# Chapter 2. Causal diagrams to evaluate sources of bias



2.1 Introduction . . . . . 24
2.2 Causal DAGs to evaluate sources of bias . . . . . 24
2.3 Example: building a DAG for opium consumption and lung cancer . . . . . 39
2.4 DAGs and specific sources of bias . . . . . 41
2.5 DAGs and multiple sources of bias . . . . . 51
2.6 Signed DAGs . . . . . 53
2.7 Use of DAGs in evidence synthesis . . . . . 56
2.8 Summary . . . . . 57

CHAPTER 2.

# Causal diagrams to evaluate sources of bias

*Matthew Fox, Onyebuchi A. Arah, Sonja Swanson, and Vivian Viallon*

## 2.1 Introduction

A key concern in studies of causal effects is identifying factors that prevent an observed association between an exposure and a disease from being equal to the true underlying causal effect of that exposure on the disease in the target population of interest. Although all analyses are subject to some systematic and random error, causal graphs, including causal directed acyclic graphs, can be used to attempt to understand which sources of bias may exist in studies and when sources of bias may prevent the identification of causation, such as the carcinogenic effect of an exposure.

Systematic error, or bias, occurs whenever the estimates generated in the study differ from the true causal effect for reasons other than random error. A key feature of systematic error is that, unlike random error, as the sample size in which the putative causal effect of interest is being studied increases, the systematic error is not expected to decrease. Epidemiologists generally focus on three main sources of systematic error: confounding, information bias, and selection bias, all defined in the Preface. These sources of bias are demonstrated in this chapter using causal diagrams and are discussed more extensively in the subsequent three chapters.

## 2.2 Causal DAGs to evaluate sources of bias

### *2.2.1 Introduction*

Causal diagrams represent hypothesized relations between variables (Pearl and Mackenzie, 2018; Lipsky and Greenland, 2022). This section describes one type of causal graph, the directed acyclic graph (DAG), and discusses how to use DAGs to reason about bias, as well as ways in which they may be useful to *IARC Monographs* reviewers and to those evaluating research studies of causal effects. DAGs can also be used to identify a set of variables that is sufficient to control for confounding. Side Box 2.1 provides a brief history of causal diagrams; Side Box 2.2 gives their relation to the concepts of counterfactuals.

Why should those studying cause and effect learn about causal DAGs? And why specifically would an *IARC Monographs* Working Group reviewer want to learn to use DAGs? Because epidemiological studies are being used to inform public health policy decision-making, including specifically

**Side Box 2.1.** History of causal diagrams

Graphical methods have a long history in science; they can be traced back to Sewall Wright's path tracing approach (Wright, 1960) and to structural equation modelling, and were developed further by Glymour and Scheines (1986) and Pearl (2009). However, their use within epidemiology increased substantially after the publication of a seminal article by Greenland et al. (1999). They are related to but separate from counterfactuals (or counterfactual variables). A detailed explanation of counterfactuals is outside the scope of this chapter; however, a brief explanation of the relation between DAGs and counterfactuals is given in Side Box 2.2. (text continues on page 24)

**Side Box 2.2.** Relation between causal diagrams and counterfactuals

Causal DAGs are one formal language for causal inference, in which causal effects are defined in terms of counterfactual or potential outcomes. In brief, to understand the effect of a binary exposure $X$ on a binary outcome $Y$, $Y^x$ can be defined as the counterfactual outcome that would have occurred given exposure level $x$, and counterfactual contrasts of interest can be described as being about those counterfactuals (e.g. a causal risk ratio in a given population is $E[Y^{X=1}]/E[Y^{X=0}]$). Causal DAGs do not reference counterfactuals explicitly, because they encode the way in which data are realized (i.e. the data-generation process) rather than counterfactual worlds. The indirect link between causal DAGs and counterfactuals is that the absence of an arrow $X \rightarrow Y$ in a causal DAG encodes the sharp causal null that $Y^{X=1} = Y^{X=0} = Y$ for all individuals in the study (to put it simply, the exposure has no effect on the outcome for any individual investigated in the study). Pearl (2009) depicts the potential outcome as the outcome $Y$ resulting from a mutilated DAG in which the arrow pointing into $X$ is deleted and $X$ is set to a specific value $x$ depicting an intervention on $X$. Such mutilated or augmented DAGs are sometimes called post-intervention DAGs; they can be used to identify potential outcomes as the consequences of an intervention on an exposure $X$. Note that other types of causal diagram, including twin networks and single-world intervention graphs (SWIGs), have more explicit links to counterfactual theory. The focus here is on DAGs because of their ubiquity in practice, but it should be acknowledged that there are relative strengths and limitations to other formalizations of causal inference and causal graphs. (text continues on page 24)

in the context of *IARC Monographs* hazard identification, there is a need to communicate the findings of studies among researchers unambiguously, across the disciplines with which epidemiologists collaborate (e.g. toxicology and exposure science), and to stakeholders and decision-makers (Swanson, 2015). In practice, causal graphs facilitate communication between colleagues versed in causal graphs. Experience shows that disagreements within scientific teams over appropriate analyses often come down to the team members each assuming different causal structures that underlie the data in their minds. When these assumed structures are expressed as DAGs, they illuminate which questions are most important. For example, suppose the disagreement is over whether a particular covariate should or should not have been adjusted for in a study. When the competing DAGs are drawn, it may become obvious that the scientific consensus favours one graph over the other, thus ending the disagreement and clarifying which consensus evaluation (e.g. *sufficient*, *limited*, or *inadequate*, as described in Chapter 1) best describes the available evidence.

In cancer epidemiology, DAGs are used to summarize and formalize assumptions about the causal relations that may exist among variables relevant to the assessment of the carcinogenicity of the exposure under study. These DAGs represent our understanding of the data-generation process, meaning the set of variables, both measured and unmeasured, as well as the relations between them, that lead to the observed data we have to investigate for assessment of

a causal effect of an exposure on an outcome (Example 2.1a).

Causal diagrams are made of nodes (represented by the variables named in the diagram) and arrows. Each node represents one variable, and a single-headed arrow between two nodes represents an assumption of a possible causal effect between the corresponding two variables. A single-headed arrow is sometimes referred to as a directed edge, because the direction of the arrow is intended to indicate the direction of causation.

Technical details about the implications of the arrows are given in Side Box 2.3, but two details are noted here. First, arrows in DAGs can only be single-headed (i.e. directed). This means that there can be no feedback loops; therefore, bidirectional relations where two variables both seem to affect each other must be represented with time-dependent variables, which affect each other over time (examples are given in Section 2.2.4 and in Fig. 2.3 in Side Box 2.3, as well as in Section 3.2.4(a) and Example 7.6). Second, the graph must be acyclic, meaning that there is no place in the graph where one can start and trace a path following the direction of the arrows and get back to where one started. This is necessary because arrows encode time, given that for *A* to cause *B*, *A* must precede *B* (temporality). Thus, *A* cannot cause *B* in the future and have this, in turn, affect itself in the past. (If one thinks that both *A* and *B* can cause each other, this should be depicted in a DAG using several instances of *A* and *B*, indexed over time.) Satisfying the conditions of being directed (i.e. having only single-headed arrows) and being acyclic creates a DAG and allows for assessments of bias in published research and possible strategies to mitigate bias when designing and analysing studies (Example 2.1b).

As Example 2.1 illustrates, DAGs must be created by people using the best knowledge they have of the (causal) associations between the variables involved that lead to the observed data. Thus, DAGs do not by themselves indicate whether or not a variable is a confounder or a mediator; those creating the DAG must decide on what they believe to be the causal structure that created the data, and then use the rules of DAGs (see Section 2.2.4) to assess, under the assumptions encoded in the DAG, whether confounding or mediation (or some other of the structural relations described below) is present.

