CHAPTER 12

Causal Diagrams

M. Maria Glymour and Sander Greenland

INTRODUCTION

Diagrams of causal pathways have tong been used to visually summarize hypothetical relations among variables of interest. Modern causal diagrams, or causal graphs, were more recently devel oped from a mcrgerof graphical probability theory with path diagrams. The resulting theory provides a powerful yet intuitive device for deducing the statistical associations implied by causal relations. Conversely, given a set of observed statistical relations. a researcher armed with causal graph theory can systematically characterize all causal structures compatible with the observations. The theory also provides a visual representation of key concepts in the more general theory or longitudinal causality or Robins (1997); see Chapter 21 for further discussion and references on the tatter topic.

The graphical rules linking causal relations to statistical associations are grounded in mathe. matics. Hence, one way to think of causal diagrams is that they allow nonmathematicians to draw logically sound conclusions about certain types of statistical relations. Learning the rules for reading statistical associations from causal diagrams may take a little time and practice. Once these rules are mastered, though, they facilitate many tasks, such as understanding confounding and selection

Il •

bias, choosing covariates for adjustment and for regression analyses, understanding analyses of direct effects and instrumental-variable analyses, and assessing "natural experiments." In particular, diagrams help researchers recognize and avoid common mistakes in causal analysis.

This chapter begins with the basic definitions and assumptions used in causal graph theory. It then describes constnction of causal diagrams and the graphical separation rules linking the causal assumptions encoded in a diagram to the statistical relations implied by the diagram. The chapter concludes by presenting some examples of applications. Some readers may prefer to begin with the examples and refer back to the definitions and rules for causal diagrams as needed. The section on Graphical Models, however, is essential to understanding the examples. Full technical details of causal diagrams and their relation to causal inference can be found in Pearl (2000) and Spines et al. (2001), while Greenland and Pearl (2008) provide a short technical review. Less technical articles geared toward health scientists include Greenland et al. (1999a), Robins (2001), Greenland and Brumback (2002), Hernán et al. (2002), Jewell (2004), and Glymour (2006b).

PRELIMINARIES FOR CAUSAL GRAPHS

Consider two variables X and Y for which we wish to represent a causal connection from X to Y, often phrased as "X causes Y" or "X affects Y." Causal diagrams may be constructed with almost any definition of cause and effect in mind. Nonetheless, as emphasized in Chapter 4, it is crucial to distinguish causation from mere association. For this purpose we use the potential-outcome (counterfactual) concept of causation. We say that X affects Y in a population of units (which may be people, families, neighborhoods, etc.) if and only if there is at least one unit for which changing (intervening on) X will change Y (Chapter 4).

STATISTICAL INDEPENDENCE

Association of X and Y corresponds to statistical dependence of Y and X, whereby the distribution of Y differs across population strata defined by levels of X, When the distribution of Y does not differ across strata of X, we say that X and Y are statistically independent, or unassociated. If X and Y are unassociated (independent), knowing the value of X gives us no information about the value of Y. Association refers to differences in Y between units with different X values. Such between-unit differences do not necessarily imply that changing the value of X for any single unit will result in a change in Y (which is causation),

It is helpful to rephrase the above ideas more formally. Let Pr(Y = y) be the expected proportion of people in the population who have y for the value of Y; this expected proportion is more often called the probability that Y = y. If we examine the proportion who have Y = y within levels or strata of a second variable X. we say that we are examining the probability of Y given or conditional on X. We use a vertical line "l" to denote "given" or "conditional on." For example, Pr(Y = ylX = x) denotes the proportion with Y = y in the subpopulation with X = x. Independence of X and Y then corresponds to saying that for any pair of values x and y for X and Y,

$$\Pr(Y = y|X = x) = \Pr(Y = y) \qquad (12\text{-}11$$

which means that the distribution of Y values does not differ across different subpopulations defined by the X values. In other words, the equation says that the distribution of Y given (or conditional on) a particular value of X always equals the total population (marginal or unconditional) distribution of Y. As stated earlier,

if X and Y are independent, knowing the value of X and nothing more about a unit provides no information about the Y value of the unit.

