Published in final edited form as:

Annu Rev Sociol. 2014 July ; 40: 31–53. doi:10.1146/annurev-soc-071913-043455.

Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable

Felix Elwert¹ and Christopher Winship²

¹Department of Sociology, University of Wisconsin, Madison, Wisconsin 53706

²Department of Sociology, Harvard University, Cambridge, Massachusetts 02138

Abstract

Endogenous selection bias is a central problem for causal inference. Recognizing the problem, however, can be difficult in practice. This article introduces a purely graphical way of characterizing endogenous selection bias and of understanding its consequences (Hernán et al. 2004). We use causal graphs (direct acyclic graphs, or DAGs) to highlight that endogenous selection bias stems from conditioning (e.g., controlling, stratifying, or selecting) on a so-called collider variable, i.e., a variable that is itself caused by two other variables, one that is (or is associated with) the treatment and another that is (or is associated with) the outcome. Endogenous selection bias can result from direct conditioning on the outcome variable, a post-outcome variable, a post-treatment variable, and even a pre-treatment variable. We highlight the difference between endogenous selection bias, common-cause confounding, and overcontrol bias and discuss numerous examples from social stratification, cultural sociology, social network analysis, political sociology, social demography, and the sociology of education.

Keywords

causality; directed acyclic graphs; identification; confounding; selection

1. INTRODUCTION

A large literature in economics, sociology, epidemiology, and statistics deals with sample selection bias. Sample selection bias occurs when the selection of observations into a sample is not independent of the outcome (Winship & Mare 1992, p. 328). For example, estimates of the effect of education on wages will be biased if the sample is restricted to low earners (Hausman & Wise 1977), and estimates for the effect of motherhood on wages will be biased if the sample is restricted to employed women (Gronau 1974, Heckman 1974).

Surveys of sample selection bias in sociology tend to emphasize examples where selection is affected by the treatment or the outcome, i.e., where sample selection occurs after treatment (Berk 1983, Winship & Mare 1992, Stolzenberg & Relles 1997, Fu et al. 2004). The

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

underlying problem, which we call endogenous selection bias, is more general than this, however, in two respects. First, bias can also arise if sample selection occurs prior to treatment. Second, bias can occur from simple conditioning (e.g., controlling or stratifying) on certain variables even if no observations are discarded from the population (i.e., if no sample is selected).

Recognizing endogenous selection bias can be difficult, especially when sample selection occurs prior to treatment or when no selection (in the sense of discarding observations) takes place. This article introduces sociologists to a new, purely graphical, way of characterizing endogenous selection bias that subsumes all of these scenarios. The goal is to facilitate recognition of the problem in its many forms in practice. In contrast to previous surveys, which primarily draw on algebraic formulations from econometrics, our presentation draws on graphical approaches from computer science (Pearl 1995, 2009, 2010) and epidemiology (Greenland et al. 1999a, and especially Hernán et al. 2004). This paper makes five points:

First, we present a definition of endogenous selection bias. Endogenous selection bias results from conditioning on an endogenous variable that is caused by two other variables, one that is (or is associated with) the treatment and one that is (or is associated with) the outcome (Hernán et al. 2002, 2004). In the graphical approach introduced below, such endogenous variables are called collider variables. Collider variables may occur before or after the treatment and before or after the outcome.¹

Second, we describe the difference between confounding, or omitted variable bias, and endogenous selection bias (Pearl 1995; Greenland et al. 1999a; Hernán et al. 2002, 2004). Confounding originates from common causes, whereas endogenous selection originates from common outcomes. Confounding bias results from failing to control for a common cause variable that one should have controlled for, whereas endogenous selection bias results from controlling for a common outcome variable that one should not have controlled for.

Third, we highlight the importance of endogenous selection bias for causal inference by noting that all nonparametric identification problems can be classified as one of three underlying problems: overcontrol bias, confounding bias, and endogenous selection bias. Endogenous selection captures the common structure of a large number of biases usually known by different names, including selective nonresponse, ascertainment bias, attrition bias, informative censoring, Heckman selection bias, sample selection bias, homophily bias, and others.

Fourth, we illustrate that the commonplace advice not to select (i.e., condition) on the dependent variable of an analysis is too narrow. Endogenous selection bias can originate in many places. Endogenous selection can result from conditioning on an outcome, but it can also result from conditioning on variables that have been

¹We avoid the simpler term “selection bias” because it has multiple meanings across literatures. “Endogenous selection bias” as defined in Section 4 of this article encompasses “sample selection bias” from econometrics (Vella 1998) and “Berkson’s (1946) bias” and “M-bias” (Greenland 2003) from epidemiology. Our definition is identical to those of “collider stratification bias” (Greenland 2003), “selection bias” (Hernán et al. 2004), “explaining away effect” (Kim & Pearl 1983), and “conditioning bias” (Morgan & Winship 2007).

affected by the outcome, variables that temporally precede the outcome, and even variables that temporally precede the treatment (Greenland 2003).

Fifth, and perhaps most importantly, we illustrate the ubiquity of endogenous selection bias by reviewing numerous examples from stratification research, cultural sociology, social network analysis, political sociology, social demography, and the sociology of education.

The specific methodological points of this paper have been made before in various places in econometrics, statistics, biostatistics, and computer science. Our innovation lies in offering a unified graphical presentation of these points for sociologists. Specifically, our presentation relies on direct acyclic graphs (DAGs) (Pearl 1995, 2009; Elwert 2013), a generalization of conventional linear path diagrams (Wright 1934, Blalock 1964, Duncan 1966, Bollen 1989) that is entirely nonparametric (making no distributional or functional form assumptions).² DAGs are formally rigorous yet accessible because they rely on graphical rules instead of algebra. We hope that this presentation may empower even sociologists with little mathematical training to recognize and understand precisely when and why endogenous selection bias may trouble an empirical investigation.

We begin by reviewing the required technical material. Section 2 describes the notion of nonparametric identification. Section 3 introduces DAGs. Section 4 explains the three structural sources of association between two variables—causation, confounding, and endogenous selection—and describes how DAGs can be used to distinguish causal associations from spurious associations. Section 5—the heart of this article—examines a large number of applied examples in which endogenous selection is a problem. Section 6 concludes by reflecting on dealing with endogenous selection bias in practice.

2. IDENTIFICATION VERSUS ESTIMATION

We investigate endogenous selection bias as a problem for the identification of causal effects. Following the modern literature on causal inference, we define causal effects as contrasts between potential outcomes (Rubin 1974, Morgan & Winship 2007). Recovering causal effects from data is difficult because causal effects can never be observed directly (Holland 1986). Rather, causal effects have to be inferred from observed associations, which are generally a mixture of the desired causal effect and various noncausal, or spurious, components. Neatly dividing associations into their causal and spurious components is the task of identification analysis. A causal effect is said to be *identified* if it is possible, with ideal data (infinite sample size and no measurement error), to purge an observed association of all noncausal components such that only the causal effect of interest remains.

Identification analysis requires a theory (i.e., a model or assumptions) of how the data were generated. Consider, for example, the identification of the causal effect of education on wages. If one assumes that ability positively influences both education and wages, then the observed marginal (i.e., unadjusted) association between education and wages does not

²The equivalence of DAGs and nonparametric structural equation models is discussed in Pearl (2009, 2010). The relationship between DAGs and conventional structural equation models is discussed by Bollen & Pearl (2013).

identify the causal effect of education on wages. But if one can assume that ability is the only factor influencing both education and wages, then the conditional association between education and wages after perfectly adjusting for ability would identify the causal effect of education on wages. Unfortunately, the analyst's theory of data generation can never be fully tested empirically (Robins & Wasserman 1999). Therefore, it is important that the analyst's theory of data generation is stated explicitly for other scholars to inspect and judge. We use DAGs to clearly state the data-generating model.

Throughout, we focus on nonparametric identification, i.e., identification results that can be established solely from qualitative causal assumptions about the data-generating process without relying on parametric assumptions about the distribution of variables (e.g., joint normality) or the functional form of causal effects (e.g., linearity). We consider this nonparametric focus an important strength because cause-effect statements are claims about social reality that sociologists deal with on a daily basis, whereas distributional and functional form assumptions often lack sociological justification.

Identification (assuming ideal data) is a necessary precondition for unbiased estimation (with real data). Because there is, at present, no universal solution to estimating causal effects in the presence of endogenous selection (Stolzenberg & Relles 1997, p. 495), we focus here on identification and deal with estimation only in passing. Reviews of the large econometric literature on estimating causal effects under various scenarios of endogenous selection can be found in Pagan & Ullah (1997), Vella (1998), Grasdøl (2001), and Christofides et al. (2003), which cover Hausman & Wise's (1977) truncated regression models and Heckman's (1976 and Heckman's (1979) classic two-step estimator, as well as numerous extensions. Bareinboim & Pearl (2012) discuss estimation under endogenous selection within the graphical framework adopted here. A growing literature in epidemiology and biostatistics builds on Robins (1986, 1994, 1999) to deal with the endogenous selection problems that occur with time-varying treatments.

3. DAGs IN BRIEF³

DAGs encode the analyst's qualitative causal assumptions about the data-generating process in the population (thus abstracting from sampling variability; see Pearl 2009 for technical details). Collectively, these assumptions comprise the model, or theory, on which the analyst can base her identification analysis. DAGs consist of three elements (Figure 1). First, letters represent random variables, which may be observed or unobserved. Variables can have any distribution (continuous or discrete). Second, arrows represent direct causal effects.⁴ Direct effects may have any functional form (linear or nonlinear), they may differ in magnitude across individuals (effect heterogeneity), and—within certain constraints (Elwert 2013)—they may vary with the value of other variables (effect interaction or modification). Because the future cannot cause the past, the arrows order pairs of variables in time such that DAGs do not contain directed cycles. Third, missing arrows represent the sharp assumption of no causal effect between variables. In the parlance of economics, missing arrows represent

³This and the following section draw on Elwert's (2013) survey of DAGs for social scientists.

