



Education Corner

Educational Note: Paradoxical collider effect in the analysis of non-communicable disease epidemiological data: a reproducible illustration and web application

Miguel Angel Luque-Fernandez,^{1,2,3,4,5*} Michael Schomaker,⁶
 Daniel Redondo-Sanchez,^{1,5} Maria Jose Sanchez Perez,^{1,5}
 Anand Vaidya⁷ and Mireille E Schnitzer^{8,9}

¹Biomedical Research Institute, Non-Communicable and Cancer Epidemiology Group (ibs.Granada), Andalusian School of Public Health, University of Granada, Granada, Spain, ²Department of Non-Communicable Disease Epidemiology, London School of Hygiene and Tropical Medicine, London, UK, ³Centre de Recherche en Epidemiologie, Biostatistique et Recherche Clinique Ecole de Sante Publique, Universite Libre de Bruxelles, Brussels, Belgium, ⁴Department of Epidemiology, Harvard School of Public Health, Harvard University, Boston, MA, USA, ⁵Biomedical Network Research Centers of Epidemiology and Public Health (CIBERESP), ISCIII, Madrid, Spain, ⁶Centre of Infectious Disease Epidemiology and Research, University of Cape Town, Cape Town, South Africa, ⁷Brigham and Women's Hospital, Harvard Medical School, Harvard University, Boston, MA, USA, ⁸Faculty of Pharmacy and Department of Social and Preventive Medicine, University of Montreal, Montreal, QC, Canada and ⁹Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Montreal, QC, Canada

*Corresponding author. Biomedical Research Institute of Granada (ibs.Granada), Non-communicable Disease and Cancer Epidemiology Group, University of Granada, Campus Universitario de Cartuja, C/Cuesta del Observatorio 4, 18080 Granada, Spain. E-mail: miguel.luque.easp@juntadeandalucia.es

Editorial decision 6 November 2018; Accepted 30 November 2018

Abstract

Classical epidemiology has focused on the control of confounding, but it is only recently that epidemiologists have started to focus on the bias produced by colliders. A collider for a certain pair of variables (e.g. an outcome Y and an exposure A) is a third variable (C) that is caused by both. In a directed acyclic graph (DAG), a collider is the variable in the middle of an inverted fork (i.e. the variable C in $A \rightarrow C \leftarrow Y$). Controlling for, or conditioning an analysis on a collider (i.e. through stratification or regression) can introduce a spurious association between its causes. This potentially explains many paradoxical findings in the medical literature, where established risk factors for a particular outcome appear protective. We use an example from non-communicable disease epidemiology to contextualize and explain the effect of conditioning on a collider. We generate a dataset with 1000 observations, and run Monte-Carlo simulations to estimate the effect of 24-h dietary sodium intake on systolic blood pressure, controlling for age, which acts as a

confounder, and 24-h urinary protein excretion, which acts as a collider. We illustrate how adding a collider to a regression model introduces bias. Thus, to prevent paradoxical associations, epidemiologists estimating causal effects should be wary of conditioning on colliders. We provide R code in easy-to-read boxes throughout the manuscript, and a GitHub repository [<https://github.com/migariane/ColliderApp>] for the reader to reproduce our example. We also provide an educational web application allowing real-time interaction to visualize the paradoxical effect of conditioning on a collider [<http://wat.zilei.com/shiny/collider/>].

Key words: Epidemiological methods, causality, non-communicable disease epidemiology

Key Messages

- Paradoxical associations between an outcome and exposure are common in epidemiological studies using observational data.
- A collider is a variable that is causally influenced by two other variables.
- Controlling for a collider in multivariable regression analyses can introduce a spurious association between its causes (e.g. exposure and outcome).
- Directed acyclic graphs based on existing subject-matter knowledge can help to identify colliders.
- Whether or not it is advisable to adjust for a collider depends on the main analytical objective. For instance, a predictive model may condition on a collider to increase prediction accuracy, whereas one should typically not condition on it when estimating causal effects to prevent bias.

Introduction

During the past 30 years, classical epidemiology has focused on the control of confounding.¹ It is only recently that epidemiologists have started to focus on the bias produced by colliders in addition to confounders.^{2,3} Directed acyclic graphs (DAGs) can help to visualize the assumed structural relationships between the variables under analysis. With this framework, we can distinguish between biases resulting from: (i) not conditioning on common causes of exposure and outcome (unadjusted confounding); or (ii) conditioning on common effects (collider bias).^{4,5} Epidemiologists use DAGs to determine the set of variables that are necessary to control for confounding and to summarize the subject-matter knowledge of the data-generating process. Using the DAGs terminology, variables including A (exposure) and Y (outcome) are ‘nodes’ connected by an arrow (a.k.a. directed edge), and a ‘path’ is a way to get from one node to another travelling along its arrows. The directed arrow (\rightarrow) from A to Y means that one does not exclude the possibility that A causes Y.^{6–8}

A collider for a certain pair of variables (e.g. outcome and exposure) is a third variable that is caused by both of them. In DAG terminology, a collider is the variable in the middle of an inverted fork (i.e. variable C in $A \rightarrow C \leftarrow Y$).^{6–8} Using

regression to control for a collider, or stratifying the analysis with respect to a collider, can introduce a spurious association between its causes, which can potentially introduce non-causal associations between the exposure and the outcome. This has been used to explain why the medical literature contains many paradoxical findings, where established risk factors appear protective for the outcome.^{9–12} For instance, numerous studies have reported a paradoxical protective effect of maternal cigarette smoking during pregnancy on pre-eclampsia, which has been named the pre-eclampsia smoking paradox. This paradox is due to gestational age at delivery, which is a collider between smoking (exposure) and pre-eclampsia (outcome).⁹ However, the magnitude of the resulting bias will depend on the associations between the collider and the two parent variables.

We hope that this methodological note will contribute to the increasing awareness of ‘colliders’ and an understanding of the potential magnitude of collider bias among applied epidemiologists. The remainder of this note is structured as follows:

- We review terminology related to DAGs and the rules one can follow to determine whether a causal effect is estimable.

- ii. We demonstrate the statistical structure of collider bias using a simulated dataset.
- iii. We illustrate the effect of conditioning on a collider using a realistic non-communicable disease epidemiology example (hypertension and dietary sodium intake).
- iv. We provide R code in easy-to-read boxes throughout the manuscript and in a GitHub repository: [<https://github.com/migariane/ColliderApp>].
- v. We provide readers with an educational web application allowing real-time interaction to visualize the paradoxical effect of conditioning on a collider [<http://watzilei.com/shiny/collider/>].