Equation 12—1 involves no variable other than X and Y, and is the definition of marginal independence of X and Y. When we examine the relations between two variables within levels of a third variable—for example, the relation between income and mortality within levels of education—we say that we are examining the conditional relation. We examine conditional relationships in many contexts in epidemiology. We may intentionally condition on a variable(s) through features of study design such as restriction or matching, or analytic decisions, such as stratification or regression modeling. Conditioning may arise inadvertently as well, for example due to refusal to participate or

loss to follow-up, These events essentially force conditioning on variables that determine participation and ascertainment. Informally, it is sometimes said that conditioning on a variable is "holding the variable constant," but this phrase is misleading because it suggests we are actively intervening on the value of the variable, when all we are doing is separating the data into groups based on observed values of the variable and estimating the effects within these groups (and then, in some cases, averaging these estimates over the groups, see Chapter 15).

To say that X and Y are independent given Z means that for any values x, y, : for X, Y, and Z,

$$\Pr(Y = y|X — x, z = \Pr(Y = y|Z \quad [12\text{-}2]$$

which says that, within any stratum of Z, the distribution of Y does not vary with X. In other words, within any stratum defined in terms of Z alone, we should see no association between X and Y. If X and Y are independent given Z, then once one knows the Z value of a unit, finding out the value of X provides no further information about the value of Y.

CAUSATION AND ASSOCIATION

As explained in Chapter 4, causation and association are qualitatively different concepts. Causal relations are directed; associations are undirected (symmetric). Sample associations are directly observable, but causation is not. Nonetheless, our intuition tells us that associations are the result of causal forces. Most obviously, if X causes Y, this will generally result in an association between X and Y. The catch, ofcourse. is that even if we observe X and Y without error, many other forces (such as confounding and selection) may also affect the distribution of Y and thus induce an association between X and Y that is not due to X causing Y. Furthermore, unlike causation, association is symmetric in time (nondirectional), e.g., an association of X and Y could reflect Y causing X rather than X causing Y.

A study of causation must describe plausible explanations for observed associations in terms of causal structures, assess the logical and statistical compatibility of these suuctures with the observations, and (in some cases) develop probabilities for those structures. Causal graphs provide schematic diagrams of causal strictures, and the independencies predicted by a graph provide a means to assess the compatibility of each causal structure with the observations.

More specifically, when we see an association of X and Y, we will seek sound explanations for this observation, For example, logically, if X always precedes Y, we know that Y cannot be causing X. Given that X precedes Y, obvious explanations for the association are that X causes Y, that X and Y share a common cause (confounding), or some combination of the two (which can also lead to no association even though X affects Y). Collider bias is a third type of explanation that seems much less intuitive but is easily illustrated with graphs. We will first discuss focus on collider bias because it arises frequently in epidemiology.

COLLIDER BIAS

As described in Chapter 9, a potentially large source of bias in assessing the effect of X on Y arises when selection into the population under study or into the study sample itself is affected by both X and Y. Such selection is a source of bias even if X and Y are independent before selection. This phenomenon was first described by Joseph Berkson in 1938 (published in Berkson [1946]). Berksonian bias is an example of the more general phenomenon called collider bias, in which the association of two variables X and Y changes upon conditioning on a third variable Z if Z is affected by both X and Y. The effects of X and Y are said to "collide" somewhere along the way to producing Z.