⁴Strictly speaking, arrows represent possible (rather than certain) causal effects (see Elwert 2013). This distinction is not important for the purposes of this paper and we neglect it below.

exclusion restrictions. For example, Figure 1 excludes the arrows $U \rightarrow T$ and $X \rightarrow Y$, among others. Missing arrows are essential for inferring causal information from data.

For clarity, and without infringing upon our conceptual points, we assume that the DAGs in this review display all observed and unobserved common causes in the process under investigation. Note that this assumption implies that all other inputs into variables must be marginally independent; i.e., all variables have independent error terms. These independent error terms need not be displayed in the DAG because they never help nonparametric identification (we occasionally display them below when they harm identification).

A path is a sequence of arrows connecting two variables regardless of the direction of the arrowheads. A path may traverse a given variable at most once. A causal path between treatment and outcome is a path in which all arrows point away from the treatment and toward the outcome. Assessing existence and magnitude of causal paths is the task of causal testing and causal estimation, respectively. The set of all causal paths between a treatment and an outcome comprise the total causal effect. Unless specifically stated otherwise, this review is concerned with the identification of average total causal effects in the data represented by the DAG. In Figure 1, the total causal effect of T on Y comprises the causal paths $T \rightarrow Y$ and $T \rightarrow C \rightarrow Y$. All other paths between treatment and outcome are called noncausal (or spurious) paths, e.g., $T \rightarrow C \leftarrow U \rightarrow Y$ and $T \leftarrow X \leftarrow U \rightarrow Y$.

If two arrows along a path both point directly into the same variable, the variable is called a collider variable on the path (collider, for short). Otherwise, the variable is called a noncollider on the path. Note that the same variable may be a collider on one path and a noncollider on another path. In Figure 1, C is a collider on the path $T \rightarrow C \leftarrow U$ and a noncollider on the path $T \rightarrow C \rightarrow Y$.

We make extensive reference to conditioning on a variable. Most broadly, conditioning refers to introducing information about a variable into the analysis by some means. In sociology, conditioning usually takes the form of controlling for a variable in a regression model, but it could also take the form of stratifying on a variable, performing a group-specific analysis (thus conditioning on group membership), or collecting data selectively (e.g., excluding children, prisoners, retirees, or nonrespondents from a survey). We denote conditioning graphically by drawing a box around the conditioned variable.

4. SOURCES OF ASSOCIATION: CAUSATION, CONFOUNDING, AND ENDOGENOUS SELECTION

The Building Blocks: Causation, Confounding, and Endogenous Selection

The power of DAGs lies in their ability to reveal all marginal and conditional associations and independences implied by a qualitative causal model (Pearl 2009). This enables the analyst to discern which observable associations are solely causal and which ones are not—in other words, to conduct a formal nonparametric identification analysis. Absent sampling variation (i.e., chance), all observable associations originate from just three elementary configurations: chains ($A \rightarrow B$, $A \rightarrow C \rightarrow B$, etc.), forks ($A \leftarrow C \rightarrow B$), and inverted forks ($A \rightarrow C \leftarrow B$) (Pearl 1988, 2009; Verma & Pearl 1988). Conveniently, these three

configurations correspond exactly to the sources of causation, confounding bias, and endogenous selection bias.⁵

First, two variables are associated if one variable directly or indirectly causes the other. In Figure 2*a*, *A* indirectly causes *B* via *C*. In data generated according to this DAG, the observed marginal association between *A* and *B* would reflect pure causation—in other words, the marginal association between *A* and *B* identifies the causal effect of *A* on *B*. By contrast, conditioning on the intermediate variable *C* in Figure 2*b* blocks the causal path and hence controls away the causal effect of *A* on *B*. Thus, the conditional association between *A* and *B* given *C* (which, in this DAG, is zero) does not identify the causal effect of *A* on *B* (which is nonzero). We say that conditioning on a variable on a causal path from treatment to outcome induces overcontrol bias.

Second, two variables are associated if they share a common cause. In Figure 3*a*, *A* and *B* are associated only because they are both caused by *C*. This is the familiar situation of common-cause confounding bias—the marginal association between *A* and *B* does not identify the causal effect of *A* on *B*. Conditioning on *C* in Figure 3*b* eliminates this spurious association. Therefore, the conditional association between *A* and *B* given *C* identifies the causal effect of *A* on *B* (which, in this DAG, is zero).

The third reason two variables may be associated with each other is less widely known, but it is central to this article: Conditioning on the common outcome of two variables (i.e., a collider) induces a spurious association between them. We call this phenomenon endogenous selection bias. In Figure 4*a*, *A* and *B* are marginally independent because they do not cause each other and do not share a common cause. Thus, the marginal association between *A* and *B* identifies the causal effect (which, in this DAG, is zero). In Figure 4*b*, by contrast, conditioning on the collider *C* induces an association between its immediate causes *A* and *B*. Therefore, the conditional association between *A* and *B* given *C* is biased for the causal effect of *A* on *B*. A dotted line between the immediate causes of a conditioned collider denotes the spurious association induced between them by conditioning on the collider.⁶ This dotted line functions like a regular path segment when tracing paths between variables. Conditioning on a variable that is caused by a collider (a descendant of the collider) has qualitatively the same consequence as conditioning on the collider itself (Figure 4*c*) because the descendant is a proxy measure of the collider.

Because endogenous selection is not obvious and can be difficult to understand in the abstract, we first offer an intuitive hypothetical example to illustrate Figure 4. Consider the relationships between talent, *A*, beauty, *B*, and Hollywood success, *C*. Suppose, for argument's sake, that beauty and talent are unassociated in the general population (i.e., beauty does not cause talent, talent does not cause beauty, and beauty and talent do not share a common cause). Suppose further that beauty and talent are separately sufficient for becoming a successful Hollywood actor. Given these assumptions, success clearly is a

⁵Deriving the associational implications of causal structures requires two mild assumptions not discussed here, known as causal Markov assumption and faithfulness. See Glymour & Greenland (2008) for a nontechnical summary. See Pearl (2009) and Spirtes et al. (2000) for details.

⁶One does not draw a dotted line for the spurious association induced by common-cause confounding.

collider variable. Now condition on the collider, for example, by looking at the relationship between beauty and talent only among successful Hollywood actors. Under our model of success, knowing that a talentless person is a successful actor implies that the person must be beautiful. Conversely, knowing that a less than beautiful person is a successful actor implies that the person must be talented. Either way, conditioning on the collider (success) has created a spurious association between beauty and talent among the successful. This spurious association is endogenous selection bias.

This example may lack sociological nuance, but it illustrates the logic of endogenous selection bias that conditioning on a common outcome distorts the association between its causes.⁷ To achieve greater realism, one could loosen assumptions and, for example, allow that *A* causes *B* (maybe because talented actors invest less in self-beautification). The problem of endogenous selection, however, would not go away. The observed conditional association between talent and beauty given Hollywood success would remain biased for the true causal effect of talent on beauty. Adding arrows to a given set of variables never helps nonparametric identification (Pearl 2009). Unfortunately, sign and magnitude of endogenous selection bias cannot generally be predicted without additional assumptions (Hausman & Wise 1981, Berk 1983, VanderWeele & Robins 2009b).

In sum, overcontrol bias, confounding bias, and endogenous selection bias are distinct phenomena: They arise from different causal structures, and they call for different remedies. The contrasts among Figures 2, 3, and 4 neatly summarize this difference. Overcontrol bias results from conditioning on a variable on a causal path between treatment and outcome; the remedy is not to condition on variables on a causal path. Confounding bias arises from failure to condition on a common cause of treatment and outcome; the remedy is to condition on the common cause. Finally, endogenous selection bias results from conditioning on a collider (or the descendant of a collider) on a noncausal path linking treatment and outcome; the remedy is not to condition on such colliders or their descendants. All three biases originate from analytical mistakes, albeit from different ones: With confounding, the analyst did not condition although he should have; and with overcontrol and endogenous selection, he did condition although he should not have.

Identification in General DAGs

All DAGs are built from chains, forks, and inverted forks. Therefore, understanding the associational implications of chains (causation), forks (confounding), and inverted forks (endogenous selection) is sufficient for conducting nonparametric identification analyses in arbitrarily complicated causal models. The so-called d-separation rule summarizes these implications (Pearl 1988, 1995; Verma & Pearl 1988):

First, note that all associations in a DAG are transmitted along paths but that not all paths transmit association. A path between two variables, *A* and *B*, does not transmit association and is said to be blocked, closed, or d-separated if

⁷Conditioning on a collider always changes the association between its causes within at least one value of the collider, except for pathological situations, e.g., when the two effects going into the collider cancel each other out exactly. In real life, the probability of exact cancellation is zero, and we disregard this possibility henceforth. The technical literature refers to this as the faithfulness assumption (Spirtes et al. 2000).

1. the path contains a noncollider, C , that has been conditioned on, $A \rightarrow \boxed{C} \rightarrow B$, or $A \leftarrow \boxed{C} \rightarrow B$; or if
2. the path contains a collider, C , and neither the collider nor any of its descendants have been conditioned on, $A \rightarrow C \leftarrow B$.

Paths that are not d-separated do transmit association and are said to be open, unblocked, or d-connected. Two variables that are d-separated along all paths are statistically independent. Conversely, two variables that are d-connected along at least one path are statistically dependent, i.e., associated (Verma & Pearl 1988).

Importantly, conditioning on a variable has opposite consequences depending on whether the variable is a collider or a noncollider: Conditioning on a noncollider blocks the flow of association along a path, whereas conditioning on a collider opens the flow of association.