### 2.2.2 Paths

This section describes paths in DAGs and how they can be used to identify sources of bias. A path is defined as any sequence of consecutive arrows in the causal diagram, regardless of the directions of the arrows. In Fig. 2.1, examples of paths include red meat consumption → CRC, red meat consumption → BMI → CRC, and red meat consumption ← family history of CRC → CRC. Any path that always follows the direction of the arrows is called a directed path (e.g. red meat consumption → BMI → CRC), and any path that does not necessarily follow the direction of the arrows is called an undirected path (e.g. red meat consumption ← family history of CRC → CRC). Whereas the arrows represent causal relations between variables, the paths in the DAG can be used to identify whether we expect to see associations between any two variables; some of these associations may be causal, some of them may represent bias, and some of them will be a combination of the two.

**Example 2.1a.** Motivation for creating a DAG for red meat consumption

For illustration, suppose that a team is reviewing the literature on whether red meat consumption is a hazard for colorectal cancer (CRC) and that there is a debate about whether it is critical for studies to have adjusted for family history of CRC and body mass index (BMI) to be considered high-quality evidence as part of the review (for simplicity, assume here that these are the only critical factors). Furthermore, suppose that some team members think that adjustments for both are necessary, whereas others think that only family history of CRC should be adjusted for and that adjusting for BMI may induce bias. A causal diagram, such as the hypothetical DAG of Fig. 2.1, could be drawn to help guide the group. (text continues above)

**Side Box 2.3.** Assumptions about arrows: causality and temporality

Assumptions in causal diagrams lie in the arrows that are absent from the diagram, as well as those that are present. The presence of an arrow from node $C_1$ to node $C_3$ in a causal diagram encodes an assumption that variable $C_1$ could be a (direct) cause of variable $C_3$, while the absence of an arrow from node $C_3$ to node $C_1$ stipulates the absence of a (direct) causal effect of variable $C_3$ on variable $C_1$. Arrows in DAGs encode assumptions about the existence of possible causal effects but not about the strength of such effects or their functional forms. The DAG in Fig. 2.2 indicates that $Y$ could be influenced by $C_1$, $C_2$ (via $C_4$), $C_4$, and $X$, but not how much these variables may influence $Y$. In particular, the DAG does not reflect whether the effect of $X$ on $Y$ is assumed to depend on the levels of some combination of $C_1$, $C_2$, and $C_4$ (in such a situation, these would be effect modifiers).

**Fig. 2.2.** Illustrative example of a causal diagram for the study of a possible causal effect of exposure $X$ on cancer risk $Y$.

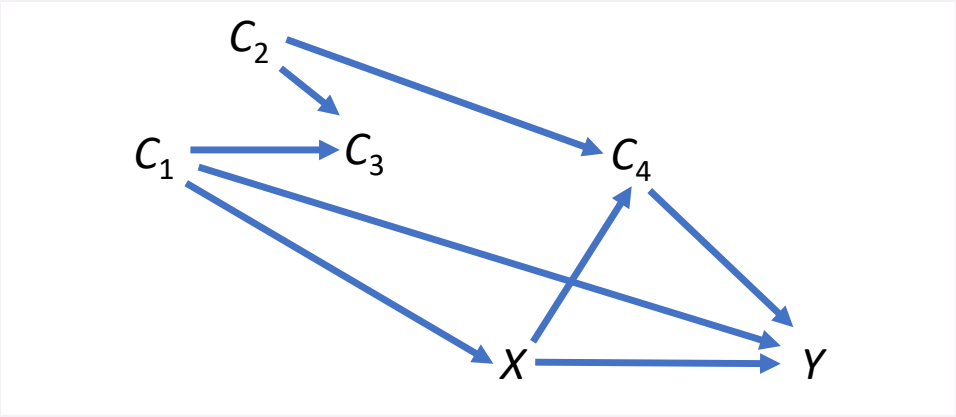


As noted by Hill (1965), temporality is a critical component of causality; for an agent to be causal, its presence must precede the development of the outcome. This implies that there can only be arrows $C_1 \rightarrow C_3$ and not also $C_3 \rightarrow C_1$ because causes ($C_1$) must precede their consequences ($C_3$). Cycles, or feedback loops, are usually prohibited in causal diagrams. If a directed path exists from $C_1$ to $C_3$, this implies that $C_1$ occurs before $C_3$; therefore, there cannot be another directed path from $C_3$ to $C_1$ because this would violate the temporality criterion. When no cycles are present in a causal diagram, it is said to be acyclic and is usually referred to as a DAG. To recap, the presence of an arrow from $C_1$ to $C_3$ in a DAG reflects an assumption that: (i) $C_1$ might have a direct causal effect on $C_3$; (ii) $C_3$ has no direct causal effect on $C_1$ (by the definition of a DAG, if there is an arrow from $C_1$ to $C_3$, there cannot be one from $C_3$ to $C_1$); and (iii) more generally, $C_3$ has no causal effect on $C_1$ (by the definition of a DAG, there cannot be any directed path from $C_3$ to $C_1$ if there is one from $C_1$ to $C_3$).

It is not always easy to determine the directionality of an arrow between two variables. Consider the example of obesity and physical activity. By increasing total energy expenditure, physical activity can help individuals to maintain their energy balance or even lose weight, so it can be inferred that lack of physical activity is probably a cause of obesity. However, excess weight also hampers physical activity, so that obesity can also be seen as a cause of lack of physical activity.

**Side Box 2.3.** Assumptions about arrows: causality and temporality (continued)

Such scenarios can be represented in DAGs by acknowledging the time-varying nature of exposures in the causal diagram (Fig. 2.3) and by drawing arrows (i) between (lack of) physical activity at any given time *t* and obesity at later times, *t* + 1, *t* + 2, … and (ii) between obesity at any given time *t* and (lack of) physical activity at later times, *t* + 1, *t* + 2, …. The expected association between obesity and (lack of) physical activity at any given time *t* can also be due to shared causes of these two variables; for example, health consciousness, although difficult to define and therefore rarely measured, may affect both (amount of) physical activity and obesity (e.g. through diet). In Fig. 2.3, it is assumed, for simplicity, that both diet and health consciousness are time-fixed, although time-varying versions could also be considered for these two variables. (text continues on page 26)

**Fig. 2.3.** Example of a longitudinal causal diagram, to illustrate a situation in which two variables affect each another but there is still no feedback loop.

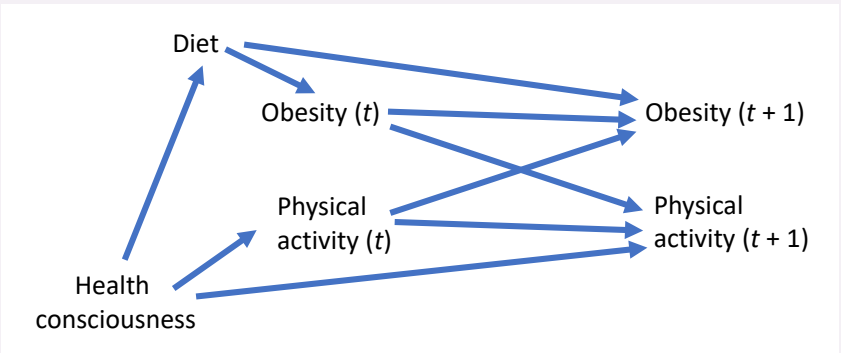


Paths in the DAG represent key information for assessing bias in a study. If a DAG is a true representation of the data-generation mechanism, some paths create associations (whether causal or non-causal) between variables, while others do not. Therefore, it is crucial to specify the paths that comprise a DAG, especially those linking the exposure (here, red meat consumption) and the outcome (here, CRC), to identify whether any observed association (here, between red meat consumption and CRC) could only result from causation or may include bias. Note that, although this may seem counterintuitive, it is possible to enumerate paths that do not follow the direction of the arrows, and it will be seen later that there are good reasons to do so.

