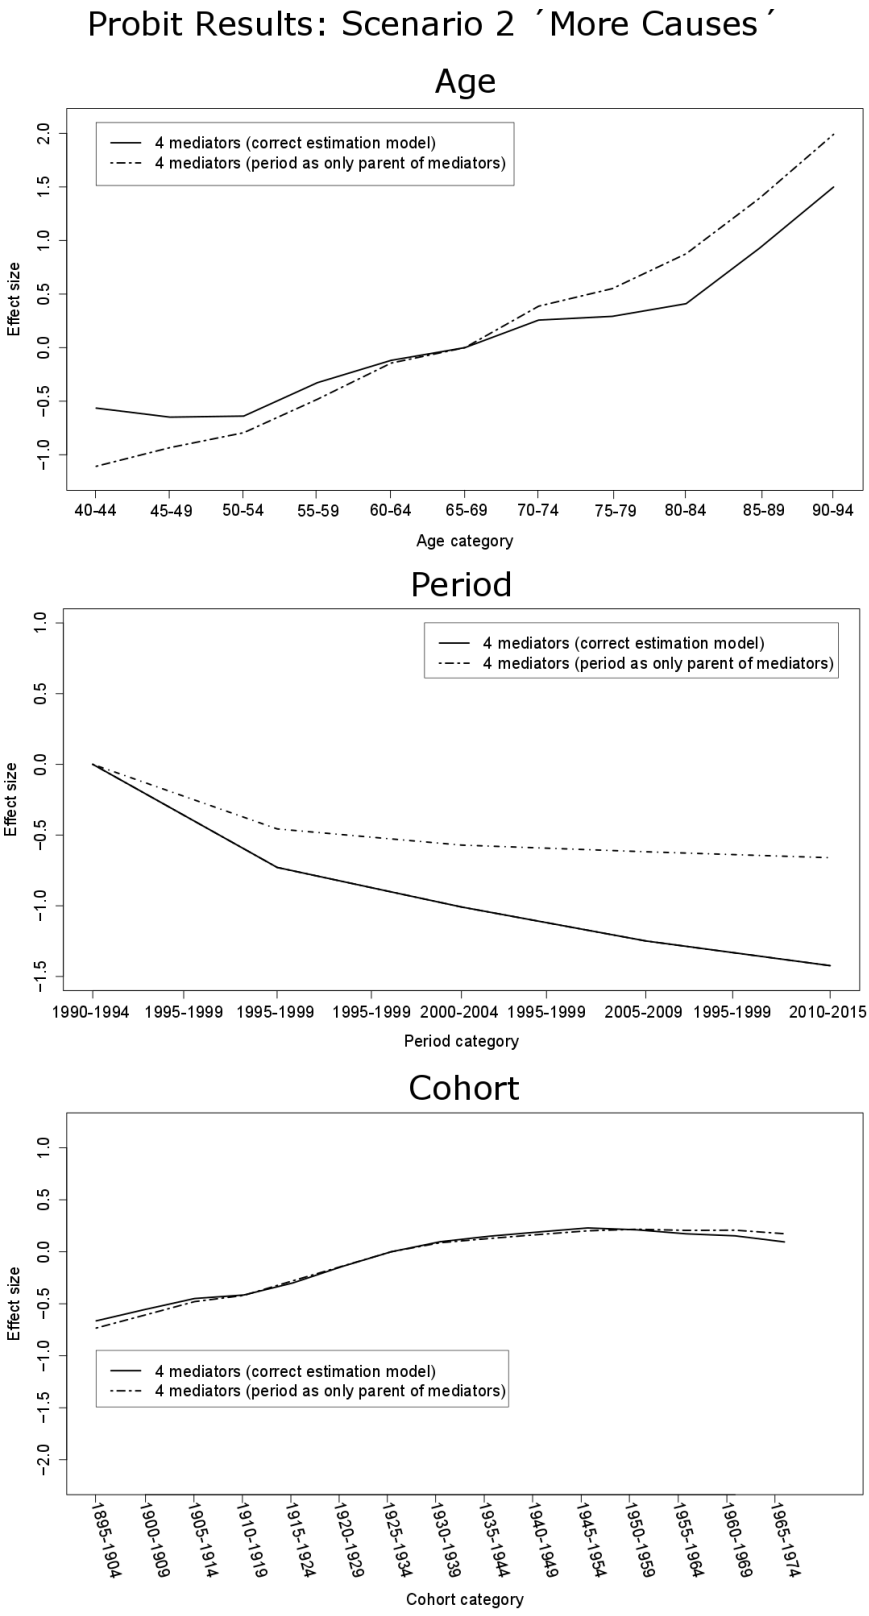

Figure 5. Average APC estimates in scenario 2 ("More causes") using the mechanism-based approach: summary of 1000 simulations.

