

Commentary

The Table 2 Fallacy: Presenting and Interpreting Confounder and Modifier Coefficients

Daniel Westreich* and Sander Greenland

* Correspondence to Dr. Daniel Westreich, Department of Obstetrics and Gynecology, Duke Global Health Institute, Duke University, DUMC 3967, Durham, NC 27710 (e-mail: daniel.westreich@duke.edu).

Initially submitted January 13, 2012; accepted for publication October 11, 2012.

It is common to present multiple adjusted effect estimates from a single model in a single table. For example, a table might show odds ratios for one or more exposures and also for several confounders from a single logistic regression. This can lead to mistaken interpretations of these estimates. We use causal diagrams to display the sources of the problems. Presentation of exposure and confounder effect estimates from a single model may lead to several interpretative difficulties, inviting confusion of direct-effect estimates with total-effect estimates for covariates in the model. These effect estimates may also be confounded even though the effect estimate for the main exposure is not confounded. Interpretation of these effect estimates is further complicated by heterogeneity (variation, modification) of the exposure effect measure across covariate levels. We offer suggestions to limit potential misunderstandings when multiple effect estimates are presented, including precise distinction between total and direct effect measures from a single model, and use of multiple models tailored to yield total-effect estimates for covariates.

causal diagrams; causal inference; confounding; direct effects; epidemiologic methods; mediation analysis; regression modeling

Abbreviation: HIV, human immunodeficiency virus.

In scientific manuscripts in which the results of population-based research are reported, “Table 1” typically includes a description of key demographic, social, and clinical characteristics of the study groups, often categorized by the levels or arms of the primary exposure (treatment). This information is often useful when thinking about both internal and external validity (generalizability) of the study results (1).

Also common is a “Table 2” that shows multivariate-adjusted associations with the outcome for variables summarized in Table 1. For example, in a study of the impact of aspirin on stroke risk in which the main study association is adjusted for age and sex, Table 2 might contain risk, rate, or odds ratios for aspirin, age, and sex. Thus, the table contains effect estimates for secondary risk factors for the outcome in addition to the effect estimate for the primary exposure. Usually all of the estimates are derived from the same model.

In the present commentary, we explain how such a table can be misleading. In particular, we illustrate how readers can be misled by effect estimates for secondary risk factors from the model also used to control for confounding of exposure. We also offer suggestions for improving tables with effect estimates for different variables. A related discussion has been published by VanderWeele and Staudt (2).

“Primary effect” will refer to the causal effect of an exposure of primary interest, “secondary effect” will refer to the causal effect of a covariate not of primary interest in the initial adjustment model (e.g., a confounder or effect-measure modifier), “total effect” will refer to the net of all associations of a variable through all causal pathways to the outcome, and “direct effect” will refer to an association after blocking or controlling some of those pathways (3–6). We presume the reader understands causal diagrams, as many introductions are now available (5–8). For simplicity, we use logistic models (so coefficients of single regressors

represent changes in the log odds of the outcome) and we assume the outcome is uncommon at all covariate levels to avoid complications arising from noncollapsibility of odds ratios (9, 10).

TABLE 2 FALLACIES

An example

A recent article (11) illustrates the issues we raise above. Page 1,359 shows a Table 2 with the main exposure (work-related violence) and covariates including gender, age, cohabitation status, education, and income. The text refers to “cause-specific hazard ratios for use of psychotropics in relation to work-related violence and the covariates gender, age, cohabitation, education, income...” (11, p. 1358) without indicating how interpretations may differ between variables listed in Table 2. Thus, the presentation may leave the impression that, for example, the hazard ratio for gender can be interpreted in the same way as the hazard ratio for the main effect of work-related violence.

Same model, different types of effect

How can Table 2 be harmful? By presenting adjusted effect estimates for secondary risk factors alongside the adjusted effect estimate for the primary exposure, Table 2 suggests implicitly that all of these estimates can be interpreted similarly, if not identically. This is often not the case.