One can then determine the identifiability of a causal effect by checking whether one can block all noncausal paths between treatment and outcome without blocking any causal paths between treatment and outcome by conditioning on a suitable set of observed variables.⁸ For example, in Figure 1, the total causal effect of T on Y can be identified by conditioning on X because X does not sit on a causal path from T to Y and conditioning on X blocks the two noncausal paths between T and Y , $T \leftarrow \boxed{X} \leftarrow U \rightarrow Y$ and $T \leftarrow \boxed{X} \leftarrow U \rightarrow C \rightarrow Y$, that would otherwise be open. (A third noncausal path between T and Y , $T \rightarrow C \leftarrow U \rightarrow Y$, is unconditionally blocked because it contains the collider C . Conditioning on C would ruin identification for two reasons: First, it would induce endogenous selection bias by opening the noncausal path $T \rightarrow \boxed{C} \leftarrow U \rightarrow Y$; second, it would induce overcontrol bias because C sits on a causal path from T to Y , $T \rightarrow \boxed{C} \rightarrow Y$.)

The motivation for this review is that it is all too easy to open noncausal paths between treatment and outcome inadvertently by conditioning on a collider, thus inducing endogenous selection bias. Examples of such bias in sociology are discussed next.

5. EXAMPLES OF ENDOGENOUS SELECTION BIAS IN SOCIOLOGY

Endogenous selection bias may originate from conditioning on a collider at any point in a causal process. Endogenous selection bias can be induced by conditioning on a collider that occurs after the outcome; by conditioning on a collider that is an intermediate variable between the treatment and the outcome; and by conditioning on a collider that occurs before the treatment (Greenland 2003). Therefore, we organize the following sociological examples of endogenous selection bias according to the timing of the collider relative to the treatment and the outcome.

⁸This is the logic of “identification by adjustment,” which underlies identification in regression and matching. Other identification strategies exist. Pearl’s *do-calculus* is a complete nonparametric graphical identification criterion that includes all types of nonparametric identification as special cases (Pearl 1995).

Conditioning on a (Post-) Outcome Collider

It is well known that outright selection on the outcome, as well as conditioning on a variable affected by the outcome, can lead to bias (e.g., Berkson 1946; Gronau 1974; Heckman 1974, 1976, 1979). Nevertheless, both problems continue to occur in empirical sociological research, where they may result from nonrandom sampling schemes during data collection or from seemingly compelling choices during data analysis. Here, we present canonical and more recent examples, explicating them as endogenous selection bias due to conditioning on a collider.

Sample truncation bias—Hausman & Wise's (1977) influential analysis of sample truncation bias furnishes a classic example of endogenous selection bias. They consider the problem of estimating the effect of education on income from a sample selected (truncated) to contain only low earners. Figure 5 describes the heart of the problem. Income, I , is affected by education, E , as well as by other factors traditionally subsumed in the error term, U . Assume, for clarity and without loss of generality, that E and U are independent in the population (i.e., share no common causes and do not cause each other). Then the marginal association between E and I in the population would identify the causal effect $E \rightarrow I$. The DAG shows, however, that I is a collider variable on the path between treatment and the error term, $E \rightarrow I \leftarrow U$. Restricting the sample to low earners amounts to conditioning on this collider. As a result, there are now two sources of association between E and I : the causal path $E \rightarrow I$, which represents the causal effect of interest, and the newly induced noncausal path $E \leftarrow U \rightarrow I$. Sample truncation has induced endogenous selection bias so that the association between E and I in the truncated sample does not identify the causal effect of E on I .

This example illustrates an important general point: If a treatment, T , has an effect on the outcome, Y , $T \rightarrow Y$, then the outcome is a collider variable on the path between the treatment and the outcome's error term, U , $T \rightarrow Y \leftarrow U$. Therefore, outright selection on an outcome, or selection on a variable affected by the outcome, is always potentially problematic.

Nonresponse bias—The same causal structure that leads to truncation bias also underlies non-response bias in retrospective studies, except that nonresponse is better understood as conditioning on a variable affected by the outcome rather than conditioning on the outcome itself. Consider a simplified version of the example discussed by Lin et al. (1999), who investigate the consequences of nonresponse for estimating the effect of divorced fathers' income on their child support payments (Figure 6). A divorced father's income, I , and the amount of child support he pays, P , both influence whether a father responds to the study, R . (Without loss of generality, we neglect that I and P may share common causes.) Response behavior is thus a collider along a non-causal path between treatment and outcome, $I \rightarrow R \leftarrow P$. Analyzing only completed interviews by dropping nonresponding fathers (listwise deletion) amounts to conditioning on this collider, which induces a noncausal association between fathers' income and the child support they pay, i.e., endogenous selection bias.

Ascertainment bias—Sociological research on elite achievement sometimes employs nonrandom samples when collecting data on the entire population is too expensive or

drawing a random sample is infeasible. Such studies may sample all individuals, organizations, or cultural products that have achieved a certain elite status (e.g., sitting on the Supreme Court, qualifying for the Fortune 500, or receiving critical accolades) and a control group of individuals, organizations, or cultural products that have distinguished themselves in some other manner (e.g., sitting on the federal bench, posting revenues in excess of \$100 million, or being commercially successful). The logic of this approach is to compare the most outstandingly successful cases to somewhat less outstandingly successful cases in order to understand what sets them apart. The problem is that studying the causes of success in a sample selected for success invites endogenous selection bias.

We suspect that this sampling strategy may account for a surprising recent finding about the determinants of critical success in the music industry (Schmutz 2005; see also Allen & Lincoln 2004, Schmutz & Faupel 2010). The study aimed, among other goals, to identify the effects of a music album's commercial success on its chances of inclusion in *Rolling Stone* magazine's vaunted list of *500 Greatest Albums of All Time*. The sample included roughly 1,700 albums: all albums in the *Rolling Stone 500* and 1,200 additional albums, all of which had earned some other elite distinction, such as topping the Billboard charts or winning a critics' poll. Among the tens of thousands of albums released in the United States over the decades, the 1,700 sampled albums clearly represent a subset that is heavily selected for success. A priori, one might expect that outstanding commercial success (i.e., topping the Billboard charts) should increase an album's chances of inclusion in the *Rolling Stone 500* (because the experts compiling the list can only consider albums that they know, probably know all chart toppers, and cannot possibly know all albums ever released). Surprisingly, however, a logistic regression of the analytic sample indicated that topping the Billboard charts was strongly negatively associated with inclusion in the *Rolling Stone 500*.

Why? As the authors noted, this result should be interpreted with caution (Schmutz 2005, p. 1520; Schmutz & Faupel 2010, p. 705). We hypothesize that the finding is due to endogenous selection bias. The DAG in Figure 7 illustrates the argument. (For clarity, we do not consider additional explanatory variables in the study.) The outcome is an album's inclusion in the *Rolling Stone 500*, R . The treatment is topping the Billboard charts, B , which presumably has a positive effect on R in the population of all albums. By construction, both R and B have a strong, strictly positive effect on inclusion in the sample, S . Analyzing this sample means that the analysis conditions on inclusion in the sample, S , which is a collider on a noncausal path between treatment and outcome, $B \rightarrow S \leftarrow R$. Conditioning on this collider changes the association between its immediate causes, R and B , and induces endogenous selection bias. If the sample is sufficiently selective of the population of all albums, then the bias may be large enough to reverse the sign on the estimated effect of B on R , producing a negative estimate for a causal effect that may in fact be positive.

Table 1 makes the same point. For binary B and R , $B = 1$ denotes topping the Billboard charts, $R = 1$ denotes inclusion in the *Rolling Stone 500*, and 0 otherwise. The cell counts a , b , c , and d define the population of all albums. In the absence of common-cause confounding (as encoded in Figure 7), the population odds ratio $OR_p = ad/bc$ gives the true causal effect of B on R . Now note that most albums ever released neither topped the charts

nor won inclusion in the *Rolling Stone 500*. To reduce the data-collection burden, the study then samples on success by including all albums with $R = 1$ and all albums with $B = 1$ in the analysis and severely undersamples less successful albums, $a_S < a$. The odds ratio from the logistic regression of R on B in this selected sample would then be downward biased compared with the true population odds ratio and possibly flip directions, $OR_S = a_S d / bc < ad / bc = OR_P$.

We note that the recovery of causal information despite direct selection on the outcome is sometimes possible if certain constraints can be imposed on the data. This is the domain of the epidemiological literature on case-control studies (Rothman et al. 2008) and the econometric literature on response- or choice-based sampling (Manski 2003, ch. 6). A key rule in case-control studies, however, is that observations should not be sampled (ascertained) as a function of both the treatment and the outcome. The consequence of violating this rule is known as ascertainment bias (Rothman et al. 2008). The endogenous selection bias in the present example is ascertainment bias because control albums are sampled as a function of (a) exclusion from the *Rolling Stone 500* and (b) commercial success. (Greenland et al. 1999 and Hernán et al. 2004 use DAGs to discuss ascertainment bias as endogenous selection bias in a number of interesting biomedical case-control studies.)

Heckman selection bias—Heckman selection bias—perhaps the best-known selection bias in the social sciences—can be explicated as endogenous selection bias. A classic example is the effect of motherhood (i.e., having a child), M , on the wages offered to women by potential employers, W_O (Gronau 1974, Heckman 1974). Figure 8 displays the relevant DAGs. Following standard labor market theory, assume that motherhood will affect a woman's reservation wage, W_R , i.e., the wage that would be necessary to draw her out of the home and into the workforce. Employment, E , is a function of both the wage offer and the reservation wage because a woman will only accept the job if the wage offer meets or exceeds her reservation wage. Therefore, employment is a collider on the path $M \rightarrow W_R \rightarrow E \leftarrow W_O$ between motherhood and offer wages. (For simplicity, assume that M and W_O share no common cause.)