As an example, suppose that X and Y are marginally independent and Z = Y — X, so Z is completely determined by X and Y. Then X and Y will exhibit perfect dependence given Z: If Z = z, then Y = X + z. As a more concrete example, body mass index (BMI) is defined as (weight in kg)/(height in meters)^2 and so is strongly affected by both height and weight. Height and weight are associated in any natural population, but not perfectly: We could not exactly tell a person's weight from his or her height. Suppose, however, we learn that the person has BMI = 25 kg/m^2 ;

then, upon being told (say) that the person is 2 m tall, we can compute his weight exactly. as BMI(height^2 j 25(4) = 100 kg,

Collider bias occurs even when the causal dependency of the collider Z on X and Y is not perfect, and when there are several intermediates between X and the collider or between Y and the collider. It can also be induced when X and Z (or Y and Z) are associated due to a common cause rather than because X influences Z.

Collider bias can result from sample selection, stratification, or covariate adjustment if X and Y affect selection or the stratifying covariates, It can be just as severe as confounding, as shown in the classic example in which X, Y, and Z were exogenous estrogen use, endometrial cancer, and uterine bleeding (Chapter 9). As discussed later, it can also can induce confounding.

SUMMARY

Four distinct causal structures can contribute to an association between X and Y: (a) X may cause Y; (b) Y may cause X; (c) X and Y may share a common cause that wc have failed to condition on (confounding); or (d) we have conditioned or selected on a variable affected by X and Y, factors influenced by such a variable, or a variable that shares causes with X and Y (collider bias). Of course, the obsened association may also have been affected by purely random events. As described in Part Ill of this book, conventional statistics focus on accounting for the resulting random variation. The remainder of this chapter focuses on the representation of causal structures via graphical models, and on the insights that these representations provide. Throughout, we focus on the causal structures underlying our observations, ignoring random influences.

GRAPHICAL MODELS

TERMINOLOGY

Causal diagrams visually encode an investigator's assumptions about causal relations among the exposure. outcomes, and covariates. We say that a variable X affects a variable Y directly (relative to the other variables in the diagram) if there is an arrow from X to Y. We say that X affects Y indirectly if there is a head-to-tail sequence of arrows (or "one-way street") from X to Y; such a sequence is called a directed path or causal path. Any variable along a causal path from X to Y is called an intermediate variable between X and Y. X may affect Y both directly and indirectly. In Figure 12—1, X affects Y directly and Z indirectly. The absence of a directed path between two variables represents the assumption that neither affects the other; in Figure 12—1, U and X do not affect each other.

Children of a variable X are variables that are affected directly by X (have an arrow pointing to them from X): conversely. parents of X are variables that directly affect X (have an arrow pointing from them to X). More generally, the desccndants of a variable X are variables affected. either directly or indirectly. by X; conversely, the ancestors of X are all the variables that affect X directly or indirectly. In Figure 12—1, Y has parents U and X, and a child Z; X has one child (Y) and two descendants (Y and Z); and Z has a parent Y and three ancestors, Y, U. and X.

It is not necessary to include all causes of variables in the diagram. If two or more variables in a graph share a cause, however, then this cause must also be shown in the graph as an ancestor of those variables, or else the graph is not considered a causal graph. A variable with no parents in a causal graph is said to be erogenous in the graph: otherwise it is endogenous, Thus, all

FIGURE 12-1 A Causal d;agram With no confounding

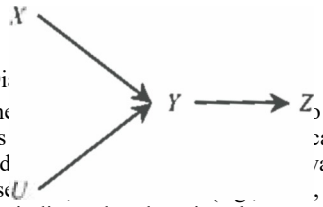

Chapter 12 • Causal Di
exogenous variables in the
unknown common causes
represented as unspecified
slight modification of these
cause with other variables in the grap
casual graph must show them: they ma
variables they are thought to influence.
, 2000) use a two-headed arrow betwee
variables as a shorthand to indicate that there is at least one unknown exogenous common cause c
two variables (e.g., X Z means that there is at least one unknown exogenous variable U such that
—+ Z). We assume in the remainder of this chapter that unknown common causes are repres
explicitly in causal diagrams, so there is no need for two-headed arrows.