Consider an observational study of the effect of human immunodeficiency virus (HIV) seroconversion on subsequent age-specific 10-year risk of stroke. Figure 1 gives a possible causal diagram for the study that implies that age and smoking status are determined before HIV status. In this figure, the open paths from HIV to stroke passing through smoking and age show that the effect of HIV (the primary exposure) on stroke risk may be confounded by these covariates. This would occur if, for example, the probability of infection with HIV increases with age and smoking, in both cases perhaps due to immunosuppression. We might account for this confounding using the following logistic model:

$$\begin{aligned} \text{logit}(\text{Stroke}|\text{HIV}, \text{Smoking}, \text{Age}) \\ = \beta_0 + \beta_1 \times \text{HIV} + \beta_2 \times \text{Smoking} + \beta_3 \times \text{Age}. \end{aligned} \quad (\text{Model 1})$$

We might report the estimated coefficients for HIV, smoking, and age (or their antilogs, which are odds ratios) in model 1 in Table 2. Many readers would assume that these 3 coefficients (or odds ratios) could all be interpreted similarly, simply, and causally; after all, they are mutually adjusted. However, even if the model is correct, these 3 coefficients represent different types of causal effects.

Given Figure 1 and model 1, β_1 can be interpreted as the conditional total effect of contracting HIV on the 10-year log odds (logit) of stroke, that is, the log odds ratio for the total effect of HIV on stroke at any given level of smoking and age. However, β_2 cannot be interpreted in the same




Figure 1. Causal diagram for the effect of human immunodeficiency virus (HIV) seroconversion on 10-year stroke risk, with confounding by smoking level and age.

way: In particular, it cannot be interpreted as a total effect of smoking. In the adjusted model, β_2 is a direct effect (a direct log odds ratio) of smoking relative to HIV; that is, β_2 is the portion of the smoking effect on the log odds of stroke that is not mediated through the smoking effect on HIV seroconversion. More precisely, it is the controlled direct effect of smoking on the log odds when HIV is held fixed at a given level, thus blocking the smoking effect on HIV. To the extent we might allow talk of effects of aging, β_3 is similarly the controlled direct effect of aging on the log odds when HIV and smoking are held fixed, thus blocking the age effects on smoking and HIV (3, 12).

To interpret β_2 as a direct effect of smoking after blocking its effect on HIV infection, we must adjust for all confounders of both the exposure-outcome relationship and the mediator-outcome relationship (3, 4, 12); given Figure 1, these assumptions are met. Given those assumptions (and the usual assumptions of no other bias source), all 3 coefficients represent certain causal effects; nonetheless, interpreting all 3 as the same type of effect is a subtle error. In particular, β_1 is a total effect; β_2 is a direct effect after blocking smoking's effect on HIV; and β_3 is a direct effect after blocking the effects of age on smoking and HIV. Yet, many readers would interpret all 3 as total effects.

Same model, different degrees of confounding

Now suppose Figure 2, which adds an unmeasured covariate U that affects only smoking and stroke, is correct. Model 1 then remains valid for obtaining an unbiased estimate of the total effect of HIV on the log odds of stroke; that is, under model 1, β_1 retains its interpretation as the total effect of HIV, as it did under Figure 1, because smoking and age satisfy the back-door criterion for sufficiency for confounding control (4, 5, 7, 13). This means that, after conditioning on (or blocking) smoking and age, there are no open paths between HIV and stroke besides the direct arrow from HIV to stroke, which represents the causal effect of interest in this analysis.

Under Figure 2, adjustment for smoking removes confounding of HIV by U and confounding of HIV by smoking because U is connected to HIV only through its




Figure 2. Causal diagram for the effect of human immunodeficiency virus (HIV) seroconversion on 10-year stroke risk, with confounding by smoking level, age, and U .

effect on smoking. Nonetheless, the interpretation of β_2 as a direct effect of smoking after blocking HIV is now incorrect, because β_2 is confounded by U . More precisely, under the causal model in Figure 2 and the regression model 1, β_2 no longer validly represents an effect of smoking (although β_1 remains a valid estimate of the total HIV effect on the log odds of stroke).

Perhaps less obviously, under Figure 2 and model 1, β_3 no longer validly represents an effect of age. This is because smoking is now a collider on the indirect path from age to smoking to U to stroke, and adjustment for smoking (needed to estimate the HIV effect) opens the U -stroke path, thus biasing β_3 as an effect measure (4, 5). In other words, we are forced to control for smoking to unbiasedly estimate the HIV effect; but under Figure 2, adjustment for smoking biases the estimated direct effect of age (if there is no adjustment for U). Estimation of the direct effect of age on stroke would require control of HIV, smoking, and U .