Many social science data sets include information on motherhood, M , but they typically do not include information on women's reservation wage, W_R . Importantly, the data typically include information on offer wages, W_O , only for those women who are actually employed. If the analyst restricts the analysis to employed women, he will thus condition on the collider E , unblock the non-causal path from motherhood to offer wages, $M \rightarrow W_R \rightarrow E \leftarrow W_O$, and induce endogenous selection bias. The analysis would thus indicate an association between motherhood and wages even if the causal effect of motherhood on wages is in fact zero (Figure 8a).

If motherhood has no effect on offer wages, as in Figure 8a, the endogenous selection problem is bad enough. Further complications are introduced by permitting the possibility that motherhood may indeed have an effect on offer wages (e.g., because of mothers' differential productivity compared with childless women or because of employer discrimination). This is shown in Figure 8b by adding the arrow $M \rightarrow W_O$. This slight

modification of the DAG creates a second endogenous selection problem. Given that, as noted above, all outcomes are colliders on the path between treatment and the error term if treatment has an effect on the outcome, conditioning on E now further amounts to conditioning on the descendant of the collider W_O on the path $M \rightarrow \boxed{W_O} \leftarrow \varepsilon$. This induces a noncausal association between motherhood and the error term on offer wages, $M \text{ --- } \varepsilon \rightarrow W_O$. Note that this second—but not the first—endogenous selection problem persists even if the noncollider W_R is measured and controlled.

The distinction between these two endogenous selection problems affords a fresh opportunity for nonparametric inference that, to our knowledge, has not yet been exploited in the applied literature on the motherhood wage penalty: The analyst could nonparametrically test the causal null hypothesis of no effect of motherhood on offer wages by observing and controlling women's reservation wages. If M and W_O are not associated after conditioning on W_R and E , then there is no causal effect of M on W_O as in Figure 8a. Nonparametric identification of the magnitude of the causal effect $M \rightarrow W_O$, however, is made impossible by its very existence, as seen in Figure 8b: Even if W_R is observed and controlled (thus blocking the flow of association along the noncausal path $M \rightarrow \boxed{W_R} \rightarrow \boxed{E} \leftarrow W_O$), conditioning on E would still create the noncausal path $M \text{ --- } \varepsilon \rightarrow W_O$. Therefore, the magnitude of the association between M and W_O after conditioning on W_R and E would be biased for the true magnitude of the causal effect.

One can predict the direction of the bias under additional assumptions, specifically that the causal effects in Figure 8 are in the same direction for all women (VanderWeele & Robins 2009b). If, for example, motherhood decreases the chances of employment (by increasing reservation wages), and if higher offer wages increase the chances of accepting employment for all women, then employed mothers must on average have received higher wage offers than employed childless women. Consequently, an empirical analysis of motherhood on offer wages that is restricted to working women would underestimate the motherhood wage penalty. This is exactly what Gangl & Ziefle's (2009) selectivity-corrected analysis of the motherhood wage penalty finds.

Conditioning on an Intermediate Variable

Methodologists have long warned that conditioning on an intermediate variable that has been affected by the treatment can lead to bias (e.g., Heckman 1976; Berk 1983; Rosenbaum 1984; Holland 1988; Robins 1989, 2001, 2003; Smith 1990; Angrist & Krueger 1999; Wooldridge 2005; Sobel 2008). This bias appears in many guises. It is the source of informative censoring and attrition bias in longitudinal studies, and it is a central problem of mediation analysis in the search for social mechanisms and the estimation of direct and indirect effects. Following Pearl (1998), Robins (2001), and Cole & Hernán (2002), we use DAGs to explicate the problem of conditioning on intermediate variables as endogenous selection bias.

Informative/dependent censoring and attrition bias in longitudinal studies—

Most prospective longitudinal (panel) studies experience attrition over time. Participants that were enrolled at baseline may die, move away, or simply refuse to answer in subsequent

waves of data collection. Over time, the number of cases available for analysis decreases, sometimes drastically so (e.g., Behr et al. 2005, Alderman et al. 2001). When faced with incomplete follow up, some sociological analyses simply analyze the available cases. Event-history or survival-analytic approaches follow cases until they drop out (censoring), whereas other analytic approaches may disregard incompletely followed cases altogether. Either practice can lead to endogenous selection bias.

Building on Hill's (1997) analysis of attrition bias, consider estimating the effect of poverty, P , on divorce, D , in a prospective study with attrition, C . Figure 9 displays various scenarios about the attrition process. Figure 9a shows the most benevolent situation. Here, P affects D , but neither variable is causally related to C —attrition is assumed to be completely random with respect to both treatment and outcome. If this assumption is correct, then analyzing complete cases only, i.e., conditioning on not having attrited, is entirely unproblematic; the causal effect $P \rightarrow D$ is identified without any adjustments for attrition. In Figure 9b, poverty is assumed to affect the risk of dropping out, $P \rightarrow C$. Nevertheless, the causal effect of $P \rightarrow D$ remains identified. As in Figure 9a, conditioning on C in Figure 9b does not lead to bias because conditioning on C does not open any noncausal paths between P and D .⁹ If, however, there are unmeasured factors, such as marital distress, U , that jointly affect attrition and divorce, as in Figure 9c, then C becomes a collider along the noncausal path between poverty and divorce, $P \rightarrow C \leftarrow U \rightarrow D$, and conditioning on C will unblock this path and distort the association between poverty and divorce. Attrition bias thus is nothing other than endogenous selection bias. Note that even sizeable attrition per se is not necessarily a problem for the identification of causal effects. Rather, attrition is problematic if conditioning on attrition opens a noncausal path between treatment and outcome, as in Figure 9c (for more elaborate attrition scenarios, see Hernán et al. 2004).

Imperfect proxy measures affected by the treatment—Much of social science research studies the effects of schooling on outcomes such as cognition (e.g., Coleman et al. 1982), marriage (e.g., Gullickson 2006, Raymo & Iwasawa 2005), and wages (e.g., Griliches & Mason 1972, Leigh & Ryan 2008, Amin 2011). Schooling, of course, is not randomly assigned to students: On average, individuals with higher innate ability are both more likely to obtain more schooling and more likely to achieve favorable outcomes later in life regardless of their level of schooling. Absent convincing measures to control for confounding by innate ability, researchers often resort to proxies for innate ability, such as IQ-type cognitive test scores. Using DAGs, we elaborate econometric discussions of this proxy-control strategy (e.g., Angrist & Krueger 1999; Wooldridge 2002, pp. 63–70) by differentiating between three separate problems, including endogenous selection bias.

To be concrete, consider the effect of schooling, S , on wages, W . Figure 10a starts by positing the usual assumption that innate ability, U , confounds the effect of schooling on wages via the unblocked noncausal path $S \leftarrow U \rightarrow W$. Because true innate ability is unobserved, the path cannot be blocked, and the effect $S \rightarrow W$ is not identified. Next, one

⁹One may ask whether conditioning on C in Figure 9b would threaten the identification of the population-average causal effect if the effect of P on D varies in the population. Under the model of Figure 9b, this is not the case because the DAG implies that the effect of P on D does not vary systematically with C (VanderWeele & Robins 2007b).

might control for measured test scores, Q , as a proxy for the unobserved U , as shown in Figure 10*b*. To the extent that U strongly affects Q , Q is a valid proxy for U . But to the extent that Q does not perfectly measure U , the effect of S on W remains partially confounded. This is the first problem of proxy control—the familiar issue of residual confounding when an indicator (i.e., test scores) imperfectly captures the desired underlying construct (i.e., ability). Nevertheless, under the assumptions of Figure 10*b*, controlling for Q will remove at least some of the confounding bias exerted by U and not introduce any new biases.

The second problem, endogenous selection bias, enters the picture in Figure 10*c*. Cognitive test scores are sometimes measured after the onset of schooling. This is a problem because schooling is known to affect students' test scores (Winship & Korenman 1997), as indicated by the addition of the arrow $S \rightarrow Q$. This makes Q into a collider along the noncausal path from schooling to wages, $S \rightarrow Q \leftarrow U \rightarrow W$. Conditioning on Q will open this noncausal path and thus lead to endogenous selection bias.

The third distinct problem results from the possibility that Q may itself exert a direct causal effect on W , $Q \rightarrow W$, perhaps because employers use test scores in the hiring process and reward high test scores with better pay regardless of the applicant's schooling. If so, the causal effect of schooling on wages would be partially mediated by test scores along the causal path $S \rightarrow Q \rightarrow W$, and controlling for Q would block this causal path, leading to overcontrol bias.