Statistical structure of confounding and collider bias

Review of confounding

Confounding arises from common causes of the exposure (A) and the outcome (Y). Note that in Figure 1A, both the outcome (Y) and the exposure (A) share a common ‘parent’ (direct cause). Y and A are both called ‘descendants’ of W as they are both caused by W. The confounder wholly or partially accounts for the observed association of the exposure (A) on the outcome (Y). The presence of a confounder can lead to ‘confounding bias’, and thus inaccurate estimates of the effect of A on Y. More precisely, bias means that the associational measure, for example the crude odds ratio, is different from the causal effect, such as the true marginal causal odds ratio (we give a clear definition of a marginal causal effect further below).

Figure 1A gives an example of a confounding structure, where the path $A \leftarrow W \rightarrow Y$ is called a ‘back-door path’ which is defined as any path from A to Y that starts with an arrow into A. Without conditioning on variables, a path is open when it does not contain colliders. An open back-door path can be blocked and confounding removed by

conditioning on non-colliders (via regression or stratification). In Figure 1A, conditioning on the confounder W blocks the open back-door path. A path that is blocked by a collider can be opened by conditioning on the collider.¹² To sufficiently control for confounding, epidemiologists must identify a set of variables in the DAG that block all open back-door paths from the exposure (A) to the outcome (Y) by conditioning on variables along each path (i.e. using stratification or regression). In statistical terms, being able to block all back-door paths is known as conditional exchangeability or ignorability.

To describe confounding and collider bias, we may use the expression ‘association is not causation’. This means that measures of association, such as the conditional mean difference in the case of a binary A, $E(Y|A=1, \mathbf{W}) - E(Y|A=0, \mathbf{W})$, is not identical to its marginal causal counterpart, the average treatment effect: $E(Y(1)) - E(Y(0))$. Causal effects are often formulated in terms of potential outcomes, as formalized by Rubin.¹³ Let A denote a continuous exposure, W a pre-exposure vector of potential confounders and Y a continuous outcome. Each individual has a potential outcome corresponding to any given level of the exposure, that is the outcome they would have received had they been exposed to $A = a$, denoted $Y(a)$. However, it is only possible to observe a single realization of the outcome for an individual. We may observe $Y(a)$ only for those who were exposed with $A = a$.¹³ If W is the set of confounding variables, then $Y(a) \perp A | \mathbf{W}$ refers to conditional exchangeability, where the symbol \perp means ‘independent’. It implies that (within the strata of W) the distribution of $Y(a)$ is the same regardless of the value of A that the individual actually received, i.e. $E(Y(a)|A, \mathbf{W})$ is the same regardless of the value of A that the individual actually received. We therefore have no systematic differences in how subjects would have performed, under any given exposure, which are not already explained by W.




Figure 1. Basic structural associations between exposure and outcome: confounding (A), collider (B), and M-bias (C).

Box 1. To generate data consistent with Figure 2A

```

library(visreg) # load package to visualize regression output
library(ggplot2) # load package to visualize regression output
N <- 1000      # sample size
set.seed(777)
W <- rnorm(N)      # confounder
A <- 0.5 * W + rnorm(N) # exposure
Y <- 0.3 * A + 0.4 * W + rnorm(N) # outcome
fit1 <- lm(Y ~ A) # crude model
fit2 <- lm(Y ~ A+W) # adjusted model
# visualize crude and adjusted models
visreg(fit1, 'A', gg = TRUE, line = list(col = 'blue'),
points = list(size = 2, pch = 1, col = 'black')) + theme_classic()
visreg(fit2, 'A', gg = TRUE, line = list(col = 'blue'),
points = list(size = 2, pch = 1, col = 'black')) + theme_classic()

```

Demonstration of confounding and regression adjustment

We now demonstrate adjustment for confounding via linear regression models. In [Box 1](#) we show how to generate data consistent with the DAG from [Figure 2A](#), after which we run two different regression models. The confounder W is generated as a standard normal random variable, i.e. with mean 0 ($\mu = 0$) and variance 1 ($\sigma^2 = 1$). The generation of A depends on the value of W plus an error term, and Y is generated depending on both A and W plus an error term, where both error terms have independent standard normal distributions. Note that the simulation assumes linear relationships between the variables, and that the true simulated causal effect of the exposure A on Y is 0.3 (the coefficient in the linear regression model). Then, we fit unadjusted ($fit1$) and adjusted ($fit2$: adjusted for W) linear regression models to estimate associations between A and Y . We visualize the fit of both models using the R software package `visreg`, where we used R version 3.5.1 (R Foundation for Statistical Computing, Vienna).

Note that our confounder W is the only variable that does not have parents in [Figure 1A](#), i.e. it is not caused by any variable in the DAG. Therefore, in the code, it is the only variable that is generated independently of the other variables in the model. However, both A and Y depend on a common cause W (their parent) which is the source of the open back-door path between A and Y . As an

illustration of the confounding bias due to W , [Table 1](#) (columns 1, 2) shows the coefficients of A and W from the fitted regression models. The first regression does not condition on W and therefore has an upwards bias in the coefficient of A (0.471). However, the second regression closes the open back-door path by including the confounder W in the regression model. Thus, it estimates the causal effect as 0.289, close to the true coefficient (0.3) ([Figure 2A](#), [Table 1](#): columns 1, 2), the residual difference being entirely due to sampling variability.