All the graphs we will consider are acyclic, which means that they contain no feedback loops
means that no variable is an ancestor or descendant of itself, so if X causes Y, Y cannot also cause X a
same moment, If a prior value of Y affects X, and then X affects a subsequent value of Y, these must eac
shown as separate variables (e.g., Yo Xl —+ h) (for discussions of extensions to causal structures including feedback, see Spines [19951, Pearl and Dechter 11996), and Lauritzen and Richardson 120021). In most causal graphs the only connectors between variables are one-headed arrows (—+), although some graphs use an

undirected dashed line (to indicate associations induced by collider bias. Connectors, whether arrows or dashed lines, are also known as edges, and variables are often called nodes or vertices of the graph, Two variables joined by a connector are said to be adjacent or neighbors. If the only connectors in the graph are one-headed arrows, the graph is called directed. A directed acyclic graph or DAG is thus a graph with onty arrows between variables and with no feedback loops. The remainder of our discussion applies to DAGs and graphs that result from conditioning on variables in DAGs.

A path between X and Y is any noncrossing and nonrepeating sequence traced out along connec tors (also called edges) starting with X and ending with Y, regardless ofthc direction ofarrowhcads. A variable along the path from X to Y is said to intercept the path. Directed paths are the special case in which all thc connectors in the path flow head to tail. Any other path is an undirected path. In Figure 12—1 , U -+ Y 4-- X is an undirected path from U to X, and Y intcrccpts the patlb

When tracing out a path, a variable on the path where two arrowheads meet is called a collider on that path. In Figure 12—1, Y is a collider on the path U —Y Y +- X from U to X. Thus, a collider on a path is a direct effect (child) of both the variable just before it and the variable just after it on the path. A directed path cannot contain a collider. If a variable on a path has neighbors on both sides but is not a collider, then the variable must be either an intermediate (X —+ Y —¥ Z or X Y +- Z) or a cause (X 4— Y -+ Z) of its immediate neighbors on the path.

Being a collider is specific to a path. In the same DAG, a variable may be a collidcr on one path but an intermcdiale on another path; e.g., in Figure 12-1, Y is an intermediate rather than a collider on the path X -+ Y -+ Z. Nonetheless, a variable with two or more parents (direct causes) is called a collider in the graph, to indicate that it is a collider on at least one path. As we will see, paths Wilh colliders can turn out to be sources of confounding and selection bias.

RULES LINKING ABSENCE OF OPEN PATHS TO STATISTICAL INDEPENDENCIES

Given a causal diagram, we can apply the d-separation criteria (or directed-graph separation rules) to deduce indepcndencies implied by the diagram. We first focus on tiles for determining whether two variables are d-separated unconditionally, and then examine how conditioning on variables may d.separatc or d•connect other variables in the graph. Wc emphasize that the deduced relations apply only "in expectation," meaning that they apply to the expected data distribution if the causal structure represented by the graph is correct. They do not describe the associations that may arise as a result of purely random events, such as those produced by randomization or random sampling.

Unconditional d-Separation

A path is said to be open or unblocked or active unconditionally if there is no collider on thc path. Otherwise, if there is a collidcr on the path, it is said to be closed or blocked or inactive, and we say that the coltider blocks the path. By definition a directed path has no collider, so every directed path is open, although not every open path is directed. Two variables X and Y are said to be d.separatcd if there is no open path between them; otherwise they are d-connected. In Figure 12—2, the only  path from X to Y is open at Zl and but closed at W, and hence it is closed overall; thus X and Y

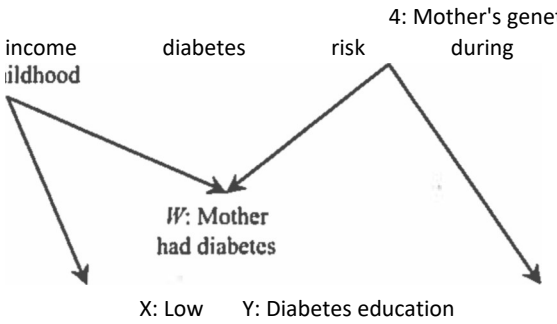

FIGURE 12—2 • ADAG under which traditional confounder-identification rules fail (an ᵉM diagram"),

are d-separated. When using these terms we will usually drop the "d." prefix and just say that they are separated or connected as appropriate.