We have thus illustrated how a model sufficient for estimating the average effect of the primary study exposure can be insufficient to provide an unbiased estimate of secondary effects. Under Figure 2, to obtain an unbiased estimate of the direct effect of smoking after blocking its effect on HIV and the direct effect of age after blocking its effect on smoking and HIV, we would have to adjust for U , for example, using the following model:

$$\begin{aligned} \text{logit}(\text{Stroke}|\text{HIV}, \text{Smoking}, \text{Age}, U) \\ = \beta_0 + \beta_1 \times \text{HIV} + \beta_2 \times \text{Smoking} + \\ \beta_3 \times \text{Age} + \beta_4 \times U. \end{aligned} \quad (\text{Model 2})$$

Under Figure 2, model 2, and low risk of the outcome, it follows from basic collapsibility results (9, 10) that β_1 would be approximately equal in models 1 and 2, because HIV is unassociated with U given smoking and age (because of the noncollapsibility of rate ratios and odds ratios, this equality would fail using a log-linear rate or logistic risk models with common outcomes (9, 10)). Nonetheless, the smoking and age coefficients (β_2 and β_3) may

differ considerably between the models because smoking is associated with U given HIV and age, and age is associated with U given HIV and smoking. This difference translates into a bias in estimates of β_2 and β_3 under model 1 when considered as estimates of smoking and age effects.

INTERPRETATION GETS HARDER WITH HETEROGENEITY

Variation (heterogeneity) of effect measures across covariate levels can severely complicate separation of direct and indirect effects (3, 14). In what follows, we will need to distinguish between 2 sources of variation of the effect measure for the study exposure across levels of an adjustment (model) covariate, depending on whether one views the covariate as a secondary intervention variable or merely as a passive stratification factor (13, 15).

The first source is variation in the effect measure that is attributable in a precise causal sense to effects of a modeled covariate. When neither exposure nor the modeled covariate is confounded by uncontrolled covariates, the observed associations of the outcome with exposure and the modeled covariate are attributable entirely to the joint effects of the exposure and the covariate; in particular, interventions on the modeled covariate (e.g., smoking cessation) would in this case alter the measure of exposure effect. This source of effect-measure variation has been called “causal interaction” (13, 15), although some find that term objectionable because the variation might reflect nothing more than model choice rather than interaction defined in terms of biologic mechanisms.

The second source of variation is that attributable to effects of uncontrolled covariates, which need not be confounders of the exposure effect but may still be confounders of covariate effects. In this situation, intervention on a modeled covariate might not change the exposure effect measure at all, or at least not to the extent suggested by mere descriptive comparison of the exposure effect measure across the covariate’s levels. This sort of effect variation (variation without a specified source) has been termed “effect heterogeneity” or “effect measure modification” (5, 13, 15), although the term “effect measure modification” (or worse, “effect modification”) is problematic because it evokes the more narrow concept of causal interaction (in which changing the covariate would change the effect measure).

Interpretation of product terms with no uncontrolled confounding of any variable in the model

Consider first the simpler case in which the only possible causes of effect-measure variation are the covariates in the model, as in Figure 1. Causal diagrams are nonparametric and thus silent about the degree of variation because that is a parametric property represented by coefficients of product terms (“statistical interactions”). There is no reason to expect homogeneity (as assumed in models 1 and 2) in most applications; the best we can hope is that our use of a homogeneous model (one without product terms, such as model 1 or model 2) provides roughly unbiased estimates

of average exposure effects across covariates and so is not misleading for marginal (standardized) effects (16).

However, sometimes the heterogeneity is severe enough that it needs to be modeled, as would be the case for studying HIV if the HIV odds ratio were the targeted effect measure and varied considerably with age and smoking; that is, if β_4 or β_5 were important to retain in the model:

$$\begin{aligned} \text{logit(Stroke|HIV, Smoking, Age)} \\ = \beta_0 + \beta_1 \times \text{HIV} + \beta_2 \times \text{Smoking} + \beta_3 \times \text{Age} \\ + \beta_4 \times \text{HIV} \times \text{Smoking} + \beta_5 \times \text{HIV} \times \text{Age} \\ + \beta_6 \times \text{Smoking} \times \text{Age.} \end{aligned} \quad (\text{Model 3})$$

Under Figure 1 and model 3, the log odds ratio (the change in the log odds of stroke) associated with moving from HIV = 0 to HIV = 1 is $\beta_1 + \beta_4 \times \text{Smoking} + \beta_5 \times \text{Age}$, which is a function of smoking and age. Typical presentations would tabulate estimates of this quantity (or its antilog) for various choices of smoking and age. We will consider the causal interpretation of this quantity below.