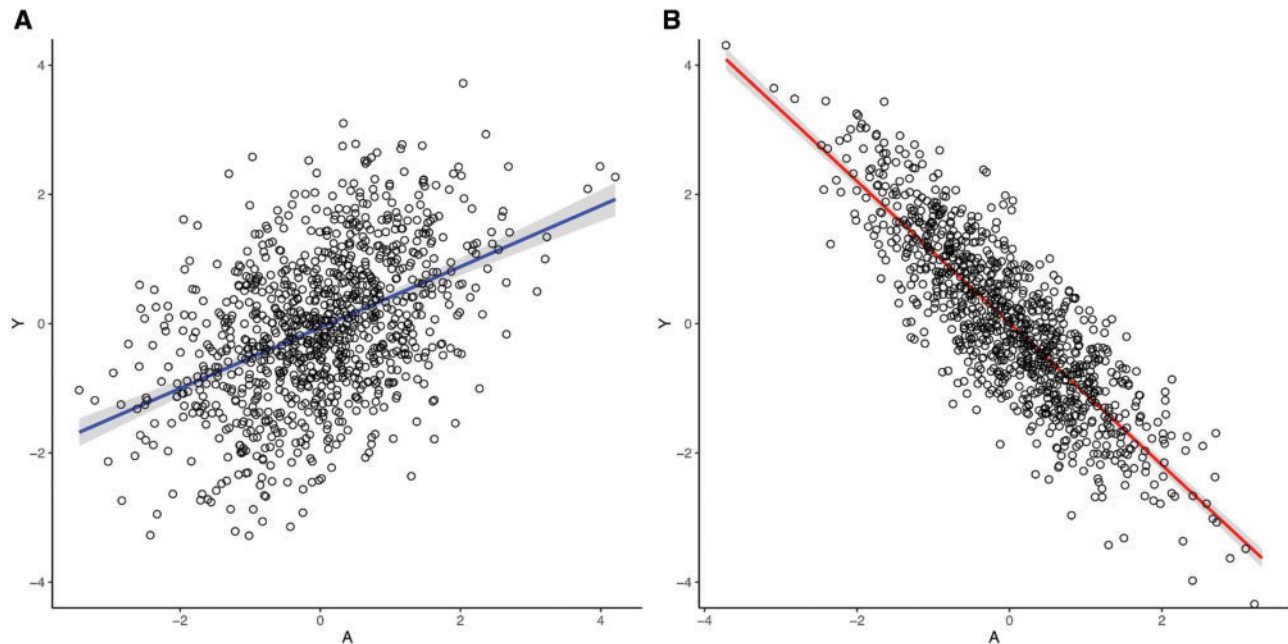
Collider structure

Unlike in [Figure 1A](#), where the causal arrows start from W , in [Figure 1B](#) they now point towards C from A and Y . If we condition on C (e.g. using regression or stratification), we will create collider bias. The common effect C is referred to as a collider on the path $A \rightarrow C \leftarrow Y$ because two arrow heads collide on this node. For intuition, suppose that rain (A) and a sprinkler (Y) are the only two causes of a wet ground (C). We also assume that the sprinkler is on a daily timer, and not related to the weather. Then, if the ground is wet, knowing that it has not rained implies that the sprinkler must be on. If we ignore the colliding structure, we may conclude that rain has a negative effect on the sprinkler even when we know a priori that this is not the case.⁸

Table 1. Coefficients and standard errors of the linear association between Y (outcome) and A (exposure) illustrating confounding and collider effects, $n = 1000$

| | Dependent variable (Y) | | | | |
|-----------|---|--|---|--|--------------------------|
| | W (confounder) | | C (collider) | | |
| | Unadjusted coefficient (standard error) (Fit 1) | Adjusted coefficients (standard error) (Fit 2) | Unadjusted coefficient (standard error) (Fit 3) | Adjusted coefficients (standard error) (Fit 4) | |
| | | | | | |
| A | 0.471 (0.030) | 0.289 (0.032) | A | 0.326 (0.031) | -0.416 (0.035) |
| W | | 0.425 (0.035) | C | | 0.491 (0.018) |
| Intercept | -0.061 (0.033) | -0.060 (0.031) | | 0.010 (0.031) | 0.035 (0.023) |
| AIC | 100.420 | -31.992 | | -55.369 | -626.824 |

Note: lower AIC is better.

**Figure 2.** Visualization of the collider effect. A: model fit2 (Box 1). B: model fit4 (Box 2).

Conditioning on the collider induces an association between the potential outcomes $Y(a)$ and the exposure (A), and conditional ignorability $Y(a) \perp A | W, C$ no longer holds. In other words: in Figure 1B and C, conditioning on the collider C opens the back-door path between A and Y which was previously blocked by the collider itself ($A \rightarrow C \leftarrow Y$). Thus, the association between A and Y would be a mixture of the association due to the effect of A on Y and the association due to the open back-door path. Thus, association would not be causation any more.

Figure 1C gives another, more complex collider structure usually known as M-bias, in which the collider (C) is the

effect of a common cause (W1) of the exposure (A) and a common cause (W2) of the outcome (Y). There is only one back-door path, and it is already blocked by the collider (C); thus we do not need to control for anything. This is the difference between confounders and colliders: a path will be open if one does not adjust for confounders, but blocked if adjustment is made; for colliders, it is the other way around. However, some could consider C to be a classical confounder as it is associated with both A, via ($A \leftarrow W1 \rightarrow C$), and with Y, via a path that does not go through A ($C \leftarrow W2 \rightarrow Y$), and it is not in the causal pathway between A and Y. However, controlling for C will introduce a collider bias.

Box 2. To generate data consistent with Figure 2B

```

library(visreg) # load package to visualize regression output
library(ggplot2) # load package to visualize regression output
N <- 1000      # sample size
set.seed(777)
A <- rnorm(N)  # exposure
Y <- -1.2 * A + rnorm(N) # outcome
C <- 1 * A + 1 * Y + rnorm(N) # collider
fit3 <- lm(Y ~ A) # crude model
fit4 <- lm(Y ~ A + C) # adjusted model
# visualize adjusted model
g2 <- visreg(fit4, 'A', gg = TRUE, line = list(col = 'red'),
  points = list(size = 2, pch = 1, col = 'black')) + theme_classic() +
  coord_cartesian(ylim = c(-4, 4)) +
  ggtitle("Figure 2B")

```

Note that if you use the traditional characteristics used to identify confounders ([i.e. a third variable (W) associated with both the exposure (A) and the outcome (Y) that is not in the causal pathway between A and Y], you can confuse a collider with a confounder.

To simulate the scenario portrayed in Figure 2B, we generate data again using a simple linear data generating mechanism (Box 2). First, we simulate A as a standard normally distributed variable. Y equals the value of A plus an error term, and C is generated depending on both A and Y, plus error. Note that as shown in Figure 1B, now the exposure A and the outcome Y are the parents of C (their common effect). We fit the unadjusted model excluding the collider (fit3) and then the model including the collider (fit4: collider model). The true causal coefficient of the exposure A is -1.2 , and the coefficients for the association of the collider C with the exposure A and the outcome Y are 1.0 and 1.0, respectively (Box 2).