If X and Y are separated in a causal graph, then the causal assumptions encoded by the graph imply that X and Y will be unassociated. Thus, if every path from X to Y is closed, the graph predicts that X and Y will be marginally independent; i.e., for any values x and y of X and Y, Pr(Y = ylX = x) = Pr(Y = y). More generally and informally we can say this: In a causal graph, the only sources of marginal association between variables are the open paths between them. Consider Table 12—1, which lists the causal assumptions represented by the diagram of Figure 12—1 , and the associations implied by those causal assumptions. For example, the causal diagram implies that U and X are marginally independent because the only path between them passes through a collider, Y. This idea is formalized later when we define compatibility.

Conditional d-Separation

We also need the concept of graphical conditioning. Consider first conditioning on a noncollider Z on a path. Because it is a noncollider, Z must either be an intermediate between its neighbors on the path (X Z -+ Y or X +- Z 4— Y) or a cause of its neighbors (X Z -+ Y). In these cases the path is open at Z, but conditioning on Z closes the path and removes Z as a source of association between X and Y. These phenomena reflect the first criterion for blocking paths by conditioning on covariates:

- Conditioning on a noncollider Z on a path blocks the path at Z.

In contrast, conditioning on a collider requires reverse reasoning. If two variables X and Y are marginally independent, we expect them to become associated upon conditioning (stratifying) on a shared effect W. In particular, suppose we are tracing a path from X to Y and reach a segment on the path with a collider, X -+ W Y. The path is blocked at W, so no association between X and Y passes through W. Nonetheless, conditioning on W or any descendant of W opens the path at W. In other words, we expect conditioning on W or any descendant to create an X—Y association via W. We thus come to the second criterion for blocking paths by conditioning on covariates:

- Conditioning on a collider W on a path, or any descendant of W, or any combination of W or its descendants, opens the path at W.

Combining these criteria, we see that conditioning on a variable reverses its status on a path: Conditioning closes noncolliders (which are open unconditionally) but opens colliders (which are closed unconditionally).

We say that a set of variables S blocks a path from X to Y if, after conditioning on S, the path is closed (regardless of whether it was closed or open to begin with). Conversely, we say that a set of variables S unblocks a path if, after conditioning on S, the path is open (regardless of whether it was closed or open to begin with), The criteria for a set of variables to block or unblock a path are summarized in Table 12—2.

TABLE

Assumptions Represented in the Directed Acyclic Graph in Figure 12-1, and Statistical Implications of These Assumptions

	Marginal Associations	Conditional Associations

Causal Assumptions Represented in Figure 12-1	Independencies Implied by Figure 12-1	Expected under Figure 12-1 (Assuming Faithfulness)	Expected under Figure 12-1 (Assuming Faithfulness)
• X and U are each direct causes of Y (direct With respect to other variables In the diagram). • Y isa directcauseofZ. • X is not a direct cause of Z but X is an indirect cause of Z Via Y. • X is not a cause of U and U is not a cause of X. • U is not a direct cause of Z, butU is an indirect cause of Z Via Y. • No two vartables in the diagram (X, U, Y, or Z) share a prior cause not shown jn the diagram, e g t no variable causes both X and Y. or both X and U.	• X and l.] are independent (the only path between them es blocked by the collider Y}. • X and Z are independent conditiona\ on Y (conditmng on Y blocks the path between X and Z h • U and Z are independent condit onal on Y.	• X and Y are associated. • U and Y are associated. • Y and Z are • X and Z are assoctated, • U and Z are assoctated.	• X andU are associated conditonal on Y {conduntng on a coilider unblocks the path). • X and V are associated conditional on Z (Z IS a descendant of the co!llder Y).