As a preliminary, for the single-variable (“main effect”) terms to be interpretable when product terms are used, zero must be a meaningful value in the study for each covariate. With product terms present, single-variable exposure terms represent effects only when all the covariates that appear in product terms with exposure are zero. Thus, by adding a product term with age, we have to center age around a reference value common in the sample, so that “Age = 0” corresponds to this value, not actual age. Here, we will center age by subtracting 40 years (so that age represents a deviation from 40 years of age). We will also assume that non-smokers (Smoking = 0) are included in the sample. It is important that the units used represent meaningful differences; otherwise, product coefficients may be misleadingly tiny even in the presence of substantial effect variation. Hence, we will assume that age is measured in decades and smoking in packs/day (rather than the misleadingly small units of years and cigarettes/day).

Under Figure 1 and model 3, the smoking-and-age-specific total effect of HIV on the log odds of stroke is $\beta_1 + \beta_4 \times \text{Smoking} + \beta_5 \times \text{Age}$, which varies with smoking and age. Such unconfounded variation in an exposure effect measure represents the deviation from additivity of the joint causal effects of HIV, smoking, and aging on the log odds of stroke. Model 3 implies that β_1 is the HIV-stroke log odds ratio when the covariates are zero, whereas the HIV-smoking and HIV-aging product coefficients β_4 and β_5 represent variation in this log odds ratio across smoking and age. Figure 1 further implies that β_1 is the total effect of HIV on the log odds of stroke among 40-year-old non-smokers (Age = 0); β_2 is the direct effect of smoking a pack of cigarettes per day on the log odds among persons 40 years of age when HIV is prevented (HIV is set to 0); and β_3 is the direct effect of aging a decade on the log odds among people in whom both HIV and smoking are prevented (HIV and Smoking both set to 0).

Figure 1 also implies that β_4 is the change in the HIV effect on the log odds produced by smoking one additional

pack per day; specifically, smoking an additional pack per day adds β_4 to the total HIV log odds ratio at any age. What may be less anticipated, however, is that β_4 is also the change in the direct smoking effect on the log odds produced by HIV; specifically, HIV adds β_4 to the direct effect of smoking a pack per day on the log odds, at any given age. Thus β_4 represents modification of the *total* HIV and the *direct* smoking log odds ratios on stroke, both conditional on age.

Next, β_5 is the change in the total HIV log odds ratio from aging an additional decade, when smoking is set (held) to a fixed level; specifically, each decade of aging adds β_5 to the HIV-stroke log odds ratio when smoking is left unchanged. We thus might say that β_5 is the direct modification of the HIV-stroke log odds ratio produced by a decade of aging when smoking level is held constant. However, β_5 is also the change in the log odds ratio for the direct effect of a decade of aging produced by HIV, when smoking level is held constant. Similarly, β_6 is the change in the log odds ratio for the direct effect of smoking a pack per day produced by an additional decade of aging when HIV status is unchanged and is also the change in the log odds ratio for the direct effect of aging from an additional pack per day when HIV status is set to a fixed level.

Model 3 forces all of the pairwise modifying effects just described (e.g., modification of HIV-stroke log odds ratio by age) to remain the same across the variable held constant (e.g., smoking), so that smoking-and-age-specific log odds ratios for the HIV effect equal $\beta_1 + \beta_4 \times \text{Smoking} + \beta_5 \times \text{Age}$. Modeling heterogeneity of modification (nonadditivity beyond 2-way products) would require at least a triple product in the model (such as HIV \times smoking \times age); the interpretation of the coefficient of this triple product would vary depending on which effect was targeted.