Table 1 (columns 3, 4) shows the coefficient of A in the unadjusted model (fit3) and the coefficients of A and C in the model adjusting for the collider (fit4). Unlike in the previous section, the simpler regression without C approximately recovers the true coefficient of A (0.3) with an estimate of 0.326, whereas the regression adjusting for C is substantially biased (-0.416). The model which includes the collider (fit4) is not unequivocally inferior from a predictive point of view, where the main focus is to improve the model's predictive performance. For instance, the model containing the collider has a much lower Akaike Information Criterion (AIC) than

the one without the collider (Table 1). However, conditioning on the collider C has paradoxically changed the direction of the association between A and Y (Figure 2B, Table 1: column 4). Thus in this case, conditioning on the collider in the regression model introduces a bias whereas ignoring the collider does not add bias. The paradoxical negative association occurs when both A and Y are positively correlated with the collider.

From this demonstration, it is clear that subject-matter knowledge (i.e. plausible biological mechanisms in clinical epidemiological settings) is necessary to perform causal estimation.¹⁴ Thus, using DAGs to communicate causal structural relationships between variables helps in identifying variables that act as colliders, and identify where conditioning may create non-causal associations between the exposure (A) and outcome (Y).^{14–16}

Motivating example

Data generation

Based on a motivating example in non-communicable disease epidemiology, we generated a dataset with 1000 observations to contextualize the effect of conditioning on a collider. Nearly one in three Americans suffer from hypertension and more than half do not have it under control.¹⁷ Increased levels of systolic blood pressure over time are associated with increased cardiovascular morbidity and mortality.¹⁸

Summative evidence shows that exceeding the recommendations for 24-h dietary sodium intake in grams (g) is associated with increased levels of systolic blood pressure (SBP) in mmHg.¹⁹ Furthermore, with advancing age, the kidney undergoes several anatomical and physiological changes that limit the adaptive mechanism responsible for maintaining the composition and volume of the extracellular fluid. These include a decline in glomerular filtration rate and the impaired ability to maintain water and sodium homeostasis in response to dietary and environmental changes.²⁰ Likewise, age is associated with structural changes in the arteries and thus SBP.¹⁸

Age is a common cause of both high SBP and impaired sodium homeostasis. Thus, age acts as a confounder for the association between sodium intake (SOD) and SBP (i.e. age is on the back-door path between sodium intake and SBP) as depicted in Figure 3. However, high levels of 24-h excretion of urinary protein (proteinuria) are caused by sustained high SBP and increased 24-h dietary sodium intake. Therefore, as depicted in Figure 3, proteinuria acts as a collider (via the path $SOD \rightarrow PRO \leftarrow SBP$). In a realistic scenario, one might control for proteinuria if physiological factors influencing SBP are not completely understood by the researcher, the relationships between variables are not depicted in a DAG or proteinuria is conceptualized as a confounder. Controlling for proteinuria (PRO) introduces collider bias.

We are interested in estimating the effect of 24-h dietary sodium intake (in grams) on SBP, adjusting for age. The objective of the illustration is to show the paradoxical effect of 24-h dietary sodium intake on SBP after conditioning on a collider (proteinuria). Box 3 shows the data

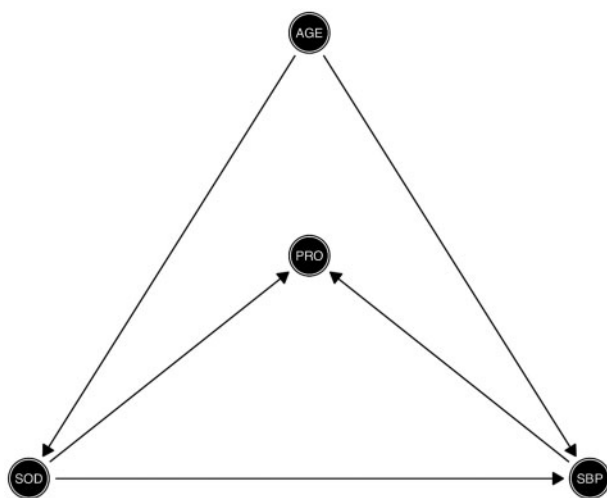


Figure 3. Directed acyclic graph depicting the structural causal relationship of the exposure and outcome, confounding and collider effects. Exposure: 24-h sodium dietary intake in g (SOD); outcome: systolic blood pressure in mmHg (SBP); confounder: age in years (AGE); collider: 24-h urinary protein excretion, proteinuria (PRO).

generation for the simulated data based on the structural relationship between the variables depicted in the DAG from Figure 3. We assumed that SBP is a common cause of age and dietary sodium intake. We also simulated 24-h excretion of urinary protein as a function of age, SBP and sodium intake. We aimed to have a range of values of the simulated data which was biologically plausible and as close to reality as possible.^{21,22}

Supplementary Table 1 (available as Supplementary data at *IJE* online) shows the descriptive statistics (minimum, maximum, mean, median, first and third quartiles) of the generated data. Note that for educational purposes, we present the code and results for a single dataset simulated by our data-generating mechanism. However at the end of the illustration, we also present the results of 1000 Monte-Carlo simulations with a sample size of 10 000 patients, aiming to quantify the bias associated with conditioning on a collider.

The simulation assumes linear relationships between the variables. Thus, the interpretation of the beta coefficients in the formulae of the code in Box 3 is straightforward. The true causal effect of sodium intake on SBP is 1.05 (i.e. systolic blood pressure = $\beta_1 \times \text{sodium} + \beta_2 \times \text{age} + \varepsilon$; where $\beta_1 = 1.05$, $\beta_2 = 2.00$ and ε is a standard normally distributed error). The coefficients for the association of PRO with SBP and sodium intake are 2.0 and 2.8, respectively (i.e. Proteinuria = $\beta_1 \times \text{SBP} + \beta_2 \times \text{Sodium} + \varepsilon$; where $\beta_1 = 2.0$, $\beta_2 = 2.8$ and ε is a standard normally distributed error) (Box 3). Supplementary Figure 1 (available as Supplementary data at *IJE* online) shows the functional form for each variable and the multivariable Spearman's correlation matrix.

We fit three different linear regression models (Box 4) to evaluate the effect of sodium intake on SBP: (i) unadjusted model; (ii) model adjusted for age; (iii) model adjusted for age and the collider (proteinuria). The model specifications are shown here below; in Box 4 we show how to fit and visualize the corresponding models in R.