The confluence of these three problems leaves the analyst in a quandary. If Q is measured after the onset of schooling, should the analyst control for Q to remove (by proxy) some of the confounding owed to U ? Or should she not control for Q to avoid inducing endogenous selection bias and overcontrol bias? Absent detailed knowledge of the relative strengths of all effects in the DAG—which is rarely available, especially where unobserved variables are involved—it is difficult to determine which problem will dominate empirically and hence which course of action should be taken.¹⁰

Direct effects, indirect effects, causal mechanisms, and mediation analysis—

Common approaches to estimating direct and indirect effects (Baron & Kenny 1986) are highly susceptible to endogenous selection bias. We first look at the estimation of direct effects. To fix ideas, consider an analysis of the Project STAR class-size experiment that asks whether class size in first grade, T , has a direct effect on high school graduation, Y , via some mechanism other than boosting student achievement in third grade, M (Finn et al. 2005). Figure 11*a* gives the corresponding DAG. Because class size is randomized, T and Y share no common cause, and the total effect of treatment, T , on the outcome, Y , is identified by their marginal association. The post-treatment mediator, M , however, is not randomized and may therefore share an unmeasured cause, U , with the outcome, Y . Candidates for U might include parental education, student motivation, and any other confounders of M and Y not explicitly controlled in the study. The existence of such a variable U would make M a

¹⁰If Q is measured before the onset of schooling, then the analyst can probably assume the DAG in Figure 10*b* (and certainly rule out Figure 10*c* and *d*). If Figure 10*b* is true, then conditioning on Q is safe and advisable.

collider variable along the noncausal path $T \rightarrow \overline{M} \leftarrow U \rightarrow Y$ between treatment and outcome. Conditioning on M (e.g., by regressing Y on T and M) in order to estimate the direct causal effect $T \rightarrow Y$ would unblock this noncausal path and induce endogenous selection bias. Hence, the direct effect of T on Y is not identified.¹¹

Attempts to detect the presence of indirect, or mediated, effects by comparing estimates for the total effect of T on Y and estimates for the direct effect of T on Y (Baron & Kenny 1986) are similarly susceptible to endogenous selection bias. Suppose, for argument's sake, that achievement in third grade has no causal effect on high school graduation (Figure 11b). The total effect of T on Y would then be identical with the direct effect of T on Y because no indirect effect of T on Y via M exists. The total effect is identified by the marginal association between T and Y . The conditional association between T and Y given M , however, would differ from the total effect of T on Y because M is a collider, and conditioning on the collider induces a noncausal association between T and Y via $T \rightarrow \overline{M} \leftarrow U \rightarrow Y$. Thus, the (correct) estimate for the total causal effect and the (naive and biased) estimate for the direct effect of T on Y would differ, and the analysts would falsely conclude that an indirect effect exists even if it does not.

The endogenous selection problem of mediation analysis is of particular concern for current empirical practice in sociology because conventional efforts to control for unobserved confounding typically focus on the confounders of the main treatment but not on the confounders of the mediator. The modern literature on causal mediation analysis discusses estimands and nonparametric identification conditions (Robins & Greenland 1992; Pearl 2001, 2012; Sobel 2008; Shpitser & Vander-Weele 2011) as well as parametric and nonparametric estimation strategies (VanderWeele 2009a, 2011a; Imai et al. 2010; Pearl 2012).

Conditioning on a Pre-Treatment Collider

Controlling for pre-treatment variables can sometimes increase rather than decrease bias. One class of pre-treatment variables that analysts should approach with caution is pretreatment colliders (Pearl 1995, Greenland et al. 1999b, Hernán et al. 2002, Greenland 2003, Hernán et al. 2004, Elwert 2013).

Homophily bias in social network analysis—A classic problem of causal inference in social networks is that socially connected individuals may exhibit similar behaviors not because one individual influences the other (causation) but because individuals who are similar tend to form social ties with each other (homophily) (Farr 1858). Distinguishing between homophily and interpersonal causal effects (also called peer effects, social contagion, network influence, induction, or diffusion) is especially challenging if the sources of homophily are unobserved (latent). Shalizi & Thomas (2011) recently showed that latent homophily bias constitutes a previously unrecognized example of endogenous selection bias, where the social tie itself plays the role of the pre-treatment collider.

¹¹In this DAG, it is similarly impossible to identify nonparametrically the direct effect $T \rightarrow Y$ from an estimate of the total causal effect of T on Y minus the product of the estimates for the direct effects $T \rightarrow M$ and $M \rightarrow Y$. This difference method would fail because the effect $M \rightarrow Y$ is confounded by U and hence not identified in its own right.

Consider, for example, the spread of civic engagement between friends in dyads generically indexed (i, j) (Figure 12). The question is whether the civic engagement of individual i at time t , $Y_i(t)$, causes the subsequent civic engagement of individual j , $Y_j(t+1)$. For expositional clarity, assume that civic engagement does not spread between friends [i.e., no arrow $Y_i(t) \rightarrow Y_j(t+1)$]. Instead, each person's civic engagement is caused by their own characteristics, U , such as altruism, which are at least partially unobserved [giving arrows $Y_i(t) \leftarrow U_i \rightarrow Y_i(t+1)$ and $Y_j(t) \leftarrow U_j \rightarrow Y_j(t+1)$]. Tie formation is homophilous, such that altruistic people form preferential attachments, $U_i \rightarrow F_{i,j} \leftarrow U_j$. Hence, $F_{i,j}$ is a collider. The problem, then, is that the very act of computing the association between friends means that the analysis conditions on the existence of the social tie, $F_{i,j} = 1$. Conditioning on $F_{i,j}$ opens a non-causal path between treatment and outcome, $Y_i(t) \leftarrow U_i \rightarrow \boxed{F_{i,j}} \leftarrow U_j \rightarrow Y_j(t+1)$, which results in an association between i 's and j 's civic engagement even if one does not cause the other.

In sum, if tie formation or tie dissolution is affected by unobserved variables that are associated with, respectively, the treatment variable in one individual and the outcome variable in another individual, then searching for interpersonal effects will induce a spurious association between individuals in the network, i.e., endogenous selection bias.

One solution to this problem is to model tie formation (and dissolution) explicitly (Shalizi & Thomas 2011). In the model of Figure 12, this can be accomplished by measuring and conditioning on the tie-forming characteristics of individual i , U_i , or individual j , U_j , or both. Other solutions for latent homophily bias are proliferating in the literature. For example, Elwert & Christakis (2008) introduce a proxy strategy to measure and subtract homophily bias; Ver Steeg & Galstyan (2011) develop a formal test for latent homophily; VanderWeele (2011b) introduces a formal sensitivity analysis; O'Malley et al. (2014) explore instrumental variables solutions; and VanderWeele & An (2013) provide an extensive survey.

What pre-treatment variables should be controlled?—Pre-treatment colliders have practical and conceptual implications beyond social network analysis. Empirical examples exist where controlling for pre-treatment variables demonstrably increases rather than decreases bias. For example, Steiner et al. (2010) report a within-study comparison between an observational study and an experimental benchmark of the effect of an educational intervention on student test scores. In the observational study, the bias in the estimated effect of treatment on math scores increased by about 30% when controlling only for pre-treatment psychological disposition or only for pre-treatment vocabulary test scores. Furthermore, they found that although controlling for additional pre-treatment variables generally reduced bias, sometimes including additional controls increased bias. It is possible that these increases in bias resulted from controlling for pre-treatment colliders.

This raises larger questions about control-variable selection in observational studies. Which pre-treatment variables should an analyst control, and which pre-treatment variables are better left alone? An inspection of the DAGs in Figure 13 demonstrates that these questions cannot be settled without recourse to a theory of how the data were generated. Start by comparing the DAGs in Figure 13a and b. In both models, the pre-treatment variable X meets the commonsense, purely associational, definition of confounding: (a) X temporally

precedes treatment, T ; (b) X is associated with the treatment; and (c) X is associated with the outcome, Y . The conventional advice would be to condition on X in both situations.

But this conventional advice is incorrect. The proper course of action differs sharply depending on how the data were generated (Greenland & Robins 1986, Weinberg 1993, Cole & Hernán 2002; Hernán et al. 2002, Pearl 2009). In Figure 13a, X is indeed a common-cause confounder as it sits on an open noncausal path between treatment and outcome, $T \leftarrow X \rightarrow Y$. Conditioning on X will block this noncausal path and identify the causal effect of T on Y . By contrast, in Figure 13b, X is a collider that blocks the noncausal path between treatment and outcome, $T \leftarrow U_2 \rightarrow X \leftarrow U_1 \rightarrow Y$. Conditioning on X would open this noncausal path and induce endogenous selection bias. Therefore, the analyst should condition on X if the data were generated according to Figure 13a, but she should not condition on X if the data were generated according to Figure 13b.

If U_1 and U_2 are unobserved, the analyst cannot distinguish empirically between Figure 13a and b because both models have identical observational implications. In both models, all combinations of X , T , and Y are marginally and conditionally associated with each other. Therefore, the only way to decide whether or not to control for X is to decide, on a priori grounds, which of the two causal models accurately represents the data-generating process. In other words, the selection of control variables requires a theoretical commitment to a model of data generation.

In practice, sociologists may lack a fully articulated theory of data generation. It turns out, however, that partial theory suffices to determine the proper set of pre-treatment control variables. If the analyst is willing to assume that controlling for some combination of observed pre-treatment variables is sufficient to identify the total causal effect of treatment, then it is sufficient to control for all variables that (directly or indirectly) cause treatment or outcome or both (VanderWeele & Shpitser 2011).

The virtue of this advice is that it prevents the analyst from conditioning on a pretreatment collider. In Figure 13a, it accurately counsels controlling for X (because X is a cause of both treatment and outcome), and in Figure 13b it accurately counsels not controlling for X (because X causes neither treatment nor outcome).

Nevertheless, the advice is still conditional on some theory of data generation, however coarsely articulated, specifically that there exists some sufficient set of observed controls. Note that no such set exists in Figure 13c. Here, X is both a confounder (on the noncausal paths $T \leftarrow X \leftarrow U_1 \rightarrow Y$ and $T \leftarrow X \leftarrow U_2 \rightarrow Y$) and a collider (on the noncausal path $T \leftarrow U_2 \rightarrow X \leftarrow U_1 \rightarrow Y$). Controlling for X would remove confounding but induce endogenous selection, and not controlling for X would do the opposite. In short, there exists no set of pre-treatment controls that is sufficient for identifying the causal effect of T on Y , and hence the VanderWeele & Shpitser (2011) rule does not apply. When all variables in the DAG are binary, Greenland (2003) suggests that controlling for a variable that is both a confounder and a collider, as in Figure 13c, likely (but not necessarily) removes more bias than it creates.¹²

6. DISCUSSION AND CONCLUSION

In this review, we have argued that understanding endogenous selection bias is central for understanding the identification of causal effects in sociology. Drawing on recent work in theoretical computer science (Pearl 1995, 2009) and epidemiology (Hernán et al. 2004), we have explicated endogenous selection as conditioning on a collider variable in a DAG. Interrogating the simple steps necessary for separating causal from noncausal associations, we next saw that endogenous selection bias is logically on par with confounding bias and overcontrol bias in the sense that all nonparametric identification problems can be reduced to confounding (i.e., not conditioning on a common cause), overcontrol (i.e., conditioning on a variable on the causal pathway), endogenous selection (i.e., conditioning on a collider), or a mixture of the three (Verma & Pearl 1988).