Interpretation changes when there are uncontrolled confounders of a covariate

Interpretation of effect-measure variation across modeled covariates changes if some of those covariates are themselves confounded by uncontrolled covariates, even if the primary study exposure is not confounded (15). In this case, some or all of the variation in the exposure effect measure may be due to those uncontrolled covariates. For example, U confounds smoking in Figure 2 and confounds the effects of age given smoking; thus, U could be responsible for some of the variation in an HIV effect measure across combined smoking and age levels. In the extreme case in which the HIV effect measure was constant across smoking and age given U , U would be the only remaining cause of variation in the HIV effect measure across smoking and age levels; thus, there would be no modification of the HIV effect measure by smoking or age, so that neither alteration of smoking habits nor aging would change the HIV effect measure. The product terms in model 3 are commonly called “interaction terms” in statistics, but in this example these terms would not represent any causal interaction among the variables in the products.

In less extreme examples conforming to Figure 2, some variation of the HIV effect measure across smoking and

age would remain upon control of U , so that alteration of smoking habits would change the HIV effect measure. Nonetheless, if U were not controlled, the observed variation of the HIV effect measure across smoking and age would not equal the change in the HIV odds ratio that would be produced by a change in smoking or by aging. In other words, U could confound the apparent modification of the HIV effect measure by smoking and age even if it did not account for it entirely. We note that in Figure 2, the confounding of age modification of the HIV effect measure by U would be entirely due to conditioning on smoking (which is a collider for U and age (17)); again, however, conditioning on smoking was necessary to remove confounding by smoking and by U of the HIV effect measure.

To illustrate how this affects Table 2 interpretation, suppose Figure 2 is the correct causal model and that model 3 remains the correct regression model for the effect of HIV on stroke. The model terms involving smoking and age would suffice to control for confounding of the HIV-stroke relationship by smoking, age, and U because all backdoor paths from HIV to stroke that include U are blocked by conditioning on smoking and age. Thus, the total effect of HIV on the cohort would still be given by smoking-and-age standardization, and smoking-and-age-specific effects of HIV on the log odds of stroke would still equal $\beta_1 + \beta_4 \times \text{Smoking} + \beta_5 \times \text{Age}$.

Nonetheless, we would expect confounding of smoking by U to bias most interpretations given above in terms of the effects of smoking or aging on stroke or HIV effect measures. Specifically, we could no longer interpret β_2 and β_3 as direct effects of smoking and aging on the log odds of stroke, β_4 and β_5 as alterations of the HIV log odds ratio produced by smoking or aging, or β_6 as alteration of smoking log odds ratios by aging or vice-versa. Thus, under Figure 2, β_4 , β_5 , and β_6 no longer represent deviation from additivity of the joint causal effects on the log odds (causal interactions). For example, because of confounding by U , we could no longer interpret β_4 as the change in the HIV-stroke log odds ratio produced by each additional pack per day of smoking. In particular, β_1 and $\beta_1 + \beta_4$ would remain the total HIV effects on the log odds among the nonsmokers and pack-per-day smokers, yet β_4 would be biased by U effects as a measure of modification of the HIV-stroke log odds ratio by smoking. Put another way, β_4 no longer equals the departure from additivity of the joint causal effects of HIV and smoking on the stroke log odds (their causal interaction on the logit scale), even though β_1 and $\beta_1 + \beta_4$ still equal the effects of HIV on the log odds of stroke among nonsmokers and pack-per-day smokers given age.

Phrased more generally, under Figure 2, U could confound the apparent modification of the HIV effect measure by smoking and age, even though U would not confound the smoking- and age-specific total effects of HIV. The variation in an HIV effect measure across smoking and age would be real but could be produced entirely or partly by the unobserved differences in U across levels of smoking or age, rather than by smoking and age alone. Because the graph is nonparametric (and in particular not dependent the chosen effect measure), this conclusion would apply

regardless of the measure or model form used to quantify effects.

Table 2 may discourage realistic modeling

The desire for simple estimates to present in Table 2 may discourage realistically flexible variable specifications. For example, spline coding of variables can improve model fit over linear or categorical codings and produce more credible smooth models of complex dose-response functions (5, 18–20). However, splines require one to recode a variable into several functions, the coefficients of which lack simple interpretation. Similarly, the use of “black-box” machine learning techniques (21, 22) or many product terms among multiple variables can improve model fit and validity but does not produce easily interpretable coefficients. Outputs of flexible regressions can be presented in terms of interpretable effect measures (e.g., risk differences or ratios) comparing specific values of a variable (e.g., 1 pack per day smokers vs. nonsmokers), but the extra labor and commentary required may lead some investigators to opt for models that are less flexible and less valid than what can be easily fit with modern software.