Models specification

Model 0: Systolic Blood Pressure in mmHg = $\beta_0 + \beta_1 \times \text{Sodium in g} + \varepsilon$

Model 1: Systolic Blood Pressure in mmHg = $\beta_0 + \beta_1 \times \text{Sodium in g} + \beta_2 \times \text{Age in years} + \varepsilon$

Model 2: Systolic Blood Pressure in mmHg = $\beta_0 + \beta_1 \times \text{Sodium in g} + \beta_2 \times \text{Age in years} + \beta_3 \times \text{Proteinuria in mg} + \varepsilon$

We also fit three logistic regression models to evaluate the effect of sodium intake on hypertension defined as a binary outcome ($SBP \geq 140 \text{ mmHg} = 1$, $SBP < 140 \text{ mmHg} = 0$): (i) an unadjusted model; (ii) a model adjusted for age; and (iii) a model adjusted for age and the collider (proteinuria).

Box 3. Data generation consistent with Figure 3

```

generateData <- function(n, seed) {
  set.seed(seed)
  Age_years <- rnorm(n, 65, 5)
  Sodium_gr <- Age_years / 18 + rnorm(n)
  sbp_in_mmHg <- 1.05 * Sodium_gr + 2.00 * Age_years + rnorm(n)
  hypertension <- ifelse(sbp_in_mmHg > 140, 1, 0)
  Proteinuria_in_mg <- 2.00 * sbp_in_mmHg + 2.80 * Sodium_gr + rnorm(n)
  data.frame(sbp_in_mmHg, hypertension, Sodium_gr, Age_years, Proteinuria_in_mg)
}
ObsData <- generateData(n=1000, seed = 777)

```

Box 4. Linear regression models in R

```

library(broom) # load packages to visualize regression model's output
library(visreg)
## Models Fit
fit0 <- lm(sbp_in_mmHg ~ Sodium_gr, data = ObsData); tidy(fit0)
fit1 <- lm(sbp_in_mmHg ~ Sodium_gr + Age_years, data = ObsData); tidy(fit1)
fit2 <- lm(sbp_in_mmHg ~ Sodium_gr + Age_years + Proteinuria_in_mg, data = ObsData); tidy(fit2)

## Models visualization
par(mfrow = c(1, 3))
visreg(fit0, ylab = 'SBP in mmHg', line = list(col = 'blue'),
  points = list(cex = 1.5, pch = 1), jitter = 10, bty = 'n')
visreg(fit1, ylab = 'SBP in mmHg', line = list(col = 'blue'),
  points = list(cex = 1.5, pch = 1), jitter = 10, bty = 'n')
visreg(fit2, ylab = 'SBP in mmHg', line = list(col = 'red'),
  points = list(cex = 1.5, pch = 1), jitter = 10, bty = 'n')

```

The model specifications are the same as described above, but now with a binary outcome (hypertension); in [Box 5](#) we show how to fit and visualize the corresponding models in R using a forest plot function.

Effect of conditioning on a collider

[Table 2](#) shows the model coefficients and goodness of fit from the linear regression models, and [Figure 5](#) shows odds ratios from the logistic regression models. [Figure 4](#)

shows the regression line and 95% confidence interval for the predicted level of SBP, illustrating the effect of conditioning on a collider. The adjusted regression line was derived as the predicted estimate of SBP, conditional on the median value of age for [Figure 4B](#) and age and proteinuria for [Figure 4C](#).²³ As opposed to the unadjusted and bivariate models ([Figure 4A](#) and [B](#)), the collider model ([Figure 4C](#)) suggests a negative relationship between sodium intake and SBP (i.e. for one unit increase in sodium intake, the expected SBP decreases by 0.9 mmHg).

Box 5. Multiplicative scale visualization using a forest plot function

```
## Models fit on multiplicative scale

library(dplyr)
library(forestplot)

fit3 <- glm(hypertension ~ Sodium_gr, family=binomial(link='logit'), data=ObsData)
or <- round(exp(fit3$coef) [2], 3) # conditional odds ratio from logistic model
ci95 <- exp(confint(fit3)) [-1,] # 95% CI of odds ratio

fit4 <- glm(hypertension ~ Sodium_gr + Age_years, family = binomial(link = 'logit'), data = ObsData)
or <- round(exp(fit4$coef) [2], 3)
ci95 <- exp(confint(fit4)) [2,]

fit5 <- glm(hypertension ~ Sodium_gr + Age_years + Proteinuria_in_mg, family = binomial(link =
'logit'), data = ObsData)
or <- round(exp(fit5$coef) [2], 3)
ci95 <- exp(confint(fit5)) [2,]

## Forest plot (see supplementary material for accessing the complete code)
fp <- rbind(result1, result2, result3); fp %>% or_graph()
```

Table 2. Univariate, bivariate and multivariate coefficients and standard errors for the linear association between systolic blood pressure and 24-h sodium dietary intake, adjusted for age acting as a confounder and proteinuria acting as a collider, $n = 1000$

| | | Dependent variable: systolic blood pressure in mmHg | | |
|-----------------------------------|-------------|---|--|--|
| | | Univariate coefficient (standard error) | Bivariate coefficients (standard error) | Multivariate collider coefficients (standard error) |
| True effect of sodium in g | 1.05 | | | |
| Sodium in g | | 3.960 (0.298) | 1.039 (0.032) | -0.902 (0.036) |
| Age in years | | | 2.004 (0.007) | 0.416 (0.027) |
| Proteinuria in mg | | | | 0.396 (0.007) |
| Intercept | | 119.420 (1.122) | -0.311 (0.407) | -0.091 (0.192) |
| AIC | | 7363.45 | 2807.89 | 1302.66 |

Note: lower AIC is better.

The odds ratio for the effect of sodium on hypertension similarly suggests that it is protective (i.e. for one unit increase in sodium intake, the risk of hypertension decreases by 98%) (Figure 5).

Monte-Carlo simulation results

Box 6 shows the code used to run the Monte-Carlo simulation on the additive scale, using the same setting as in Box 3.

The true simulated causal effect of 24-h sodium intake on SBP was 1.05 mmHg in the linear model, and the coefficients for the association of PRO with SBP and sodium intake were 2.0 and 2.8, respectively. After 1000 simulation runs, the estimated additive effect of 24-h sodium intake on SBP was -0.91 mmHg (i.e. for one unit increase in sodium intake, there was a decrease of -0.91 units in SBP). The relative bias due to conditioning on proteinuria (the collider) was 13.3%.