There are three basic path structures in causal diagrams: chains, forks, and colliders. These are each discussed next, along with their implications with respect to associations between two variables.

Chains and forks induce an association between the nodes at the opposite ends of the path, whereas colliders do not. Conditioning on (e.g. adjusting for) nodes lying within a path can change the observed associations between variables, depending on the structure type. These are each described in Table 2.1. Although this may seem an abstract discussion, these three structures can be used to help solve disagreements about which variables should be adjusted for to obtain a valid estimate of the causal effect of an exposure *X* on an outcome *Y*, and about which variables should not be adjusted for or controlled (Example 2.2).

**Example 2.1b.** Motivation for creating a DAG for red meat consumption (continued)

In Fig. 2.1, the creator of the DAG is representing a view that a family history of CRC affects both red meat consumption (e.g. having a family history of CRC might cause a person to consume less red meat) and risk of CRC (because genetic causes of CRC can be inherited). A more formal explanation of this is given later, but the DAG shows that family history of CRC is what we would typically think of as a confounder and that it would need to be adjusted for to validly estimate the causal effect of red meat consumption on CRC.

In Fig. 2.1, the creator of the DAG is also representing a view that red meat consumption can affect one's BMI and a finding that having a high BMI can cause CRC. This would be an illustration of a mediating pathway; part of the way in which red meat consumption causes CRC is by increasing one's BMI. Accordingly, adjusting for BMI in a statistical model would have the effect of removing some of the effect of red meat consumption on CRC from effect estimates and would lead to an inaccurate assessment of the true total causal effect of red meat consumption on CRC (again, this will be explained in more detail later). Thus, if this DAG is a correct representation of the way in which the data were generated, BMI is not a confounder but a mediator, and therefore should not be adjusted for analytically. As this example illustrates, DAGs can help groups, including *IARC Monographs* Working Groups, to clarify their thinking about how to infer causality and to communicate with each other about which variables should be adjusted for to determine whether something is a cancer hazard. Note that this example graph is somewhat simplified and may not be the consensus graph; if others were to draw another graph and justify their differences, then perhaps a different conclusion about adjustment for BMI or family history of CRC would be reached. (text continues on page 26)

**Fig. 2.1.** Illustrative example of a causal diagram for a study of a possible causal effect of red meat consumption on risk of colorectal cancer (CRC). BMI, body mass index.

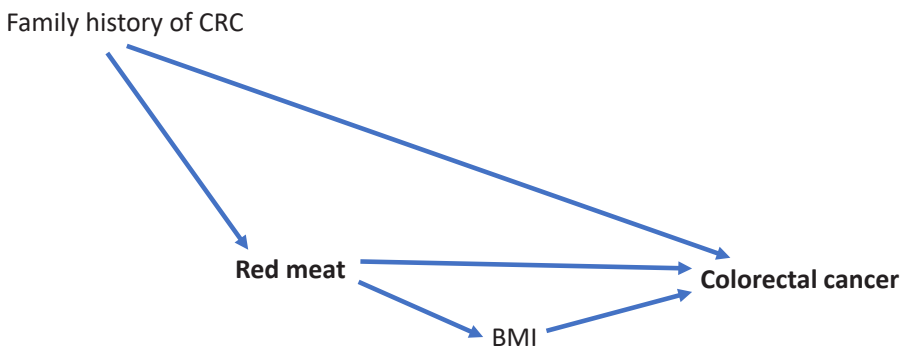

**Table 2.1.** Three basic structures in a directed acyclic graph and their implications for association and causation

| Name | Structure[a] | Example in Fig. 2.4 | Explanation | Implications using example |
|---|---|---|---|---|
| Chain | $X \rightarrow Y$<br>$X \rightarrow M \rightarrow Y$ | RM → BMI → CRC | A directed path in which all the arrows follow the same direction; the path from *X* to *Y* is open. | Creates a causal association between RM and CRC; BMI should not be adjusted for to estimate the (overall) causal effect of RM on CRC. |
| Fork | $X \leftarrow C \rightarrow Y$<br>$X \leftarrow Z \leftarrow C \rightarrow B \rightarrow Y$ | RM ← FH → CRC | An undirected path in which there is a directed path from one node to two others (*C* to *X* and *C* to *Y*); the path from *X* to *Y* is open. | Creates a non-causal association between RM and CRC; path must be blocked (e.g. by adjusting for FH) to estimate the causal effect of RM on CRC. |
| Collider | $X \rightarrow S \leftarrow Y$<br>$X \leftarrow A \rightarrow S \leftarrow B \rightarrow Y$ | RM → H ← CRC | An undirected path in which there are two directed paths from the outer nodes to a node in the centre (*X* to *S* and *Y* to *S*); the path from *X* to *Y* is blocked by collider *S*. | Creates no association between RM and CRC unless the collider is conditioned on (e.g. by adjusting for H). Controlling for H creates a non-causal association between RM and CRC (bias). |

BMI, body mass index; CRC, colorectal cancer; FH, family history of colorectal cancer; H, hospitalization; RM, red meat consumption.
[a] *M* is a mediator of *X* and *Y*, *S* is selection, *X* is an exposure, and *Y* is an outcome; other variable letter names have no specific meaning and are used to illustrate the causal structure.

**Example 2.2.** Chains, forks, and colliders used in DAGs

The three basic structures used in causal diagrams are shown in Fig. 2.4, using red meat consumption as an example. (text continues on page 28)

**Fig. 2.4.** Three basic structures in a causal diagram illustrating a study of a possible causal effect of red meat consumption on risk of colorectal cancer (CRC): (a) chains; (b) forks; and (c) colliders. BMI, body mass index. Blue arrows indicate the structure being described; grey arrows are intended to make it easier to view the structure being illustrated.

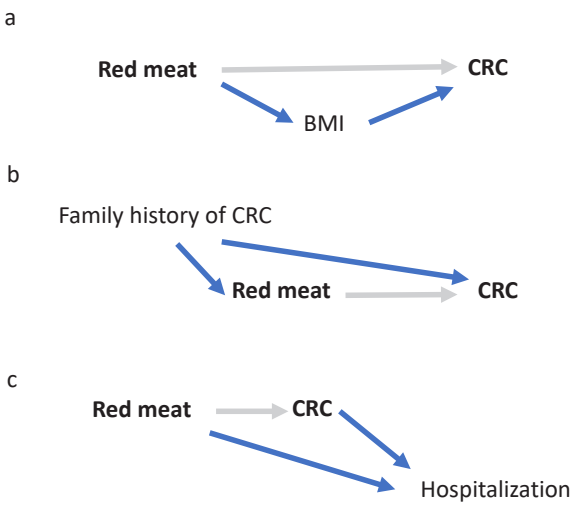

### *(a) Chains*

Directed paths are also called chains (Fig. 2.4a). In Fig. 2.1, the paths red meat consumption → CRC and red meat consumption → BMI → CRC generally imply an association between red meat consumption and CRC because in both cases red meat consumption may cause CRC, either directly (red meat consumption → CRC) or indirectly (red meat consumption → BMI → CRC). Chains represent causation; therefore, we are interested in the chains that follow pathways from the exposure to the outcome, whereas a chain from the outcome to the exposure would represent reverse causation (meaning a situation in which the variable that was assumed to be the independent variable was in fact the dependent variable, and vice versa). Because chains from the exposure to the outcome represent the causal effects whose effect sizes are to be estimated in causal epidemiological research, we do not want to disrupt these chains in a study design or data analysis.