Methodological warnings are most useful if they are phrased in an accessible language (Pearl 2009, p. 352). We believe that DAGs provide such a language because they focus the analyst's attention on what matters for nonparametric identification of causal effects: qualitative assumptions about the causal relationships between variables. This in no way denies the place of algebra and parametric assumptions in social science methodology. But inasmuch as algebraic presentations act as barriers to entry for applied researchers, phrasing methodological problems graphically may increase awareness of these problems in practice.

In the main part of this review, we have shown how numerous problems in causal inference can be explicated as endogenous selection bias. A key insight is that conditioning on an endogenous variable (a collider variable), no matter where it temporally falls relative to treatment and outcome, can induce new, noncausal, associations that are likely to result in biased estimates. We have distinguished three groups of endogenous selection problems by the timing of the collider relative to treatment and outcome. The most general advice to be derived from the first two groups—post-outcome and intermediate collider examples—is straightforward: When estimating total causal effects, avoid conditioning on post-treatment variables (Rosenbaum 1984). Of course, this advice can be hard to follow in practice when conditioning is implicit in the data-collection scheme or the result of selective nonresponse and attrition. Even where it is easy to follow, it can be hard to swallow if the avowed scientific interest concerns the identification of causal mechanisms, mediation, or direct or indirect effects, which require peering into the post-treatment black box. New solutions for the estimation of causal mechanisms that rely on a combination of great conceptual care, a new awareness of underlying assumptions, and powerful sensitivity analyses, however, are fast becoming available (e.g., Robins & Greenland 1992; Pearl 2001; Frangakis & Rubin 2002; Sobel 2008; VanderWeele 2008a, 2009b, 2010; Imai et al. 2010; Shpitser & VanderWeele 2011).

Advice for handling pre-treatment colliders is harder to come by. If it is known that the pre-treatment collider is not also a confounder, then the solution is not to condition on it when possible (this is not possible in social network analysis, where conditioning on the

¹²Causal inference for time-varying treatments often encounters variables that are both colliders and confounders (Pearl & Robins 1995). See Elwert (2013), Sharkey & Elwert (2011), and Wodtke et al. (2011) for sociological examples of time-varying treatments in education and neighborhood effects research that have statistical solutions (Robins 1999).

pretreatment social tie is implicit in the research question). If the pre-treatment collider is also a confounder, then conditioning on it will generally, but not always, remove more bias than it induces (Greenland 2003), especially if the analysis further conditions on many other pre-treatment variables (Steiner et al. 2010). In some circumstances, instrumental variable estimation may be feasible, though researchers should be aware of its limitations (Angrist et al. 1996).

One of the remaining difficulties in dealing with endogenous selection (if it cannot be entirely avoided) is that sign and size of the bias are difficult to predict. Absent strong parametric assumptions, size and sign of the bias will depend on the exact structure of the DAG, the sizes of the effects and their variation across individuals, and the distribution of the variables (Greenland 2003).¹³ Conventional intuition about size and sign of bias in causal systems can break down in unpredictable ways outside of the stylized assumptions of linear models.¹⁴ Nevertheless, statisticians have in recent years derived important (and often complicated) results that have proven quite powerful in applied work. For example, VanderWeele & Robins (2007a, 2009a,b) derive the sign of endogenous selection bias for DAGs with binary variables and monotonic effects with and without interactions. Hudson et al. (2008) apply these results productively in a study of familial coaggregation of diseases.

Greenland (2003), in particular, has derived useful analytic results and approximations for the relative size of endogenous selection bias and confounding bias in DAGs with binary variables (see also Kaufman et al. 2009). His three rules of thumb match nicely to our examples of post-outcome, intermediate, and pre-treatment colliders. First, conditioning on post-outcome colliders that are affected by treatment and outcome, $T \rightarrow C \leftarrow Y$, induces the same bias as would the corresponding confounding bias if the directions of the arrows into the collider were reversed, $T \leftarrow C \rightarrow Y$. Second, conditioning on an intermediate collider may lead to bias that is larger than the effects pointing into the collider, although the bias will usually be smaller than those effects if the intermediate collider does not itself exert a causal effect on the outcome. Finally, conditioning on a pre-treatment collider (as in Figures 12 and 13) will usually lead to a small bias, which is why recognizing latent homophily bias as endogenous selection bias (as opposed to confounding bias) should be good news to network analysts. Furthermore, if the pre-treatment variable is both a collider and a confounder (as in Figure 13c), then conditioning on it may remove more bias than it induces, even though nonparametric point identification is still not feasible. We caution, however, that these rules of thumb for binary-variable DAGs do not necessarily generalize to DAGs with differently distributed variables. More research is needed to understand the direction and magnitude of endogenous selection bias in realistic social science settings.

Empirical research in sociology pursues a variety of goals, including description, interpretation, and causal inference. Given that assumption-free causal inference is

¹³The same difficulties pertain to nonparametrically predicting size and sign of confounding bias (e.g., VanderWeele et al. 2008, VanderWeele 2008b). The conventional omitted variables formula in linear regression merely obscures these difficulties via the assumptions of linearity and constant effects (e.g., Wooldridge 2002).

¹⁴For example, the sign of direct effects is not always transitive (VanderWeele et al. 2008). In the DAG $A \rightarrow B \rightarrow C$, the average causal effect $A \rightarrow B$ and the average causal effect $B \rightarrow C$ may both be positive, and yet the average causal effect $A \rightarrow C$ could be negative. Sign transitivity in conventional structural equation models (Alwin & Hauser 1975) is an artifact of the linearity and constant-effects assumptions.

impossible, sociologists must rely on theory, prudent substantive judgment, and well-founded prior knowledge to assess the threat of bias. Using DAGs, sociologists can formally analyze the consequences of their theoretical assumptions and determine whether causal inference is possible. Such analyses will strengthen confidence in a causal claim when it is warranted and explicitly indicate sources of potential bias—including endogenous selection bias—when a causal claim is not appropriate.

Acknowledgments

We thank Miguel Hernán and Jamie Robins for inspiration and Jennie Brand, Richard Breen, David Harding, Kristian Karlson, Therese Leung, Steve Morgan, Robert Mare, Chris Muller, Gabriel Rossman, Nora Cate Schaeffer, Michael Sobel, Peter Steiner, Tyler VanderWeele, James Walker, Geoffrey Wodtke, and Elizabeth Wrigley-Field for discussions and advice. Laurie Silverberg, Genevieve Butler, and Janet Clear provided editorial assistance.

Glossary

Endogenous selection bias

this type of bias results from conditioning on a collider (or its descendant) on a noncausal path linking treatment and outcome

Conditioning

introducing information about a variable into an analysis, e.g., via sample selection, stratification, regression control

Collider variable

a common outcome of two other variables along a path; colliders are path-specific

Common-cause confounding bias

this type of bias results from failure to condition on a common cause of treatment and outcome

Identification

the possibility of recovering a causal effect from ideal data generated by the assumed data-generating process

Overcontrol bias

this type of bias results from conditioning on a variable on a causal path between treatment and outcome

Direct acyclic graph (DAG)

a graphical representation of the causal assumptions in the data-generating process

Path

a sequence of arrows linking two variables regardless of the direction of the arrows

Causal path

a path linking treatment and outcome where all arrows point away from the treatment and toward the outcome

Noncausal path

a path linking treatment and outcome that is not a causal path

d-separation

translates between the causal assumptions about the data-generating process and observable associations

LITERATURE CITED

- Alderman H, Behrman J, Kohler H, Maluccio JA, Watkins SC. Attrition in longitudinal household survey data: some tests for three developing-country samples. *Demogr Res.* 2001; 5:79–124.
- Allen MP, Lincoln A. Critical discourse and the cultural consecration of American films. *Soc Forces.* 2004; 82(3):871–94.
- Alwin DH, Hauser RM. The decomposition of effects in path analysis. *Am Sociol Rev.* 1975; 40:37–47.
- Amin V. Returns to education: evidence from UK twins: comment. *Am Econ Rev.* 2011; 101(4):1629–35.
- Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. *J Am Stat Assoc.* 1996; 8:328–36.
- Angrist JD, Krueger AB. Empirical strategies in labor economics. In: Ashenfelter O, Card D, editors *Handbook of Labor Economics*. Vol. 3. Amsterdam: Elsevier; 1999. 1277–366.
- Bareinboim E, Pearl J. Controlling selection bias in causal inference. In: Lawrence N, Girolami M, editors *UCLA Cogn Syst Lab, Tech Rep R-381 Proc 15th Int Conf Artif Intell Stat (AISTATS)*; April 21–23, 2012; La Palma, Canary Islands. Brookline, MA: Microtome; 2012. 100–8. http://ftp.cs.ucla.edu/pub/stat_ser/r381.pdf
- Baron RM, Kenny DA. The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *J Personal Soc Psychol.* 1986; 51:1173–82.
- Behr A, Bellgardt E, Rendtel U. Extent and determinants of panel attrition in the European Community Household Panel. *Eur Sociol Rev.* 2005; 21(5):489–512.
- Berk RA. An introduction to sample selection bias in sociological data. *Am Sociol Rev.* 1983; 48(3):386–98.
- Berkson J. Limitations of the application of fourfold tables to hospital data. *Biometr Bull.* 1946; 2(3):47–53. [PubMed: 21001024]
- Blalock H. *Causal Inferences in Nonexperimental Research*. Chapel Hill: Univ. N.C. Press; 1964.
- Bollen KA. *Structural Equations with Latent Variables*. New York: Wiley; 1989.
- Bollen KA, Pearl J. Eight myths about causality and structural equation models. In: Morgan SL, editor *Handbook of Causal Analysis for Social Research*. Dordrecht, Neth: Springer; 2013. 301–28.
- Christofides LN, Li Q, Liu Z, Min I. Recent two-stage sample selection procedures with an application to the gender wage gap. *J Bus Econ Stat.* 2003; 21:396–405.
- Cole SR, Hernán MA. Fallibility in the estimation of direct effects. *Int J Epidemiol.* 2002; 31:163–65. [PubMed: 11914314]
- Coleman JS, Hoffer T, Kilgore S. *High School Achievement: Public, Catholic, and Private Schools Compared*. New York: Basic Books; 1982.
- Duncan OD. Path analysis: sociological examples. *Am J Sociol.* 1966; 72(1):1–16.
- Elwert F. Graphical causal models. In: Morgan SL, editor *Handbook of Causal Analysis for Social Research*. Dordrecht, Neth: Springer; 2013. 245–73.
- Elwert F, Christakis NA. Wives and ex-wives: a new test for homogamy bias in the widowhood effect. *Demography.* 2008; 45(4):851–73. [PubMed: 19110901]
- Farr W. Influence of marriage on the mortality of the French people. In: Hastings GW, editor *Transaction National Association Promotion Social Science*. London: John W. Park & Son; 1858. 504–13.