If S blocks every path from X to Y, we say that X and Y are d.separatcd by S, or that S separates X and Y. This definition of d.separation includes situations in which there was no open path before conditioning on S. For example, a set S may bc sufficient to separate X and Y even if S includes no variables: if there is no open path between X and Y to begin with, the empty set separates them.

d•Separation and Statistical Independence

We have now specified the d-separation criteria and explained how to apply them to determine whether two variables in a graph are d-separated or d-connected, either marginally or conditionally. These concepts provide a link between the causal structure depicted in a DAG and the statistical associations we expect in data generated from that causal structure, The following two rules specify Jhe relation between d.separation and statistical independence; these rules underlie the applications we wili present.

Rule 1 (compatibility). Suppose that two variables X and Y in a causal graph are separated by a set of variables S, Then if the graph is correct, X and Y will be unassociated given S. In other

Criteria for Determining Whether a Path is Blocked or Unblocked Conditional on a Set of Variables S

The Path from X to Y is Blocked Conditional on S if Both:	The Path from X to Y is Unblocked Conditional on S if Either:

A noncollider Z on the path is in S (because the path Will be blocked by S at Z)	S contains no nonco'lider on the path (so condition:ng on S blocks no noncolllder)
OR	AND
There is a collider W on the path that is not in S and has no descendant in S (because W still blocks the path after condition;ng on S).	Every collider on the path is either in S or has a descendant In S (because conditioning on S opens every collider).

words, if S separates X from Y, we will have Pr(Y VIX x. S = S) = Pr(Y VIS S) for every possible value x, y. S of X, Y, S.

Rule 2 (weak faithfulness). Suppose that S does not separate X and Y, Then, if the graph is correct, X and Y may be associated given S. In other words, if X and Y are connected given S, then without further information we should not assume that X and Y are independent given S.

As an illustration, consider again Figure 12—1. U and X are unassociated. Because Y is a collider, however, we expect U and X to become associated after conditioning on Y or Z or both (that is, S unblocks the path whether S = {Y S = {Z). or S = {Y, Z}). In contrast, X and Z are marginally associated, but become independent after conditioning on Y or S = {U. Y}.

ASSUMPTIONS AND INTUITIONS UNDERLYING THE RULES

Although informal diagrams of causal paths go back at least to the 1920s, the mathematical theory of graphs (including DAGs) developed separately and did not at first involve causal inference. By the 1980s, however, graphs were being used to represent the structure ofjoint probability distributions. with d-separation being used to encode "stable" conditional independence relations (Pearl, 1988). One feature of this use of graphs is that a given distribution will have more than one graph that encodes these relations. In other words. graphical representations of probability distributions are not unique. For example, in probabilistic (associational) terms. A —+ B and B -+ A have the same implication, that A and B are dependent. By the 1990s, however. several research groups had adapted these probability graphs to causal inference by letting the arrows represent cause-effect relation.% as they had in path diagrams. Many graphical representations that are probabilistically equivalent are not causally equivalent. For example, if A precedes B temporally, then B A can be ruled out as a representation for the relation of A and B.

The compatibility and faithfulness rules define what wc mean when we say that a causal model for a set of variables is consistent with a probability model for the distribution of those variables. In practice, the rules are used to identify causal graphs consistent with the observed probability distributions of the graphed variables. and, conversely, to identify distributions that are consistent with a given causal graph. When the arrows in probability graphs represent causal processes, the compatibility rule above (rule I) is equivalent to the causal Markov assumption (CMA), which formalizes the idea that (apart from chance) all unconditional associations arise from ancestral causal relations. Causal explanations of an association between two variables invoke some combination of shared common causes, collider bias, and one of the variables affecting the other. These relations fortn the basis for Rule l.

Specifically, the CMA states that for any variable X, conditional upon its direct causes (parents). X is independent of all other variables that it docs not affect (its nondescendants). This condition asserts that if we can hold constant the direct causes of X, then X will be independent of any other variable that is not itself affected by X. Thus, assuming X precedes Y temporally. in a DAG without

conditioning there are only two sources of association between X and Y: Effects of X on Y (directed paths from X to Y), or common causes (shared ancestors) of X and Y, which introduce confounding. We will make use of this fact when we discuss control of bias.