One way to think about paths is as avenues for associations to flow along. Thus, if we are interested in the total effect of red meat consumption on CRC, conditioning on BMI (e.g. through restriction or covariate adjustment using methods such as stratification and regression) would create biased estimates in most situations in which the total effect is of interest, because this would block the causal path through BMI (i.e. stop the flow of the association from red meat consumption to CRC through BMI), hence eliminating part of the causal association. If the only way in which red meat consumption affected CRC was through changes in BMI, it would be expected that adjusting for BMI (or matching on it in the study design) would lead to an estimated null association between red meat consumption and CRC, when in fact there truly was a causal effect. Thus, if the hypothesized relations in the DAG are correct, reviewers would be wise to be concerned about a null result from a study of the effect of red meat consumption on CRC that adjusted for BMI, because it is possible that the reason the study showed a null result was not because there is no effect but rather because the authors inappropriately removed the effect by adjusting for BMI.

#### Key message

Although there may be circumstances in which we want to estimate the effect of an exposure that is not mediated through a specific pathway (in which case we might want to control for a variable on the causal pathway from exposure to outcome), we usually want to ensure that studies used for cancer hazard identification do not adjust for variables that lie on the causal chain from the exposure to the outcome.

There are two other basic structures in causal diagrams: forks (e.g. red meat consumption ← family history of CRC → CRC) and colliders (e.g. red meat consumption → hospitalization ← CRC). These correspond to simple forms of undirected paths.

### *(b) Forks*

In a fork (Fig. 2.4b), there is (in a path) a node that has two arrows, each pointing to one other node. A path that contains only chains or forks in which no variables in the path are controlled (e.g. adjusted for statistically), except for the first and last variables in the path, is said to be open or unblocked. Non-causal associations between two variables (i.e. dependency) flow through forks; when a fork exists in a DAG, this implies that there is an association between the two variables at the end of the fork, even though that association is not causal. In Fig. 2.1, the path red meat consumption ← family history of CRC → CRC does not, on its own, imply any causal effect of red meat consumption on CRC, but it still typically induces a spurious association between the variables red meat consumption and CRC because the two have a common cause: family history of CRC.

As described in the Preface, confounding is the entanglement of a third factor (a confounder) in the association between an exposure of interest and an outcome of interest.

#### Key message

Open forks give rise to confounding (see Section 2.4.1 and Chapter 3) and represent confounding in DAGs.

Note that forks that indicate confounding can be made up of two chains, one going from a single variable to the exposure and one going from that same variable to the outcome. These paths can comprise a single arrow or can travel across several variables to get to the exposure and the outcome, as long as the direction of the arrows continues to lead from the node to the exposure (or the outcome). Given that forks represent biasing pathways,

they need to be closed (or blocked) to remove the bias. These are the confounding paths that researchers and readers of the epidemiological literature need to be concerned about when designing studies, analysing data, or conducting literature reviews, because if these open paths are not closed, they can create bias and may lead to the conclusion that there is a hazard when there is not, or vice versa, or they can cause overestimation or underestimation of the magnitude of the effect of an exposure on an outcome.

Seen another way, forks describe shared causes (of the variables at the arrowheads) that lead to confounding (Example 2.3a).

An open path can be closed or blocked by controlling for any inner node of that path that is not a collider (described next; see Fig. 2.4c) through adjustment, stratification, matching, regression, and so on (Example 2.3b). In DAGs, conditioning through analytical control of a variable is represented by drawing a box around that variable.

Note that, in this explanation of forks, confounders have not been discussed, only biasing pathways. This is because, although it may seem that the variable that is at the apex of the fork is the confounder, the confounding pathways can be blocked by controlling any variable on the pathway. Thus, removing confounding (sometimes called de-confounding) is much more important than identifying which variable is the confounder. Nonetheless, when DAGs are used, a variable is often called a confounder if it can be used (e.g. adjusted for) to block a confounding pathway.

### *(c) Colliders*

A collider (Fig. 2.4c) is a node along a path with two arrows directly pointing to it along that path. Colliders do not, on their own, create a non-causal (biasing) pathway between the outer nodes; thus, a path that contains at least one collider that is not adjusted for is said to be blocked or closed (Example 2.4).

The bias induced by conditioning on a collider is not the most intuitive, but it can be understood by considering an idealized example, as given in Example 2.5.

A special but important case of collider stratification bias is when the collider is selection into the study. There are many reasons why people are enrolled (or choose to participate) in a study or drop out of a study. Because the study analysis can only be conducted among people who are enrolled in the study and for whom there is sufficient data, all studies are conditioned on selection. This means that if there is an effect to be estimated in a

**Example 2.3a.** Forks as depictions of shared causes in DAGs

The DAG in Fig. 2.1 indicates that red meat consumption and CRC have a shared cause, because one can trace a path following the arrows from family history of CRC to CRC and from family history of CRC to red meat consumption. Shared causes are typically thought of as confounding pathways that would need to be accounted for to find the causal effect of red meat consumption on CRC. (text continues above)

**Example 2.3b.** Conditioning or blocking of paths in DAGs

The association between red meat consumption and CRC, indicated as red meat consumption ← family history of CRC → CRC (the non-causal pathway that represents confounding), would be eliminated by conditioning on family history of CRC and could be partially removed by adjusting for any descendants of family history of CRC (i.e. variables with a directed path from family history of CRC to that variable), leaving only the associations indicated by the paths red meat consumption → BMI → CRC and red meat consumption → CRC, both of which are causal pathways. (text continues above)

**Example 2.4.** Depiction of colliders in DAGs

Suppose a group was reviewing a study in which data were collected on red meat consumption, hospitalization, and CRC and that red meat consumption increased the risk of being hospitalized (e.g. because of a heart attack), as did CRC. The path red meat consumption → hospitalization ← CRC does not create an association between red meat consumption and CRC, because this path is blocked by a collider. As long as the study design and analysis did not include conditioning on hospitalization (did not adjust for it, match on it, stratify on it, etc.), the results are likely to be valid (assuming that the DAG is correct and there are no other sources of bias). However, unlike forks, a blocked path can be unblocked by conditioning on one of the colliders (or any of its descendants, i.e. variables with an arrow towards that variable from the collider) (Berkson, 1946; Pearl, 2009). Although this is not necessarily intuitive, conditioning on colliders can induce spurious (non-causal) associations and result in collider stratification bias. Collider stratification bias is a bias that is created by conditioning on a collider, or an effect of (i.e. a descendant of) a collider. Thus, if the study adjusted for hospitalization, this would create bias in the association between red meat consumption and CRC and could make it appear that there was a hazard when there was not, or vice versa, or it could simply distort the magnitude of any real effect. Adjustment for hospitalization is elaborated on in Example 2.5, and colliders are described more intuitively in Side Box 2.4. (text continues on page 32)

population but the entire population, or a representative sample of that population, is not enrolled, it is only possible to estimate the effect in the selected sample; thus, the analysis is limited to a sample in which there are factors that lead to selection into the study. If both the exposure and the outcome are associated with selection into the study (either directly or indirectly through other forking paths), this can cause collider stratification bias. An example is provided in Section 2.4.3 (see also Chapter 5).

In conclusion, note again that for a diagram to be a DAG and therefore helpful for identifying and mitigating the impact of various sources of bias, it must be both directed and acyclic. A directed graph is one in which connections between variables must be drawn as single-headed arrows representing causality (an association cannot be implied between two variables using a dashed line without a specific cause, because any association must have a reason, as will be discussed later); furthermore, each causal path connecting more than two nodes sequentially in the DAG must contain arrows that point in the same direction. An acyclic graph is one for which there is no place in the diagram from which it is possible to trace a path following the direction of the arrows and get back to the starting point; in other words, no variable can cause itself in a DAG. Finally, for a DAG to be a causal DAG that can be used to identify sources of bias, the shared causes of any two variables in the DAG must also be represented. This means that a DAG that omits, for example, a common cause of any two variables is not a causal DAG, because the unknown causes of the two variables will not be mutually independent.