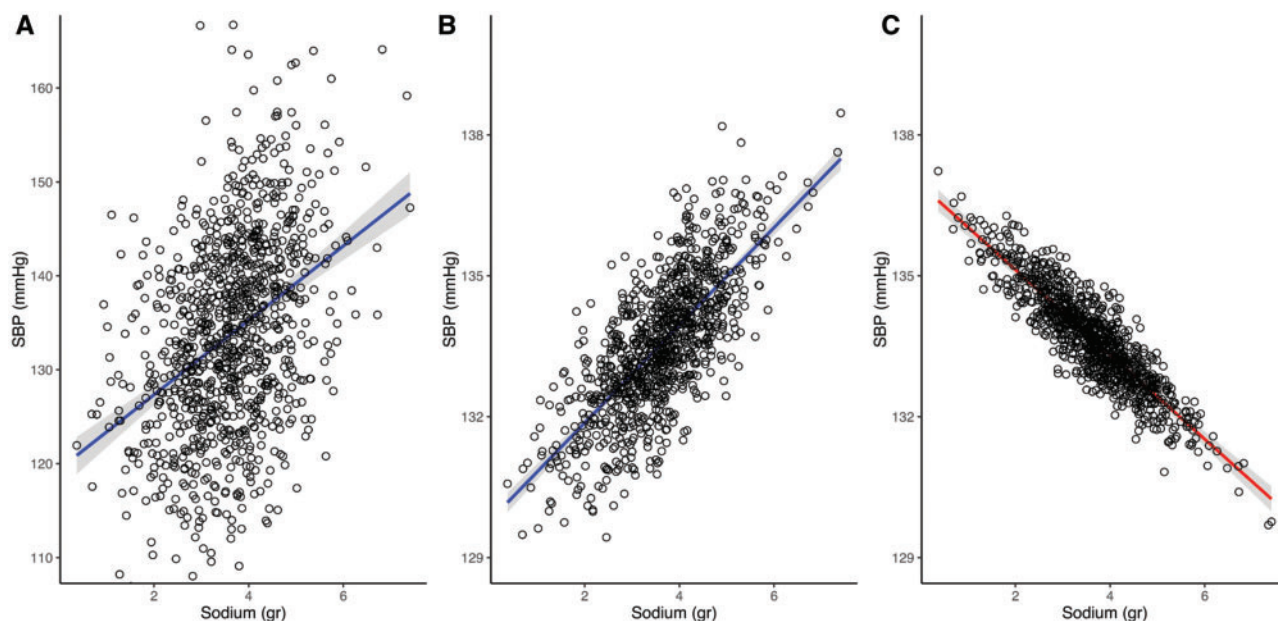


Figure 4. Collider effect for the illustration: univariate (A), bivariate (B) and multivariate (C) models fit for the linear association between systolic blood pressure and 24-h sodium dietary intake, adjusted for age acting as a confounder and proteinuria acting as a collider, $n = 1000$.

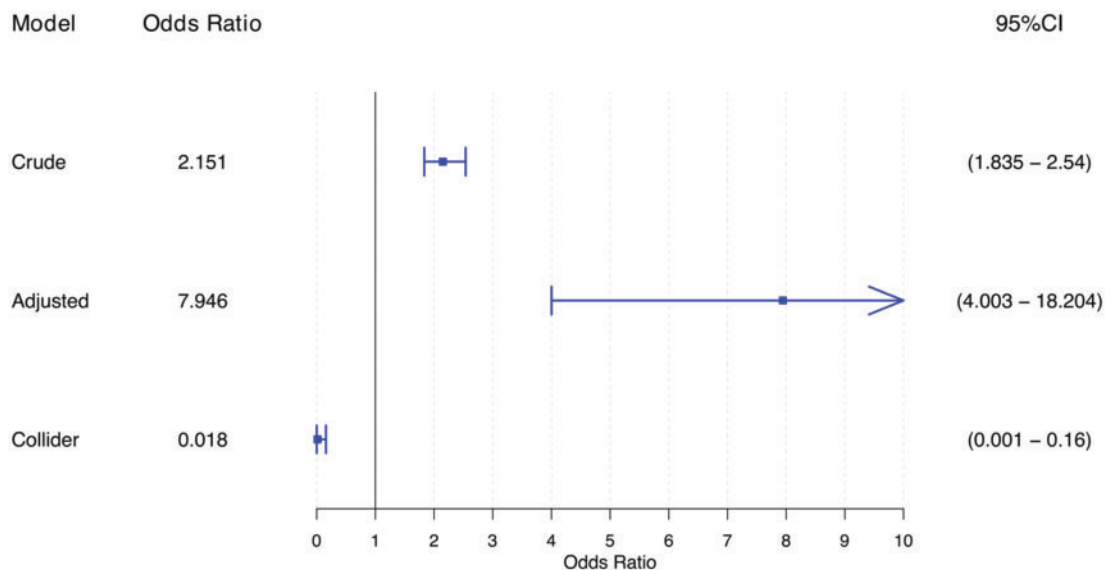


Figure 5. Collider effect for the illustration in a multiplicative scale for the effect of 24-h sodium dietary intake on systolic blood pressure, adjusted for age acting as a confounder and proteinuria acting as a collider, $n = 1000$. Crude model: unadjusted model. Adjusted model: adjusted for age acting as a confounder. Collider model: adjusted model including age and proteinuria acting as a collider.

The code included in all of the boxes is provided in a supplementary file (available as [Supplementary data](#) at *IJE* online). We also provide the link to a web application [<http://watzilei.com/shiny/collider/>] ([Supplementary Figure 2](#), available as [Supplementary data](#) at *IJE* online) where users can dynamically modify the values of the true causal effect and the coefficients in the data generation process of the collider model. The collider web application allows users to

interactively modify the range of values of the slider input and visualize the collider effect of the example. As shown in the web application, the strength of the association of the collider with both the exposure and the outcome determines the strength of the paradoxical protective effect of 24-h dietary sodium intake in grams on systolic blood pressure.