The d-separation rule (Rule 1) and equivalent conditions such as the CMA codify common intuitions about how probabilistic relations (associations) arise from causal relations. We rely implicitly on these conditions in drawing causal inferences and predicting everyday events—ranging from assessments of whether a drug in a randomized trial was effective to predictions about whether flipping a switch on the wall will suffuse a room with light. In any sequence of events, holding constant both intermediate events and confounding events (common causes) witl interrupt the causal cascades that produce associations. In both our intuition and in causal graph theory, this act of "holding constant" renders the downstream events independent of the upstream events. Conditioning on a set that d-separates upstream from downstream events corresponds to this act. This correspondence is the rationale for deducing the conditional independencies (features of a probability distribution) implied by a given causal graph from the d-separation rule.

The intuition behind Rule 2 is this: If, after conditioning on S, there is an open path between two variables, then there must bc some causal relation linking the variables, and so they ought to be associated given S, apart

from certain exceptions or special cases. An example of an exception occurs when associations transmitted along different open paths perfectly cancel each other, resulting in no association overall. Other exceptions can also occur. Rule 2 says only that we should not count on such special cases to occur, so that, in general, when we see an open path between two variables, we expect them to be associated, or at least we are not surprised if they are associated.

Some authors go beyond Rule 2 and assume that an open path between two variables means that they must be associated. This stronger assumption is catledfaithfulncss or stability and says that if S does not d-separate X and Y, then X and Y will be associated given S. Faithfulness is thus the logical converse of compatibility (Rule l). Compatibility says that if two variables are d.separatcd, then they must be independent; faithfulness says that if two variables are independent, then thcy must be d-separated. When both compatibility and faithfulness hold, we have perfect compatibility, which says that X and Y are independent given S if and only if S d-separatcs X and Y; faithfulness adds the "only if" part. For any given pattern of associations, the assumption of perfect compatibility rules out a number of possible causal structures (Spines et al., 2001 Therefore, when it is credible, perfect compatibility can help identify causal structures underlying observed data.

Nonetheless, because there are real examples of near-cancellation (e.g., when confounding obscures a real effect in a study) and other exceptions, faithfulness is controversial as a routine assumption, as are algorithms for inferring causal structure from observational data; see Robins (1997, section I l), Korb and Wallace (1997), Freedman and Humphreys (1999), Glymour et al. (1999), Robins and Wasserman (1999), and Robins et al. (2003). Because of this controversy, we discuss only uses of graphical models that do not rely on the assumption of faithfulness. Instead, we use Rule 2, which weakens the faithfulness condition by saying that the presence of open paths alerts us to the possibility of association, and so we should allow for that possibility.

The rules and assumptions just discussed should be clearly distinguished from the content. specific causal assumptions encoded in a diagram, which relate to the substantive question at hand. These rules serve only to link the assumed causal structure (which is ideally based on sound and complete contextual information) to the associations that we observe. In this fashion. they allow testing of those assumptions and estimation of the effects implied by the graph.

GRAPHICAL REPRESENTATION OF BIAS AND ITS CONTROL

A major use ofcausal graphs is to identify sources of bias in studies and proposed analyses, including biases resulting from confounding, selection, or over-adjustment. Given a causal graph, we can use the definitions and rules we have provided to determine whether a set of measured variables S is sufficient to allow us to identify (validly estimate) the causal effect of X on Y.

Suppose that X precedes Y temporally and that the objective of a study is to estimate a measure Of the effect of X on Y. We wilt call an undirected open path between X and Y a biasing path for the effect because such paths do not represent effects of X on Y, yet can contribute to the association Of X and Y. The association of X and Y is unconditionally unbiased or marginally unbiased for the effect of X on Y if the only open paths from X to Y are the directed paths.