### *2.2.3 How to create a DAG*

Researchers and reviewers often need to describe the data-generation process used in a study to assess the effect of an agent on cancer (i.e. the forces in the universe that create relations between variables, whether or not they are ever collected in a study, along with any relations created between variables in the process of study design and analysis) to decide which variables would ideally be controlled for to determine the causal effect in the study. Note that DAGs can also include several versions of a variable measured at different time points, as shown in Fig. 2.3 in Side Box 2.3. Drawing a DAG can help researchers and reviewers select such a set of variables. The drawing of DAGs requires expert knowledge of subject matter and of the data-generation process; teams researching causal relations or review panels determining the quality of existing evidence to ascribe causation (e.g.

**Side Box 2.4.** Collider bias

Consider the example where two binary {0,1} variables $C_1$ and $C_2$ have a common effect, a third binary variable $C_3$, as shown in Fig. 2.5. Conditioning on this common consequence $C_3$ usually creates a spurious association between $C_1$ and $C_2$, referred to as collider bias. For illustration, consider the situation where $C_3$ equals 1 if only one of $C_1$ and $C_2$ is equal to 1 but 0 if both are equal to 1 or neither is equal to 1, as shown in Table 2.2. If $C_1$ and $C_2$ are independent in the general population, then having information about $C_1$ for a person in the general population does not give us any information about the value of $C_2$ for that person. However, among individuals with $C_3 = 1$, if $C_1 = 0$ then necessarily $C_2 = 1$, and if $C_1 = 1$ then necessarily $C_2 = 0$. In other words, among individuals with $C_3 = 1$, having information about $C_1$ does give us information about $C_2$, highlighting that $C_1$ and $C_2$ are not independent after conditioning on $C_3$. In contrast, among individuals with $C_3 = 0$, if $C_1 = 0$ then necessarily $C_2 = 0$, and if $C_1 = 1$ then necessarily $C_2 = 1$. Therefore, within levels of $C_3$, $C_1$ and $C_2$ are perfectly inversely correlated. Conditioning on $C_3$ creates a spurious inverse association between $C_1$ and $C_2$. This is illustrated in Table 2.3 for a group of 400 individuals for whom there is no association in the total population between $C_1$ and $C_2$ but there is a perfect inverse correlation within $C_3$.

**Fig. 2.5.** Example of bias created by conditioning on a collider, $C_3$.

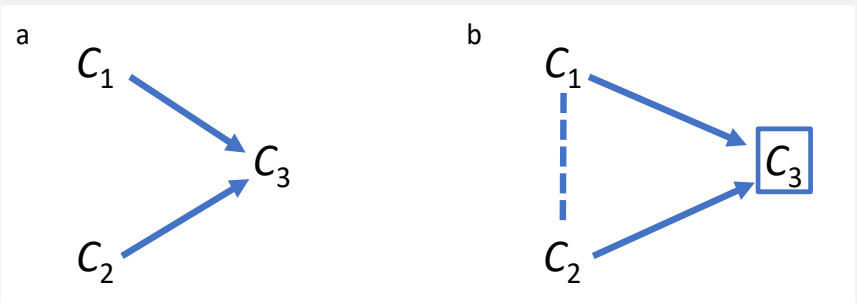


**Table 2.2.** Values of three variables, $C_1$, $C_2$, and $C_3$

| $C_1$ | $C_2$ | $C_3$ |
|---|---|---|
| 0 | 0 | 0 |
| 0 | 1 | 1 |
| 1 | 0 | 1 |
| 1 | 1 | 0 |

**Table 2.3.** Frequency of cross-tabulation of $C_1$ and $C_2$, both overall and stratified by $C_3$

| | Total | | | $C_3 = 1$ | | | $C_3 = 0$ | |
|---|---|---|---|---|---|---|---|---|
| | $C_1 = 1$ | $C_1 = 0$ | | $C_1 = 1$ | $C_1 = 0$ | | $C_1 = 1$ | $C_1 = 0$ |
| $C_2 = 1$ | 100 | 100 | $C_2 = 1$ | 0 | 100 | $C_2 = 1$ | 100 | 0 |
| $C_2 = 0$ | 100 | 100 | $C_2 = 0$ | 100 | 0 | $C_2 = 0$ | 0 | 100 |
| Total | 200 | 200 | Total | 100 | 100 | Total | 100 | 100 |
| % $C_2 = 1$ | 50% | 50% | % $C_2 = 1$ | 0% | 100% | % $C_2 = 1$ | 100% | 0% |

**Side Box 2.4.** Collider bias (continued)

Of course, collider bias is not restricted to binary variables. If smoking intensity ($C_1$) and alcohol intake ($C_2$) are both positively causally associated with the risk of a certain disease ($C_3$ = 1 if the individual develops the disease), then the association between smoking intensity and alcohol intake is typically less after conditioning on $C_3$, compared with that in the general population. This is because individuals with a low level of alcohol intake who develop the disease are more likely to have a higher smoking intensity than individuals in the general population. Similarly, individuals with a high level of alcohol intake who do not develop the disease are more likely not to smoke than individuals in the general population. (text continues on page 33)

**Example 2.5.** Example of a collider in a randomized controlled trial among downhill skiers

Suppose that we are interested in studying whether downhill skiing affects cancer risk, and in fact we even conducted a randomized trial in which participants were randomized to either never ski or ski frequently. In this randomized trial, we might expect to find no association between skiing and subsequent cancer risk, unconditionally. However, what if we restricted our analysis to only trial participants who went to the hospital sometime during the trial? Hospitalization might be a collider for a path skiing → hospitalization ← cancer, because both a cancer diagnosis and skiing accidents may lead to hospitalization (see the DAGs in Fig. 2.6). In an analysis in which only hospitalized participants were considered, skiing and cancer would probably be identified as inversely related: a person with a cancer diagnosis is less likely to be in the hospital for a skiing accident, and vice versa. Thus, we would not want to condition on hospitalization status in this trial, because it might make us wrongly conclude that skiing prevents cancer when, in this stylized example, it has no effect. (text continues on page 32)

**Fig. 2.6.** Example of a causal diagram depicting the relation between skiing (*S*) and cancer risk (*Y*) in a randomized trial: (a) full trial; (b) restriction to people who were hospitalized (*H*).

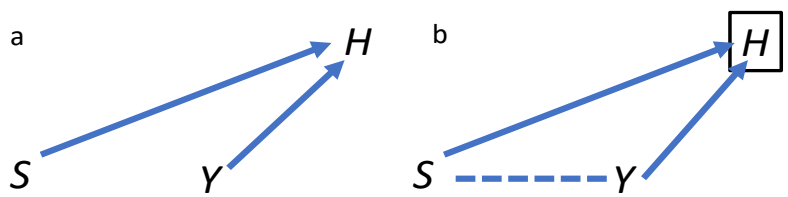

*IARC Monographs* Working Groups) typically possess such knowledge.

There are numerous approaches to creating a DAG. When creating a graph to support an analysis, it is critical to list all the variables that are considered essential in the data-generation process. For those reviewing the literature and therefore faced with the task of using the available studies (i.e. they are not conducting their own analyses or planning their own studies) to assess whether a hazard exists (but not necessarily how big the effect is), the approach could be simplified to focus on those factors that are likely to have the largest impact on creating bias. However, it is critical to note that this list should include all variables that might lead to a reasonable amount of bias, not only those measured in a study, because bias can exist even if the study did not account for it.

**Key message**

In identifying variables to be included in DAGs, priority could be given to those variables that might form part of an undirected forking path between the exposure and the outcome, because such undirected paths reflect potential sources of bias.

The next step is to link the variables with arrows, representing the possible causal relations between them, while remembering that the lack of an arrow between any two variables denotes a strong assumption about the absence of a causal relation between them. One approach to this would be to use one's best understanding of the causal relations to guide the first draft of a DAG. The DAG could then be presented to experts and stakeholders and revised based on feedback. Another approach would be to order all the nodes in time (with, say, left representing earlier time points and right representing later time points) and draw arrows from all variables that occur earlier in time to ones that occur later, only removing an arrow if there is a strong justification to do so based on expert knowledge that there is truly no causal effect of one variable on the other. Sometimes several competing DAGs must be considered when evidence- and knowledge-driven consensus remains elusive.