AVOIDING TABLE 2 FALLACIES

A reasonable starting point for causal modeling is construction of plausible causal diagrams that display the analyst’s best understanding of the literature. Those diagrams encode the causal assumptions used to select covariates for inclusion in the model (3–7, 23, 24). Table 2 problems can be avoided by limiting the table to estimates of the primary exposure effect measures under the different models, with the secondary “adjustment” covariates reported in a footnote along with how they were categorized or modeled, as is common practice in space-limited presentations. This practice leaves room for multiple estimates for the same exposures using different models (e.g., showing results from adjustment under different causal assumptions or diagrams or for difference as well as ratio measures of effect). If some of the primary exposures are intermediate with respect to other primary exposures, these intermediate primaries will have to be left out of the model used to estimate the total effects of the other primary exposures; thus, the estimates in this table may be derived from models with different covariate subsets.

Using different covariate subsets would allow Table 2 to include estimates of total effects of secondary covariates (which might be useful to other researchers). In our example, Table 2 could provide estimates of odds ratios for total effects of each variable in Figure 1, using a model with all 3 variables to estimate the total HIV effect on stroke (adjusted for smoking and age) but using a model without HIV to estimate the total smoking effect on stroke and a model with age alone to estimate the total age effect on stroke. This approach invites objections, however; among them is that the direct effect of smoking is what is needed for sensitivity analysis for other HIV-stroke analyses that did not have smoking measurements (because

confounding by smoking is transmitted only through its effects outside of any effect on HIV).

If indeed the direct effects are of interest, then Table 2 could include effect estimates from the model with all 3 variables. In that case, it seems advisable that the text description note that the smoking and age estimates are for direct effects. Under Figure 2, these interpretations require that the model include variables like U (if, unlike U , they are measured) that confound the smoking and age direct effects even if they do not confound the HIV effects (5, 6, 9).

DISCUSSION

The problem with a table presenting multiple estimated effect measures from the same model ("Table 2") is that it encourages the reader to interpret all these estimates in the same way, typically as total-effect estimates. As illustrated above, the interpretation of a confounder effect estimate may be different than for the exposure effect estimate. Of course, it is possible that some secondary reported associations are unbiased for total effects and that others are unbiased for direct effects; nonetheless, the assumption that all estimates reported in Table 2 are for total effects is not warranted. Thus, we recommend that a presentation of secondary effect estimates would best specify the type of effect being estimated.

As in all causal modeling, the interpretations described above should raise questions about the ethics and feasibility of the interventions implicit in the effect definitions. In the example, definition of the direct effect of age requires holding a person's smoking level and HIV status constant as they age, which would be unethical for smokers (among whom reduction should be encouraged) and infeasible even if desirable for those HIV negative (because some will maintain unsafe practices). More generally, definitions of direct and indirect effects involve combined interventions on both the exposure and mediators; some combinations may resemble nothing anyone would consider in reality (25–27), thus violating positivity constraints (28).

In sum, presenting estimates of effect measures for secondary risk factors (confounders and modifiers of the exposure effect measure) obtained from the same model as that used to estimate the primary exposure effects can lead readers astray in a number of ways. Extra thought and description will be needed when interpreting such secondary estimates.

ACKNOWLEDGMENTS

Author affiliations: Department of Obstetrics and Gynecology, Duke University, Durham, North Carolina (Daniel Westreich); Duke Global Health Institute, Duke University, Durham, North Carolina (Daniel Westreich); Department of Epidemiology, University of California Los Angeles, Los Angeles, California (Sander Greenland); and Department of Statistics, University of California Los Angeles, Los Angeles, California (Sander Greenland).

This work was supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development at the National Institutes of Health (grant R00 HD063961 to D.W.) and the National Institute of Allergy and Infectious Disease at the National Institutes of Health (grant 2P30 AI064518 to D.W.).