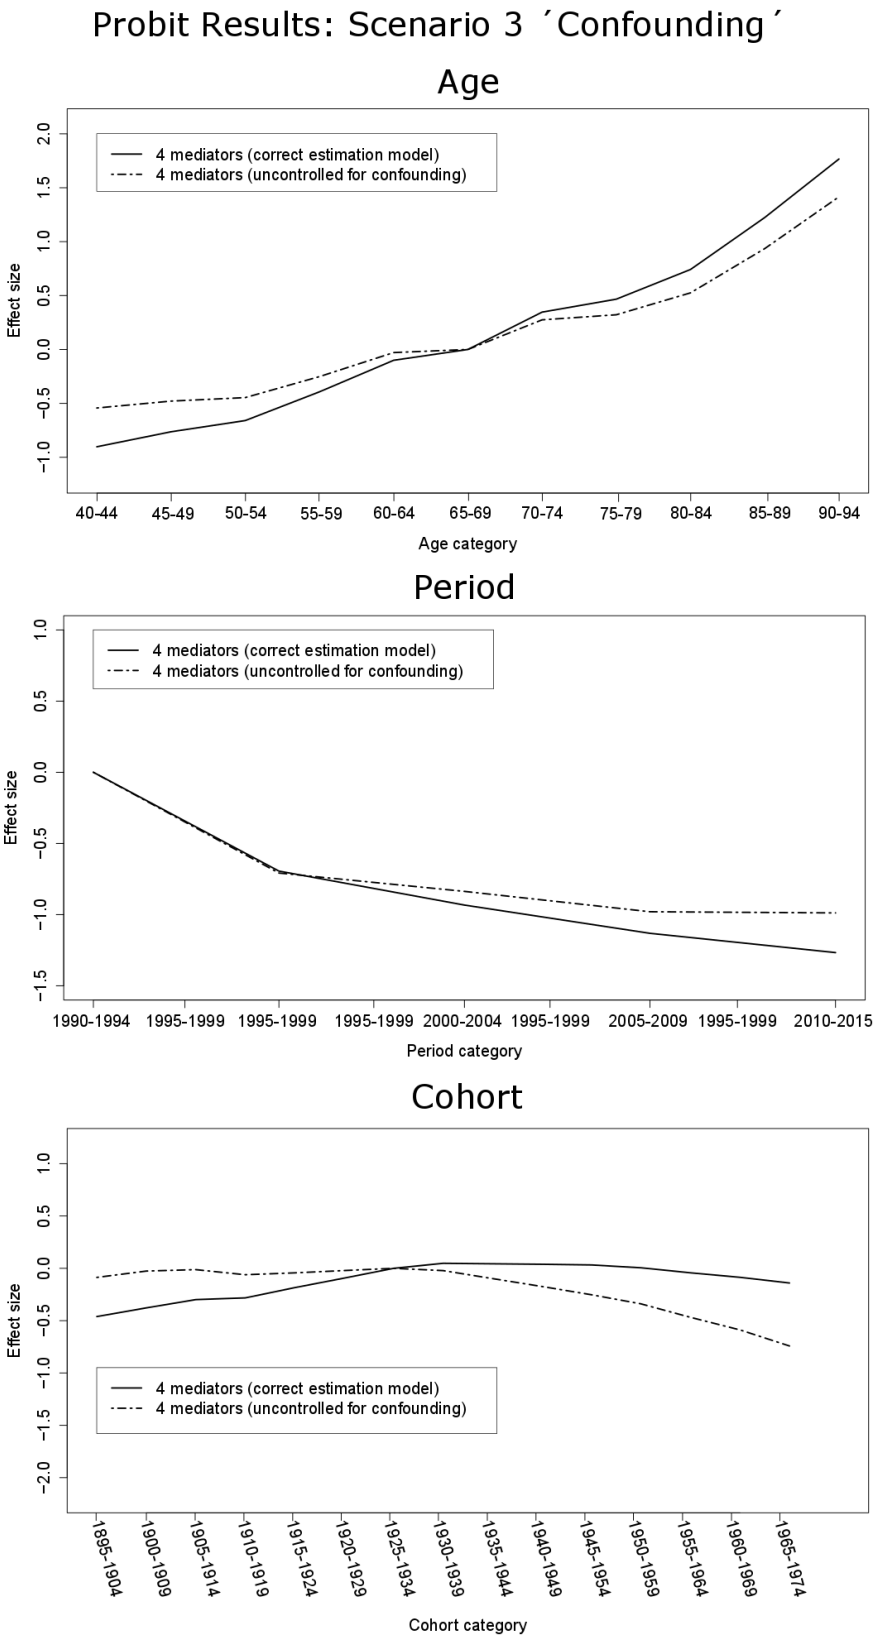

Figure 6. Average APC estimates in scenario 3 ("Confounding") using the mechanism-based approach: summary of 1000 simulations.

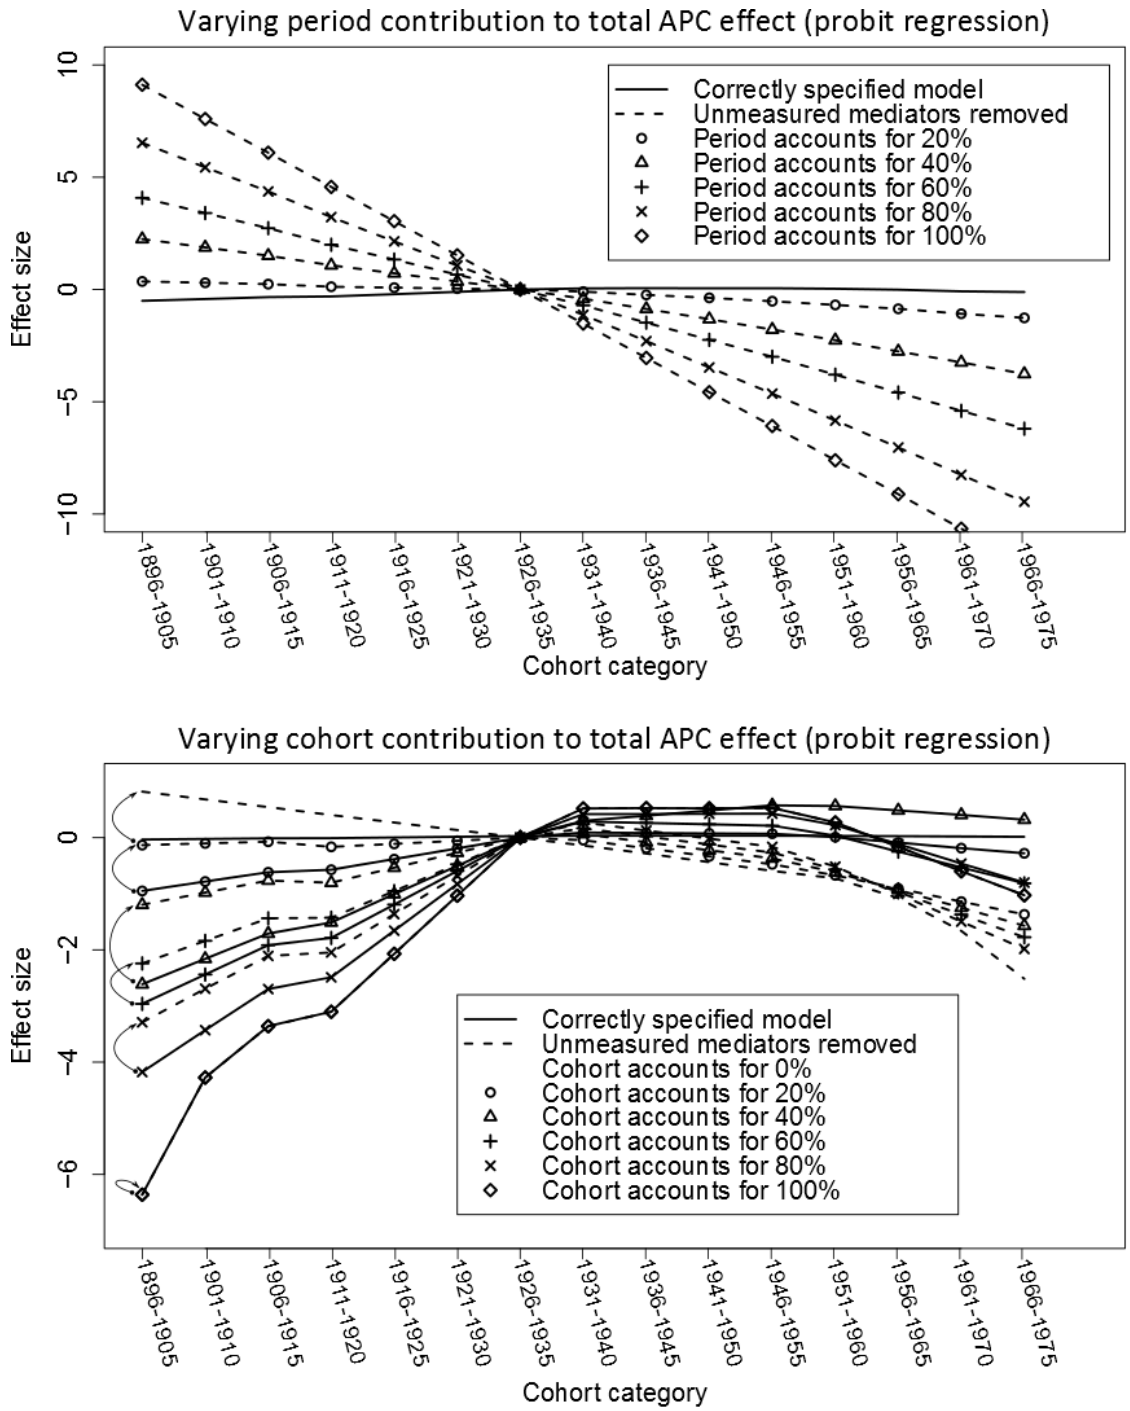

Figure 7. Varying the effect sizes of period 0 to 100% of the total APC effect in 20% increments while keeping cohort effect constant (upper), varying the effect sizes of cohort 0 to 100% of the total APC effect in 20% increments while keeping period effect constant (lower). When period accounts for 100% the correctly specified cohort trend is a horizontal line at y=0. No bias when cohort accounts for 100% because then the period effect (source of bias) accounts for 0%. Only cohort figures (probit variant) shown. Arrows in lower figure indicates the size and direction of the bias.