When the DAG describing the data-generation mechanism is complete, one can consider adding depictions of the study design and any bias that might have been created in the design process, focusing on selection bias and information bias (each described in the Preface and in detail later in this chapter). As noted in Section 2.2.2 and described further in Section 2.4.3, selection bias (collider stratification bias) can be introduced through the ways in which people are selected into or out of the study as well as into or out of analytical groups (through conditioning, matching, dropout, etc.). Selection can be represented as a node in the diagram (typically identified with the letter *S*), and the factors that are likely to cause participants to be in a study can be identified, whether they are factors determined by the design (e.g. selection of the study population in a case–control study, implementation of inclusion criteria) or factors that might determine the likelihood of participants self-selecting into a study (e.g. socioeconomic status [SES]).

Review panels assessing causation can use a DAG to determine whether the reviewers think that the analytical choices have removed all (or most) of the biases that existed. Furthermore, the DAG can be used to determine whether biases were created in the design or analysis that might prevent observation of a causal effect.

The process of drawing realistically complex DAGs can itself be complex, and some find it helpful to use software, such as DAGitty (https://dagitty.net/), which is freely available online, or the advanced Causal Fusion platform (https://www.causalfusion.net/), which is free but requires one to sign in. These tools can also provide an automated way to analyse a DAG for sets of variables, to control for confounding.

### 2.2.4 Rules of DAGs

#### (a) Causal paths

A review panel, such as an *IARC Monographs* Working Group, may wish to develop a DAG for an agent under evaluation and a type of cancer to help identify a reasonable set of variables to control for in the literature reviewed and another set that would ideally be ignored to give a valid result (Example 2.6a).

#### (b) Backdoor paths

In any study, but especially in observational studies where no randomization of the exposure occurred, it is necessary to consider that any observed association (e.g. between red meat consumption and CRC) may be a mix of any true causal effect and sources of bias (such as confounding by family history of CRC and diabetes). In DAGs, the most well known of these non-causal pathways are the open backdoor paths. These

**Example 2.6a.** A possible DAG for red meat consumption and colorectal cancer

The DAG in Fig. 2.7, which is a slightly more detailed version of the previous DAG for a study of the relation between red meat consumption and CRC, represents the DAG creator's understanding of the data-generation process. This DAG indicates that a family history of CRC is thought to affect red meat consumption (most probably by motivating one to consume less) and that a family history of CRC may also cause CRC. It also indicates that diabetes is thought to affect red meat consumption and that diabetes may cause CRC. Because there is no arrow from family history of CRC to diabetes or from diabetes to family history of CRC, the DAG also suggests that the two variables have no causal relation with each other and therefore would be expected to have no association, as revealed in the data. Note that there may also be an arrow from diabetes to hospitalization, but this is omitted for simplicity. (text continues on page 36)

**Fig. 2.7.** Directed acyclic graph for a study to assess the effect of red meat consumption on colorectal cancer (CRC).

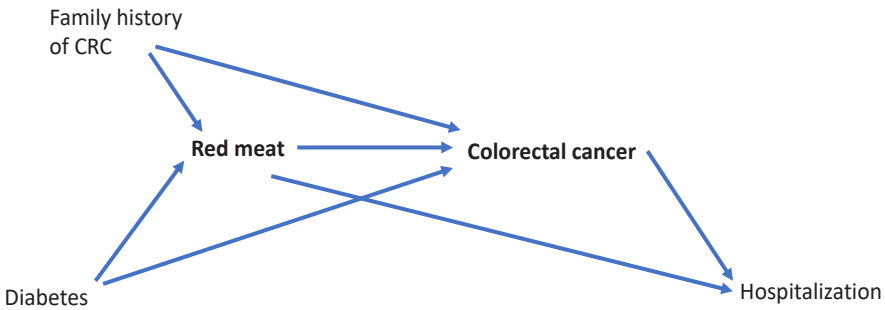


are some of the paths that reviewers of epidemiological research want to identify to determine whether they are blocked, meaning either that they contain a collider or that a variable on the path has been controlled analytically. This is necessary to ensure that a causal effect has been satisfactorily assessed in a study, especially in a study for which the reviewer believes that the arrows drawn in the DAG represent strong effects. Backdoor paths are undirected paths (meaning that it is not necessary to follow the arrow directions) that can be traced from the exposure (e.g. red meat consumption) to the outcome (e.g. CRC) by tracing a path that begins with an arrow pointing towards the exposure (hence the term *backdoor path*) and ends with an arrow pointing towards the outcome. From the head of the arrow pointing towards the exposure (e.g. red meat consumption), the path can be traced in any direction to get to the outcome (e.g. CRC).

Although some backdoor paths indicate bias, not all do. For there to be a bias, the backdoor path must be open or unblocked so that the biasing associations can flow from a variable to both the exposure and the outcome (see Example 2.6b and Side Box 2.5). If such an open backdoor path can be traced, there will be a non-causal association between the exposure and the disease that is mixed with any causal effect of the exposure on the disease and that must be accounted for through some adjustment method to identify the true causal effect of an exposure *X* on an outcome *Y*. If the backdoor path is of the first type (e.g. containing forks only), the path can be blocked through analytical control of any variable between the exposure and the outcome. As noted, if the backdoor path is of the second type (e.g. containing a collider), it is blocked naturally only if one does not condition on the collider or a descendant of the collider, through methods such as statistical adjustment or design approaches such as restriction and matching. Otherwise, the path is open and must be closed again through the

**Example 2.6b.** Backdoor paths in a DAG for red meat consumption and colorectal cancer

In Fig. 2.7, there is an open (non-causal) backdoor path from red meat consumption to CRC, i.e. red meat consumption ← family history of CRC → CRC, which is one of the forking paths described previously. No variable along this path is a collider (there is no point in the path where one can enter a variable through the head of an arrow and also exit the same variable through the head of an arrow); if we have not controlled for any variables analytically, this path will create confounding, as noted earlier. This is because red meat consumption and CRC have a shared cause: family history of CRC.

Although one could try to identify all the shared causes of red meat consumption and CRC, the backdoor approach is a systematic way of identifying all the confounding pathways. This DAG shows that there is another unblocked backdoor path, as listed in Table 2.4. As well as red meat consumption ← family history of CRC → CRC, there is red meat consumption ← diabetes → CRC. Note that the path red meat consumption → hospitalization ← CRC is not a backdoor path, because although it does start at red meat consumption and end at CRC, it does not begin by going through an arrow towards red meat consumption. This path is also not an open (unblocked) path, because it contains a collider, hospitalization; thus, it is not a biasing pathway, as long as hospitalization is not conditioned on in the design or analysis. Seen another way, there is no variable in the path that one can start at and trace a path following the arrows and get to both red meat consumption and CRC. Thus, this path does not show a shared cause of red meat consumption and CRC, and hence there is no confounding. (text continues on page 37)

**Table 2.4.** All backdoor paths from red meat consumption to colorectal cancer in Fig. 2.7

| Path | Backdoor? | Status | Reason for status | Path creates bias? |
|---|---|---|---|---|
| Red meat consumption ← family history of colorectal cancer → colorectal cancer | Yes | Open, unblocked | Fork with no collider, no variable on the path conditioned on | Yes, confounding |
| Red meat consumption ← diabetes → colorectal cancer | Yes | Open, unblocked | Fork with no collider, no variable on the path conditioned on | Yes, confounding |
| Red meat consumption → hospitalization ← colorectal cancer | No | Blocked | Path contains a collider | No, unless collider or its descendant is conditioned on |

**Side Box 2.5.** Open backdoor paths

An open or unblocked backdoor path is a backdoor path that either

- does not contain a collider and no variable along it has been conditioned on; if the path contains a collider, it is naturally blocked, and there would be no variable along this path from which one could trace a path following the arrows to get to the exposure and another path following the arrows to get to the outcome (such a path would consist only of forks); or
- contains a collider (or a descendant of a collider) that has been conditioned on and otherwise does not condition on any non-colliders; this is because adjusting for a collider or a descendant of a collider opens the path that would have otherwise been blocked. (text continues on page 37)

analytical control (i.e. conditioning) of any non-collider between the exposure and the outcome.