- Finn JD, Gerber SB, Boyd-Zaharias J. Small classes in the early grades, academic achievement, and graduating from high school. *J Educ Psychol*. 2005; 97(2):214–23.
- Frangakis CE, Rubin DB. Principal stratification in causal inference. *Biometrics*. 2002; 58:21–29. [PubMed: 11890317]
- Fu V, Winship C, Mare R. Sample selection bias models. In: Hardy M, Bryman A, editors *Handbook of Data Analysis*. London: Sage; 2004. 409–30.
- Gangl M, Ziefle A. Motherhood, labor force behavior, and women's careers: An empirical assessment of the wage penalty for motherhood in Britain, Germany, and the United States. *Demography*. 2009; 46(2):341–69. [PubMed: 21305397]
- Glymour MM, Greenland S. Causal diagrams. In: Rothman KJ, Greenland S, Lash T, editors *Modern Epidemiology*. 3rd. Philadelphia: Lippincott; 2008. 183–209.
- Grasdal A. The performance of sample selection estimators to control for attrition bias. *Health Econ*. 2001; 10(5):385–98. [PubMed: 11466801]
- Greenland S. Quantifying biases in causal models: classical confounding versus collider-stratification bias. *Epidemiology*. 2003; 14:300–6. [PubMed: 12859030]
- Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology*. 1999a; 10:37–48. [PubMed: 9888278]
- Greenland S, Robins JM. Identifiability, exchangeability and epidemiological confounding. *Int J Epidemiol*. 1986; 15:413–19. [PubMed: 3771081]
- Greenland S, Robins JM, Pearl J. Confounding and collapsibility in causal inference. *Stat Sci*. 1999b; 14:29–46.
- Griliches Z, Mason WM. Education, income, and ability. *J Polit Econ*. 1972; 80(3):S74–103.
- Gronau R. Wage comparisons—a selectivity bias. *J Polit Econ*. 1974; 82:1119–44.
- Gullickson A. Education and black-white interracial marriage. *Demography*. 2006; 43(4):673–89. [PubMed: 17236541]
- Hausman JA, Wise DA. Social experimentation, truncated distributions and efficient estimation. *Econometrica*. 1977; 45:919–38.
- Hausman JA, Wise DA. Stratification on endogenous variables and estimation. In: Manski C, McFadden D, editors *The Econometrics of Discrete Data*. Cambridge, MA: MIT Press; 1981. 365–91.
- Heckman JJ. Shadow prices, market wages and labor supply. *Econometrica*. 1974; 42(4):679–94.
- Heckman JJ. The common structure of statistical models of truncation, sample selection, and limited dependent variables and a simple estimator for such models. *Ann Econ Soc Meas*. 1976; 5:475–92.
- Heckman JJ. Selection bias as a specification error. *Econometrica*. 1979; 47:153–61.
- Hernán MA, Hernández-Díaz S, Robins JM. A structural approach to section bias. *Epidemiology*. 2004; 15:615–25. [PubMed: 15308962]
- Hernán MA, Hernández-Díaz S, Werler MM, Robins JM, Mitchell AA. Causal knowledge as a prerequisite of confounding evaluation: an application to birth defects epidemiology. *Am J Epidemiol*. 2002; 155(2):176–84. [PubMed: 11790682]
- Hill DH. Adjusting for attrition in event-history analysis. *Sociol Methodol*. 1997; 27:393–416. [PubMed: 12348200]
- Holland PW. Statistics and causal inference (with discussion). *J Am Stat Assoc*. 1986; 81:945–70.
- Holland PW. Causal inference, path analysis, and recursive structural equation models. *Sociol Methodol*. 1988; 18:449–84.
- Hudson JI, Javaras KN, Laird NM, VanderWeele TJ, Pope HG, et al. A structural approach to the familial coaggregation of disorders. *Epidemiology*. 2008; 19:431–39. [PubMed: 18379420]
- Imai K, Keele L, Yamamoto T. Identification, inference, and sensitivity analysis for causal mediation effects. *Stat Sci*. 2010; 25(1):51–71.
- Kaufman S, Kaufman J, MacLenose R. Analytic bounds on causal risk differences in directed acyclic graphs involving three observed binary variables. *J Stat Plan Inference*. 2009; 139:3473–87. [PubMed: 20161106]

- Kim JH, Pearl J. A computational model for combined causal and diagnostic reasoning in inference systems. In: Bundy A, editor *Proc 8th Int Jt Conf Artif Intell (IJCAI-83)*; Karlsruhe, FRG. Aug 8–12, 1983; San Francisco: Morgan Kaufmann; 1983. 190–93.
- Leigh A, Ryan C. Estimating returns to education using different natural experiment techniques. *Econ Educ Rev*. 2008; 27(2):149–60.
- Lin I, Schaeffer NC, Seltzer JA. Causes and effects of nonparticipation in a child support survey. *J Off Stat*. 1999; 15(2):143–66.
- Manski C. *Partial Identification of Probability Distributions*. New York: Springer; 2003.
- Morgan SL, Winship C. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. New York: Cambridge Univ. Press; 2007.
- O'Malley AJ, Elwert F, Rosenquist JN, Zaslavsky AM, Christakis NA. Estimating peer effects in longitudinal dyadic data using instrumental variables. *Biometrics*. 2014; In press. doi: 10.1111/biom.12172
- Pagan A, Ullah A. *Nonparametric Econometrics*. Cambridge, UK: Cambridge Univ. Press; 1997.
- Pearl J. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann; 1988.
- Pearl J. Causal diagrams for empirical research. *Biometrika*. 1995; 82(4):669–710.
- Pearl J. Graphs, causality, and structural equation models. *Sociol Methods Res*. 1998; 27(2):226–84.
- Pearl J. Direct and indirect effects. In: Breese J, Koller D, editors *Proc 17th Conf Uncertain Artif Intell*; Aug 2–5, 2001; Seattle, WA. San Francisco: Morgan Kaufmann; 2001. 411–20.
- Pearl J. *Causality: Models, Reasoning, and Inference*. 2nd. New York: Cambridge Univ. Press; 2009.
- Pearl J. The foundations of causal inference. *Sociol Methodol*. 2010; 40:75–149.
- Pearl J. The causal mediation formula—a guide to the assessment of pathways and mechanisms. *Prev Sci*. 2012; 13(4):426–36. [PubMed: 22419385]
- Pearl J, Robins JM. Probabilistic evaluation of sequential plans from causal models with hidden variables. In: Besnard P, Hanks S, editors *Uncertainty in Artificial Intelligence 11*. San Francisco: Morgan Kaufmann; 1995. 444–53.
- Raymo JM, Iwasawa M. Marriage market mismatches in Japan: an alternative view of the relationship between women's education and marriage. *Am Sociol Rev*. 2005; 70(5):801–22.
- Robins JM. A new approach to causal inference in mortality studies with a sustained exposure period: application to the health worker survivor effect. *Math Model*. 1986; 7:1393–512.
- Robins JM. The control of confounding by intermediate variables. *Stat Med*. 1989; 8:679–701. [PubMed: 2749074]
- Robins JM. Correcting for non-compliance in randomized trials using structural nested mean models. *Commun Stat -Theory Methods*. 1994; 23:2379–412.
- Robins JM. Association, causation, and marginal structural models. *Synthese*. 1999; 121:151–79.
- Robins JM. Data, design, and background knowledge in etiologic inference. *Epidemiology*. 2001; 23(3):313–20.
- Robins JM. Semantics of causal DAG models and the identification of direct and indirect effects. In: Green P, Hjort NL, Richardson S, editors *Highly Structured Stochastic Systems*. New York: Oxford Univ. Press; 2003. 70–81.
- Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*. 1992; 3:143–55. [PubMed: 1576220]
- Robins JM, Wasserman L. On the impossibility of inferring causation from association without background knowledge. In: Glymour CN, Cooper GG, editors *Computation, Causation, and Discovery*. Cambridge, MA: AAAI/MIT Press; 1999. 305–21.
- Rosenbaum PR. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *J R Stat Soc Ser A*. 1984; 147(5):656–66.
- Rothman KJ, Greenland S, Lash TL. Case-control studies. In: Rothman KJ, Greenland S, Lash TL, editors *Modern Epidemiology*. 3rd. Philadelphia, PA: Lippincott; 2008. 111–27.
- Rubin DB. Estimating causal effects of treatments in randomized and non-randomized studies. *J Educ Psychol*. 1974; 66:688–701.