We wish to thank Stephen R. Cole and Tyler VanderWeele for helpful comments in the preparation of this manuscript and the reviewers and editor for their extensive helpful suggestions.

Conflict of interest: none declared.

REFERENCES

1. Cole SR, Stuart EA. Generalizing evidence from randomized clinical trials to target populations: the ACTG 320 trial. *Am J Epidemiol.* 2010;172(1):107–115.
2. VanderWeele TJ, Staudt NC. Causal diagrams for empirical legal research: methodology for identifying causation, avoiding bias, and interpreting results. *Law Probab Risk.* 2011;10(4):329–354.
3. Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. *Epidemiology.* 1992;3(2):143–155.
4. Cole SR, Hernan MA. Fallibility in estimating direct effects. *Int J Epidemiol.* 2002;31(1):163–165.
5. Glymour MM, Greenland S. Causal diagrams. In: Rothman KJ, Greenland S, Lash TL, eds. *Modern Epidemiology*. 3rd ed. Philadelphia, PA: Lippincott Williams & Wilkins; 2008:183–209.
6. Pearl J. *Causality*. 2nd ed. New York, NY: Cambridge University Press; 2009.
7. Pearl J. Causal diagrams for empirical research. *Biometrika.* 1995;82(4):669–688.
8. Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology.* 1999;10(1):37–48.
9. Greenland S, Robins JM, Pearl J. Confounding and collapsibility in causal inference. *Stat Sci.* 1999;14(1):29–46.
10. Greenland S, Pearl J. Adjustments and their consequences—collapsibility analysis using graphical models. *Int Stat Rev.* 2011;79(3):401–426.
11. Madsen IE, Burr H, Diderichsen F, et al. Work-related violence and incident use of psychotropics. *Am J Epidemiol.* 2011;174(12):1354–1362.
12. VanderWeele TJ. Marginal structural models for the estimation of direct and indirect effects. *Epidemiology.* 2009;20(1):18–26.
13. VanderWeele TJ, Knol MJ. Interpretation of subgroup analyses in randomized trials: heterogeneity versus secondary interventions. *Ann Intern Med.* 2011;154(10):680–683.
14. Vanderweele TJ, Vansteelandt S. Odds ratios for mediation analysis for a dichotomous outcome. *Am J Epidemiol.* 2010;172(12):1339–1348.
15. VanderWeele TJ. On the distinction between interaction and effect modification. *Epidemiology.* 2009;20(6):863–871.
16. Greenland S, Maldonado G. The interpretation of multiplicative-model parameters as standardized parameters. *Stat Med.* 1994;13(10):989–999.

17. Greenland S. Quantifying biases in causal models: classical confounding vs collider-stratification bias. *Epidemiology*. 2003;14(3):300–306.
18. Greenland S. Dose-response and trend analysis in epidemiology: alternatives to categorical analysis. *Epidemiology*. 1995;6(4):356–365.
19. Orsini N, Greenland S. A procedure to tabulate and plot results after flexible modeling of a quantitative covariate. *Stata J*. 2011;11(1):1–29.
20. Howe CJ, Cole SR, Westreich DJ, et al. Splines for trend analysis and continuous confounder control. *Epidemiology*. 2011;22(6):874–875.
21. Breiman L. Statistical modeling: the two cultures. *Stat Sci*. 2001;16(3):199–231.
22. Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. New York: Springer; 2009.
23. Hernán MA, Hernández-Díaz S, Werler MM, et al. Causal knowledge as a prerequisite for confounding evaluation: an application to birth defects epidemiology. *Am J Epidemiol*. 2002;155(2):176–184.
24. Robins JM. Data, design, and background knowledge in etiologic inference. *Epidemiology*. 2000;11(3):313–320.
25. Robins JM, Richardson TS. Alternative graphical causal models and the identification of direct effects. In: Shroout P, ed. *Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures*. New York, NY: Oxford University Press; 2010.
26. Kaufman JS. Commentary: gilding the black box. *Int J Epidemiol*. 2009;38(3):845–847.
27. Kaufman JS. Invited commentary: decomposing with a lot of supposing. *Am J Epidemiol*. 2010;172(12):1349–1351.
28. Westreich D, Cole SR. Invited commentary: positivity in practice. *Am J Epidemiol*. 2010;171(6):674–677.