For the purposes of evaluating the literature to see whether there is an effect of an exposure on an outcome, as an *IARC Monographs* Working Group might do, the DAG should represent all variables that are likely to lead to meaningful bias, even if they were not measured in the study being reviewed. For now, assume that all the variables in the DAG were measured. Now that each of the unblocked backdoor paths has been identified, it is necessary to assess whether the set of variables that were adjusted for is sufficient to control all the confounding (i.e. all the unblocked backdoor paths were blocked). The bias from a backdoor path can be removed by conditioning (through analytical control or design approaches) on any variable along the path. Thus, it is then necessary to identify a set of variables that will close (block) all the open (unblocked) backdoor paths. In the example of red meat consumption and CRC, the study could condition on family history of CRC and diabetes through analytical control, to block all the open backdoor paths (Example 2.6c), leading to an unbiased result.

It can be seen in Fig. 2.9 that all the existing biasing pathways in the DAG (i.e. all the unblocked backdoor paths) have been successfully blocked; however, in adjusting for the collider on the pathway red meat consumption → hospitalization ← CRC, a new biasing pathway is opened up: red meat consumption --- CRC. The true effect (red meat consumption → CRC) will now be mixed with the biasing pathway (red meat consumption --- CRC), giving a biased result. The resulting bias can be large, moderate, or small, depending on the context, including the strength of the associations and the distribution of the variables in the DAG. This example demonstrates that adjustment for variables in a statistical model can sometimes create rather than remove bias.

### *(c) Importance of bias*

Before moving on to other examples, it is important to emphasize that DAGs can help to identify only whether a bias potentially exists, not its direction and magnitude. When using signed DAGs (see Section 2.6), it is sometimes possible to tell the direction of the bias; this is useful for identifying which biases can be ruled out as an explanation for an observed association (e.g. if the DAG identified a source of bias as operating downwards – a bias towards the null for a positive association – it could be concluded that an observed association is likely to be an underestimate). Where it is not possible to identify the direction or the magnitude of a bias, reviewers would need to consider using the various sensitivity analysis techniques presented in subsequent chapters before concluding that an identified source of bias would indeed be enough to change the interpretation of the study for the purposes of hazard identification.

Note also that DAGs cannot be used to solve every problem, and that there are some biases (particularly those involving the lack of a concept called faithfulness, which is beyond the scope of this chapter) that cannot easily be represented in DAGs. Another limitation of DAGs is that they do not readily depict interactions between variables.

## 2.3 Example: building a DAG for opium consumption and lung cancer

Suppose that an *IARC Monographs* Working Group comes together to evaluate whether opium consumption causes lung cancer, and that they are interested in using observational data to identify the hazard. When reviewing an observational study, such as the Golestan Cohort Study conducted in the Islamic Republic of Iran (Sheikh et al., 2020), the Working Group might be concerned that people who use opium are different from those who do not, with respect to factors that put a person at increased risk of developing lung cancer. A DAG could help the Working Group decide which variables should be controlled for in order for a study to be considered highly informative.

To begin to generate a DAG, the team would first draw the exposure and the outcome and then work through the shared causes of opium use and lung cancer as well as any other variables they think may be important in generating the data (Example 2.7a). The key is that expert knowledge is used to draw the DAG, not pure guesswork or the list of the variables that have been collected (Hernán et al., 2002).

It is critical to understand the issue hypothesized in Example 2.7a, where two variables (such as tobacco use and opium use) are associated through a third, possibly unmeasurable, latent variable (such as propensity to use substances). Having unmeasured latent factors in the DAG often creates difficulties in operationalizing what constitutes sufficient adjustment; however, such factors can still create substantial bias and

**Example 2.6c.** Conditioning to block backdoor paths in a DAG

Fig. 2.8 shows that all the unblocked backdoor paths are now blocked.

**Fig. 2.8.** Directed acyclic graph representing the data-generation process for a study of red meat consumption and colorectal cancer (CRC) with additional conditioning on family history of CRC and diabetes.

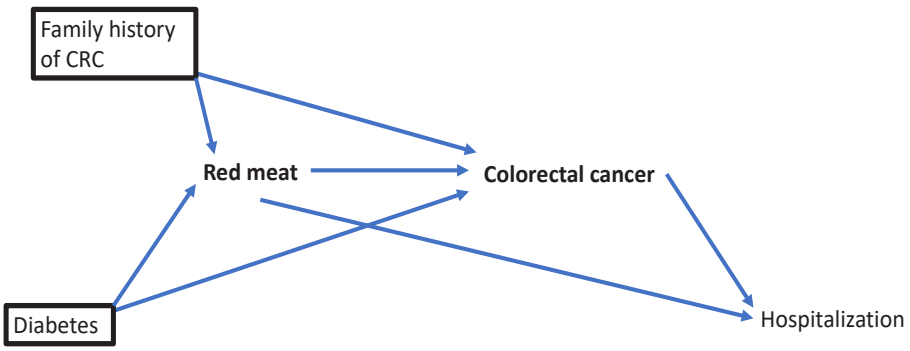


Now, suppose that the results were adjusted for family history of CRC, diabetes, and hospitalization, because it was thought that hospitalization was a confounder. As noted in Table 2.4 for the path red meat consumption → hospitalization ← CRC, the only path from red meat consumption to CRC that goes through hospitalization is not an open path; thus, it creates no bias. Would the results still be a valid estimate of the effect of red meat consumption on CRC after conditioning on hospitalization? As discussed previously, conditioning on a collider (a variable with two arrowheads into it along a pathway) creates a non-causal association between the parents (i.e. the two variables that are causes of the collider), and this pathway creates bias. This is represented in Fig. 2.9 with a box around hospitalization (representing conditioning through statistical control) and a dashed arrow from red meat consumption to CRC (representing a non-causal association that has been induced between the two variables by controlling for the collider). As noted previously, spurious associations can be induced by adjusting for variables in the analysis; these are represented with dashed lines with no arrowhead. (text continues on page 39)

**Fig. 2.9.** Directed acyclic graph representing the data-generation process for red meat consumption and colorectal cancer (CRC) with additional conditioning on family history of CRC (confounder), diabetes (confounder), and hospitalization (collider), the last of which creates a non-causal association between red meat consumption and CRC. The dashed line represents an association created by conditioning on a collider.

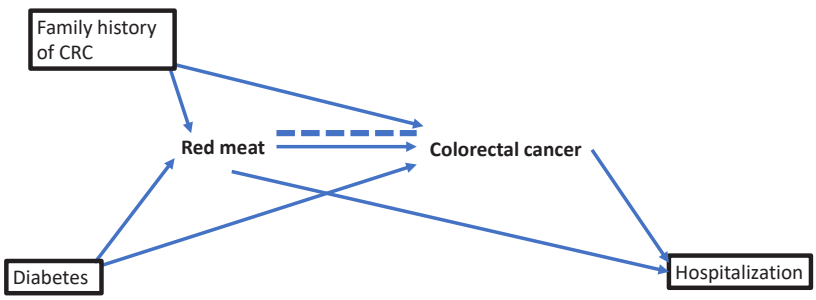

**Example 2.7a.** Depicting shared causes in a DAG for opium consumption and lung cancer

It is known that tobacco use is a cause of lung cancer, and it is observed that people who use opium are more likely to smoke tobacco. However, it is not immediately clear how this association would occur. For a DAG to help identify a set of analytical variables needed to control for confounding, it is necessary to specify how associations occur in the data. Does tobacco smoking cause opium use? Does opium use cause tobacco smoking? Although, in a minimal number of individuals, either of these could be true, the review team believes that it is more likely that there is some shared cause that links the two. This shared cause could be a propensity to use substances, or something like SES. These causes are represented in the DAG shown in Fig. 2.10, because they are both considered to be likely sources of the association between opium use and tobacco use. (text continues on page 39)

**Fig. 2.10.** Directed acyclic graph for a study of the relation between opium use and lung cancer. SES, socioeconomic status.