The magnitude of the causal effect between the exposure and the outcome, and the collider with the exposure

Box 6. Monte Carlo simulations

```
# Monte Carlo Simulations

R<-1000

true <- rep(NA, R)

collider <- rep(NA, R)

se <- rep(NA, R)

set.seed(050472)

for(r in 1: R) {

  if (r%%10 == 0) cat(paste('This is simulation run number', r, '\n'))

# Function to generate data

  generateData <- function(n) {

    Age_years <- rnorm(n, 65, 5)

    Sodium_gr <- Age_years / 18 + rnorm(n)

    sbp_in_mmHg <- 1.05 * Sodium_gr + 2.00 * Age_years + rnorm(n)

    Proteinuria_in_mg <- 2.00 * sbp_in_mmHg + 2.80 * Sodium_gr + rnorm(n)

    data.frame(sbp_in_mmHg, Sodium_gr, Age_years, Proteinuria_in_mg)

  }

ObsData <- generateData(n=10000)

# True effect

true[r] <- summary(lm(sbp_in_mmHg ~ Sodium_gr + Age_years, data = ObsData))$coef[2,1]

# Collider effect

collider[r] <- summary(lm(sbp_in_mmHg ~ Sodium_gr + Age_years + Proteinuria_in_mg, data =
ObsData))$coef[2,1]

  se[r] <- summary(lm(sbp_in_mmHg ~ Sodium_gr + Age_years + Proteinuria_in_mg, data =
ObsData))$coef[2,2]

}

# Estimate of sodium true effect

mean(true)

# Estimate of sodium biased effect in the model including the collider

mean(collider)

# simulated standard error/confidence interval of outcome regression

lci <- (mean(collider) - 1.96*mean(se)); mean(lci)

uci <- (mean(collider) + 1.96*mean(se)); mean(uci)

# Bias

Bias <- (true - abs(collider)); mean(Bias)

# % Bias

relBias <- ((true - abs(collider)) / true); mean(relBias) * 100

# Plot bias

plot(relBias)
```

Table 3. Different scenarios for the true causal effect and the magnitude of the association between the collider with the exposure (α_1) and the outcome (α_2), $n = 1000$

| Causal model True causal effect (β_1) | Collider model Magnitude of the association between the collider with the exposure (α_1) and the outcome (α_2), assuming $\alpha_1 = \alpha_2$ | Estimated causal effect | Absolute bias |
|--|--|-------------------------|---------------|
| 1 | 0.5 | 0.630 | 0.370 |
| | 1.0 | 0.033 | 0.967 |
| | 1.5 | -0.368 | 1.368 |
| | 2.0 | -0.596 | 1.596 |
| | 2.5 | -0.727 | 1.727 |
| | 3.0 | -0.807 | 1.807 |
| | 3.5 | -0.858 | 1.858 |
| | 4.0 | -0.892 | 1.892 |
| | 4.5 | -0.916 | 1.916 |
| | 5.0 | -0.933 | 1.933 |
| 2 | 0.5 | 1.453 | 0.547 |
| | 1.0 | 0.558 | 1.442 |
| | 1.5 | -0.045 | 2.045 |
| | 2.0 | -0.388 | 2.388 |
| | 2.5 | -0.586 | 2.586 |
| | 3.0 | -0.706 | 2.706 |
| | 3.5 | -0.783 | 2.783 |
| | 4.0 | -0.835 | 2.835 |
| | 4.5 | -0.871 | 2.871 |
| | 5.0 | -0.897 | 2.897 |
| 3 | 0.5 | 2.277 | 0.723 |
| | 1.0 | 1.082 | 1.918 |
| | 1.5 | 0.278 | 2.722 |
| | 2.0 | -0.181 | 3.181 |
| | 2.5 | -0.445 | 3.445 |
| | 3.0 | -0.606 | 3.606 |
| | 3.5 | -0.709 | 3.709 |
| | 4.0 | -0.778 | 3.778 |
| | 4.5 | -0.826 | 3.826 |
| | 5.0 | -0.861 | 3.861 |
| 4 | 0.5 | 3.100 | 0.900 |
| | 1.0 | 1.607 | 2.393 |
| | 1.5 | 0.600 | 3.400 |
| | 2.0 | 0.027 | 3.973 |
| | 2.5 | -0.304 | 4.304 |
| | 3.0 | -0.505 | 4.505 |
| | 3.5 | -0.634 | 4.634 |
| | 4.0 | -0.721 | 4.721 |
| | 4.5 | -0.781 | 4.781 |
| | 5.0 | -0.825 | 4.825 |
| 5 | 0.5 | 3.923 | 1.077 |
| | 1.0 | 2.132 | 2.868 |
| | 1.5 | 0.923 | 4.077 |
| | 2.0 | 0.234 | 4.766 |
| | 2.5 | -0.163 | 5.163 |
| | 3.0 | -0.405 | 5.405 |
| | 3.5 | -0.560 | 5.560 |
| | 4.0 | -0.664 | 5.664 |
| | 4.5 | -0.737 | 5.737 |
| | 5.0 | -0.789 | 5.789 |

Causal model: $SBP = \beta_0 + \beta_1 SOD + \beta_2 AGE + \beta_3 PRO$. Collider model: $PRO = \alpha_0 + \alpha_1 SOD + \alpha_2 SBP$. Absolute bias = true - estimate. AGE = age in years. SOD = 24-h dietary sodium intake (g). PRO = 24-h excretion of urinary protein (proteinuria) (mg). SBP = systolic blood pressure (mmHg).

and the outcome, determines whether paradoxical effects arise when conditioning on the collider. Table 3 shows different values for the true causal effect of sodium intake on SBP and the estimated causal effect for different values of the association between PRO (i.e. the collider) with sodium intake (α_1) and SBP (α_2) in the collider model, and assuming $\alpha_1 = \alpha_2$ (i.e. the same magnitude for the collider-exposure and the collider-outcome associations in the collider model). Overall, with this data-generating structure, the collider bias reduces the magnitude of the estimated causal effect between sodium intake and SBP. To create a paradoxical effect (i.e. the negative association between sodium intake and SBP), we found that increasing the true causal effect requires an increase of the strength of the association between collider-exposure and collider-outcome association with respect to the magnitude of the true causal effect (Table 3). Note that assuming $\alpha_1 = \alpha_2$ is not realistic, but it is a convenient simplification that helps to gain intuition about changes in the magnitude of bias.