- Schmutz V. Retrospective cultural consecration in popular music. *Am Behav Sci.* 2005; 48(11):1510–23.
- Schmutz V, Faupel A. Gender and cultural consecration in popular music. *Soc Forces.* 2010; 89(2): 685–708.
- Shalizi CR, Thomas AC. Homophily and contagion are generically confounded in observational social network studies. *Sociol Methods Res.* 2011; 40:211–39. [PubMed: 22523436]
- Sharkey P, Elwert F. The legacy of disadvantage: multigenerational neighborhood effects on cognitive ability. *Am J Sociol.* 2011; 116(6):1934–81. [PubMed: 21932471]
- Shpitser I, VanderWeele TJ. A complete graphical criterion for the adjustment formula in mediation analysis. *Int J Biostat.* 2011; 7:16. [PubMed: 21556286]
- Smith HL. Specification problems in experimental and nonexperimental social research. *Sociol Methodol.* 1990; 20:59–91.
- Sobel ME. Identification of causal parameters in randomized studies with mediating variables. *J Educ Behav Stat.* 2008; 33(2):230–51.
- Spirtes P, Glymour CN, Scheines R. *Causation, Prediction, and Search.* 2nd. New York: Springer; 2000.
- Steiner PM, Cook TD, Shadish WR, Clark MH. The importance of covariate selection in controlling for selection bias in observational studies. *Psychol Methods.* 2010; 15(3):250–67. [PubMed: 20822251]
- Stolzenberg RM, Relles DA. Tools for intuition about sample selection bias and its correction. *Am Sociol Rev.* 1997; 62(3):494–507.
- VanderWeele TJ. Simple relations between principal stratification and direct and indirect effects. *Stat Probab Lett.* 2008a; 78:2957–62.
- VanderWeele TJ. The sign of the bias of unmeasured confounding. *Biometrics.* 2008b; 64:702–6. [PubMed: 18177462]
- VanderWeele TJ. Marginal structural models for the estimation of direct and indirect effects. *Epidemiology.* 2009a; 20:18–26. [PubMed: 19234398]
- VanderWeele TJ. Mediation and mechanism. *Eur J Epidemiol.* 2009b; 24:217–24. [PubMed: 19330454]
- VanderWeele TJ. Bias formulas for sensitivity analysis for direct and indirect effects. *Epidemiology.* 2010; 21:540–51. [PubMed: 20479643]
- VanderWeele TJ. Causal mediation analysis with survival data. *Epidemiology.* 2011a; 22:582–85. [PubMed: 21642779]
- VanderWeele TJ. Sensitivity analysis for contagion effects in social networks. *Sociol Methods Res.* 2011b; 40:240–55. [PubMed: 25580037]
- VanderWeele TJ, An W. Social networks and causal inference. In: Morgan SL, editor *Handbook of Causal Analysis for Social Research.* Dordrecht, Neth: Springer; 2013. 353–74.
- VanderWeele TJ, Hernán MA, Robins JM. Causal directed acyclic graphs and the direction of unmeasured confounding bias. *Epidemiology.* 2008; 19:720–28. [PubMed: 18633331]
- VanderWeele TJ, Robins JM. Directed acyclic graphs, sufficient causes, and the properties of conditioning on a common effect. *Am J Epidemiol.* 2007a; 166(9):1096–104. [PubMed: 17702973]
- VanderWeele TJ, Robins JM. Four types of effect modification: a classification based on directed acyclic graphs. *Epidemiology.* 2007b; 18(5):561–68. [PubMed: 17700242]
- VanderWeele TJ, Robins JM. Minimal sufficient causation and directed acyclic graphs. *Ann Stat.* 2009a; 37:1437–65.
- VanderWeele TJ, Robins JM. Properties of monotonic effects on directed acyclic graphs. *J Mach Learn.* 2009b; 10:699–718.
- VanderWeele TJ, Shpitser I. A new criterion for confounder selection. *Biometrics.* 2011; 67:1406–13. [PubMed: 21627630]
- Vella F. Estimating models with sample selection bias: A survey. *J Hum Resour.* 1998; 33:127–69.

- Ver Steeg G, Galstyan A. A sequence of relaxations constraining hidden variable models. In: Cozman FG, Pfeffer A, editors *Proc. 27th Conf. Uncertain. Artif. Intell. (UAI2011)*; July 14–17, 2011; Barcelona, Spain. Corvalis, OR: AUAI Press; 2011. 717–726.
- Verma T, Pearl J. *Proc 4th Workshop Uncertain Artif Intell.* Minneapolis, MN/Mountain View, CA: AUAI Press; 1988. Causal networks: semantics and expressiveness; 352–59.
- Weinberg CR. Towards a clearer definition of confounding. *Am J Epidemiol.* 1993; 137:1–8. [PubMed: 8434568]
- Winship C, Korenman S. Does staying in school make you smarter? The effect of education on IQ in *The Bell Curve*. In: Devlin B, Fienberg SE, Resnick DP, Roeder K, editors *Intelligence, Genes, and Success: Scientists Respond to The Bell Curve*. New York: Springer; 1997. 215–34.
- Winship C, Mare RD. Models for sample selection bias. *Annu Rev Sociol.* 1992; 18:327–50.
- Wodtke G, Harding D, Elwert F. Neighborhood effects in temporal perspective: the impact of long-term exposure to concentrated disadvantage on high school graduation. *Am Sociol Rev.* 2011; 76:713–36. [PubMed: 22879678]
- Wooldridge J. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press; 2002.
- Wooldridge J. Violating ignorability of treatment by controlling for too many factors. *Econ Theory.* 2005; 21:1026–28.
- Wright S. The method of path coefficients. *Ann Math Stat.* 1934; 5(3):161–215.

Figure 1.
A directed acyclic graph (DAG).

Figure 2.

(a) *A* and *B* are associated by causation. The marginal association between *A* and *B* identifies the causal effect of *A* on *B*. (b) *A* and *B* are conditionally independent given *C*. The conditional association between *A* and *B* given *C* does not identify the causal effect of *A* on *B* (overcontrol bias).

Figure 3.

(a) *A* and *B* are associated by common cause. The marginal association does not identify the causal effect of *A* on *B* (confounding bias). (b) *A* and *B* are conditionally independent given *C*. The conditional association does identify the causal effect of *A* on *B* (which is zero in this model).

Figure 4.

(a) *A* and *B* are marginally independent. The marginal association identifies the causal effect of *A* on *B* (which is zero in this model). (b) *A* and *B* are associated conditionally on their common outcome, *C* (collider). The conditional association between *A* and *B* given *C* does not identify the causal effect of *A* on *B* (endogenous selection bias). (c) Conditioning on a descendant, *D*, of a collider, *C*, similarly induces an association between the causes of the collider. The conditional association between *A* and *B* given *D* does not identify the causal effect of *A* on *B*.

Figure 5. Endogenous selection bias due to outcome truncation, with E as education (treatment), I as income (outcome, truncated at 1.5 times poverty threshold), and U as error term on income.

Figure 6.

Endogenous selection bias due to listwise deletion [I , father's income (treatment); P , child support payments (outcome); R , survey response]. Conditioning on the post-outcome variable response behavior R (listwise deletion of missing data) induces a noncausal association between father's income, I , and his child support payments, P .

Figure 7. Endogenous selection bias due to sample selection [B , topping the Billboard charts (treatment); R , inclusion in the *Rolling Stone 500* (outcome); S , sample selection].

Figure 8.

Endogenous selection bias due to sample selection [M , motherhood (treatment); W_R , unobserved reservation wage; W_O , offer wage (outcome); E , employment; ε , error term on offer wage]. (a) Null model without effect of motherhood on offer wages. (b) Model with effect of motherhood on offer wages.

Figure 9.

Attrition and dependent/informative censoring in panel studies [P , poverty (treatment); D , divorce (outcome); C , censoring/attrition; U , unmeasured marital distress]. (a) Censoring is random with respect to poverty and divorce. (b) Censoring is affected by poverty. (c) Censoring is affected by treatment and shares a common cause with the outcome. Only in panel *c* does attrition lead to endogenous selection bias.

Figure 10.

Proxy control [S , schooling (treatment); W , wages (outcome); U , unmeasured ability; Q , test scores]. (a) U confounds the effect of S on W . (b) Q is a valid proxy for U ; conditioning on Q reduces, but does not eliminate, confounding bias. (c) Q is affected by S and U ; conditioning on Q induces endogenous selection bias. (d) Q is affected by S and affects W ; conditioning on Q induces overcontrol bias.

Figure 11.

Endogenous selection bias in mediation analysis [T , class size (randomized treatment); M , student achievement; Y , high school graduation (outcome); U , unobserved factors such as student motivation]. (a) M mediates the indirect effect of T on Y . (b) M is not a mediator.

Figure 12.

Endogenous selection bias due to latent homophily in social network analysis (i, j , index for a dyad of individuals; Y , civic engagement; U , altruism; F , friendship tie).

Figure 13.

Confounding cannot be reduced to a purely associational criterion. If U_1 and U_2 are unobserved, the DAGs in panels *a*, *b*, and *c* are observationally indistinguishable. (*a*) X is a common cause (confounder) of the treatment, T , and the outcome, Y ; conditioning on X removes confounding bias and identifies the causal effect of T on Y . (*b*) X is a pre-treatment collider on a noncausal path linking treatment and outcome; conditioning on X induces endogenous selection bias. (*c*) X is both a confounder and a collider; neither conditioning nor not conditioning on X identifies the causal effect of T on Y .

Table 1

Simplified representation of the sampling scheme of a study of the effect of topping the Billboard charts (B) on a record's inclusion in the *Rolling Stone 500* (R)

	$R = 0$	$R = 1$
$B = 0$	a	b
$B = 1$	c	d