There are two additional situations where collider bias arises, which are important to point out: (i) collider bias arising not from the choice of variables to control in the analysis, but from conditioning on a measured or unmeasured common effect of the exposure and the outcome in a sample selection; and (ii) situations where the collider is both a collider and a confounder.²⁴

Recent evidence shows that even modest influences on sample selection can generate biased and potentially misleading estimates of both phenotypic and genotypic associations.²⁵ However, the solution is not often clear, as information regarding sample selection and attrition might be unmeasured. On the other hand, in M-bias settings where the collider is also a confounder, it is useful to understand the trade-offs in bias between collider and confounder control. The size in the magnitude of collider bias may often be comparable with bias from classical confounding.²⁴ It has been shown that M-bias has a small impact unless associations between the collider and confounders are very large (relative risk >8). Generally in this situation, controlling for confounding would be prioritized over avoiding M-bias.²⁶

Conclusion

We investigated a situation where adding a certain type of variable to a linear regression model, called a ‘collider’, led to bias with respect to the regression coefficient estimates while still improving the model fit. DAGs are based on subject-matter knowledge and are vital for identifying colliders. Determining if a variable is a collider involves critical thinking about the true unobserved data-generation process and the relationship between the variables for a

given scenario.^{16,27} Then, the decision whether to include or exclude the variable in a regression model using observational data in epidemiology is based on whether the purpose of the study is prediction or explanation/causation. Under the structures we investigated here, adding a collider to a regression model is not advised when one is interested in the estimation of causal effects, as this may open a back-door path. However, if prediction is the purpose of the model, the inclusion of colliders in the models may be advisable if it reduces the model’s prediction error. Most research in epidemiology tries to explain how the world works (i.e. it is causal); thus to prevent paradoxical associations, epidemiologists estimating causal effects should be aware of such variables.

Supplementary Data

Supplementary data are available at *IJE* online.

Funding

M.A.L.F. is supported by the Spanish National Institute of Health, Carlos III Miguel Servet I Investigator Award (CP17/00206). M.P. is supported by the Andalusian Department of Health Research, Development and Innovation Office project grant PI-0152/2017. A.V. was supported by the National Institutes of Health (grants DK107407 and DK115392) and by the Doris Duke Charitable Foundation (award 2015085). M.S. is supported by a New Investigator Salary Award from the Canadian Institutes of Health Research.

Author Contributions

The article and Shiny application arise from the motivation to disseminate the principles of modern epidemiology among clinicians and applied researchers. M.A.L.F. developed the concept, designed the study, carried out the simulation, analysed the data, and wrote the article. D.R.S. and M.A.L.F. developed the Shiny application. All authors interpreted the data, and drafted and revised the manuscript, code for the manuscript and code for the Shiny application. All authors read and approved the final version of the manuscript. M.A.L.F. is the guarantor of the article.

Conflict of interest: The authors declare that they do not have any conflict of interest associated with this research and the content is solely the responsibility of the authors.

References

1. Greenland S, Morgenstern H. Confounding in health research. *Annu Rev Public Health* 2001;22:189–212.
2. Cole SR, Platt RW, Schisterman EF *et al*. Illustrating bias due to conditioning on a collider. *Int J Epidemiol* 2010;39:417–20.

3. Vanderweele TJ, Vansteelandt S. Conceptual issues concerning mediation, interventions and composition. *Stat Interface* 2009;2:457–68.
4. Hernán MA, Hernández-Díaz S, Robins JM. A structural approach to selection bias. *Epidemiology* 2004;15:615–25.
5. Robins JM, Hernán MÁ, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000;11:550–60.
6. Rohrer JM. Thinking clearly about correlations and causation: graphical causal models for observational data. *Adv Methods Pract Psychol Sci* 2018;1:27–42.
7. Pearl J. Causal diagrams for empirical research. *Biometrika* 1995;82:669–88.
8. Pearl J. *Causality: Models, Reasoning, and Inference*. 2nd edn. New York, NY: Cambridge University Press, 2009.
9. Luque-Fernandez MA, Zoega H, Valdimarsdottir U, Williams MA. Deconstructing the smoking-preeclampsia paradox through a counterfactual framework. *Eur J Epidemiol* 2016;31:613–23.
10. Hernandez-Diaz S, Schisterman EF, Hernan MA. The birth weight “paradox” uncovered? *Am J Epidemiol* 2006;164:1115–20.
11. Banack HR, Kaufman JS. The “obesity paradox” explained. *Epidemiology* 2013;24:461–62.
12. Whitcomb BW, Schisterman EF, Perkins NJ, Platt RW. Quantification of collider-stratification bias and the birthweight paradox. *Paediatr Perinat Epidemiol* 2009;23:394–402.
13. Rubin DB. Causal inference using potential outcomes. *J Am Stat Assoc* 2005;100:322–31.
14. Hernan MA. Causal knowledge as a prerequisite for confounding evaluation: an application to birth defects epidemiology. *Am J Epidemiol* 2002;155:176–84.
15. Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology* 1999;10:37–48.
16. Pearce N, Richiardi L. Commentary: three worlds collide: Berkson’s bias, selection bias and collider bias. *Int J Epidemiol* 2014;43:521–24.
17. Benjamin EJ, Blaha MJ, Chiuve SE *et al*. Heart disease and stroke statistics - 2017 update: a report from the American Heart Association. *Circulation* 2017;135:e146–603.
18. Gu Q, Burt VL, Paulose-Ram R, Yoon S, Gillum RF. High blood pressure and cardiovascular disease mortality risk among US adults: the third National Health and Nutrition Examination Survey mortality follow-up study. *Ann Epidemiol* 2008;18:302–09.
19. Sacks FM, Svetkey LP, Vollmer WM *et al*. Effects on blood pressure of reduced dietary sodium and the Dietary Approaches to Stop Hypertension (DASH) diet. *N Engl J Med* 2001;344:3–10.
20. Tareen N, Martins D, Nagami G, Levine B, Norris KC. Sodium disorders in the elderly. *J Natl Med Assoc* 2005;97:217–24.
21. Van Horn L, Carson JAS, Appel LJ *et al*. Recommended dietary pattern to achieve adherence to the American Heart Association/American College of Cardiology (AHA/ACC) guidelines: a scientific statement from the American Heart Association. *Circulation* 2016;134:e505–29.
22. Carroll MF. Proteinuria in adults: a diagnostic approach. *Am Fam Physician* 2000;62:1333–40.
23. Breheny P, Burchett W. Visualization of regression models using visreg. *R Journal* 2017;9:56–71.
24. Greenland S. Quantifying biases in causal models: classical confounding vs collider-stratification bias. *Epidemiology* 2003;14:300–06.
25. Munafo MR, Tilling K, Taylor AE, Evans DM, Davey Smith G. Collider scope: when selection bias can substantially influence observed associations. *Int J Epidemiol* 2018;47:226–35.
26. Wei L, Brookhart MA, Schneeweiss S, Mi X, Setoguchi S. Implications of M bias in epidemiologic studies: a simulation study. *Am J Epidemiol* 2012;176:938–48.
27. Pearce N, Lawlor DA. Causal inference - so much more than statistics. *Int J Epidemiol* 2016;45:1895–903.