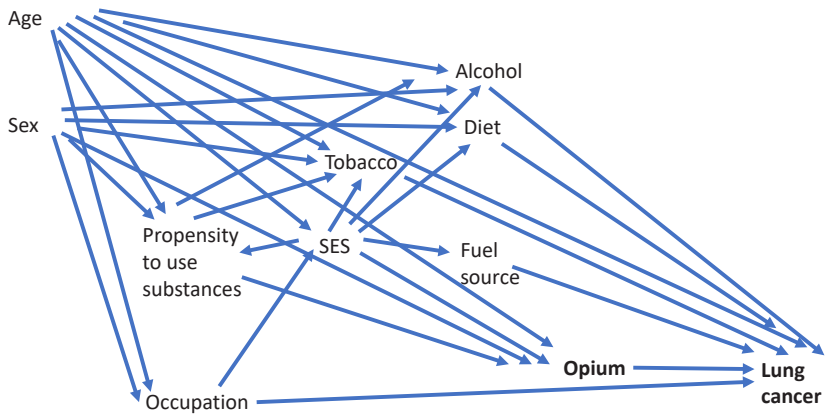


therefore should not be omitted from the DAG.

For now, selection and measurement nodes, which are discussed later in this chapter, will be ignored and left out of the DAG. If it can be assumed that the DAG is correct, this now provides a model that can be used to identify a sufficient set of variables that need to be adjusted for to control for confounding (Example 2.7b). This process will be demonstrated in the discussion on confounding in this chapter. Note here that if the study investigators have not measured all the variables in the DAG (or have not measured them well), they may not have been able to remove all the confounding directly.

## 2.4 DAGs and specific sources of bias

### *2.4.1 Confounding*

#### *(a) Identification with DAGs*

A confounder, which is a type of variable, can be distinguished from confounding, which is a bias that results from an unblocked backdoor path. A confounder is any variable that, when it has been controlled for, leads to a reduction in confounding. Given a DAG, researchers and reviewers can use simple rules to determine sets of variables whose control is sufficient to eliminate confounding bias (assuming that the variables are well measured). A set of variables is sufficient to eliminate all the confounding if (i) it blocks all open backdoor paths from an exposure *X* to an outcome *Y*, including any paths that may be opened through adjustment, and (ii) it comprises no descendant of the exposure *X* (i.e. no variable directly or indirectly influenced by the exposure) (Pearl, 2009).

Previous definitions of a confounder used statistical terminology (e.g. a confounder is a variable that is

**Example 2.7b.** Simplifying a DAG for opium consumption and lung cancer to account for confounding

Fig. 2.10 is an example of a DAG that might have been created by an *IARC Monographs* Working Group. It is again necessary to reiterate that the DAG represents the full data-generation process (i.e. the set of variables and relations between them that led to the data observed in this study), not only the variables that were measured in the study. Note also that the data-generation process for a specific exposure–outcome pair might vary across populations. For example, in some countries occupational exposures to a specific carcinogen may be prevalent, and in others they may not, so the variable would not be included.

However, suppose that the Working Group members discuss the DAG and decide that only a few of the variables have effects (represented by the arrows) that are strong enough to represent substantial bias, based on their understanding of the strength of the effect the arrows represent. The DAG might then be simplified, as in Fig. 2.11. (text continues on page 41)

**Fig. 2.11.** Simplified directed acyclic graph describing the data-generation process for the relation between opium use and lung cancer, focused on variables that are thought to be likely to cause substantial bias. SES, socioeconomic status.

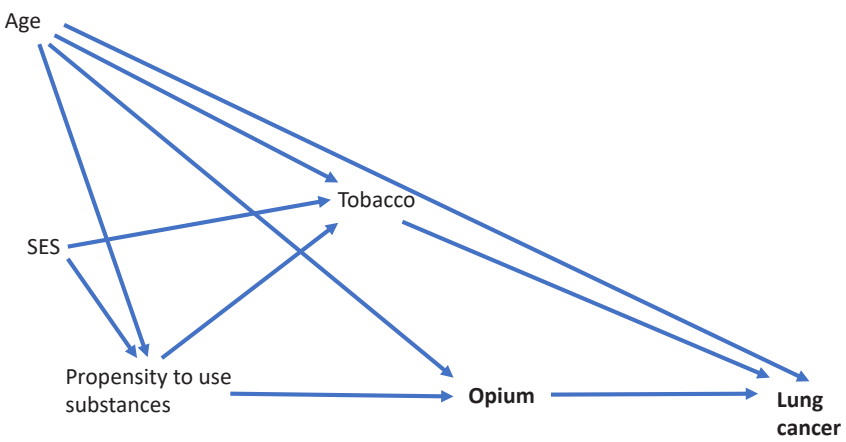


associated with both the exposure and the outcome); essentially, the definition of a confounder needs causal wording to properly distinguish it from other concepts, such as mediators (which are also variables that are associated with both the exposure and the outcome). Thus, a confounder has also been defined as any member of a minimally sufficient set of variables used for confounding control, such that dropping the variable from the sufficient set would lead to uncontrolled confounding (VanderWeele and Shpitser, 2013). Variables that are shared causes of both the exposure and the outcome qualify as confounders. Example 2.7c should make the distinction between confounders and confounding clearer.

Note, again, that free online software, such as DAGitty, can be useful here to identify all the sets of variables that would suffice to remove confounding through adjustment.

## Key message

Confounding is the bias that is created by an unblocked backdoor path. Often, different sets of variables can be used to remove the confounding; which variables are identified as the confounders depends on which are to be used to remove the confounding.

**Example 2.7c.** Confounding by backdoor paths in the DAG for opium consumption and lung cancer

As depicted in Fig. 2.11, age is a confounder of the relation between opium use and lung cancer, because it has a causal effect on both opium use (likely reducing use) and lung cancer (increasing risk). Using the DAG terminology introduced in Section 2.2.2, the path opium use ← age → lung cancer is open and may thus create a spurious association between opium use and lung cancer. In plain words, and considering the example where both opium use (yes or no) and age (old or young) are binary variables for simplicity, people who use opium are more likely to be older, so that even in the absence of a causal effect of opium use on lung cancer, an association between them is expected because of the causal effect of age on lung cancer. Of course, how much bias this creates will depend on how strongly age affects both opium use and lung cancer.

As explained in Section 2.2.2, proper control for age (e.g. by stratification, matching, or adjustment) would block the path opium use ← age → lung cancer and remove the corresponding confounding bias. However, Fig. 2.11 also shows that more-complex confounding structures can exist. For example, it might be thought that people who use opium are more likely to smoke tobacco than people who do not use opium, probably through propensity to use substances. This makes propensity to use substances a shared cause (a forking path) of both opium use and tobacco use. Because tobacco use is known to cause lung cancer, there is now an unblocked backdoor path (i.e. a confounding path): opium use ← propensity to use substances → tobacco use → lung cancer. This path can be blocked by adjusting for propensity to use substances, but this is difficult to measure. Because the path can be blocked by controlling for any variable on it, adjusting for tobacco use would also suffice to remove confounding that works through this pathway. Therefore, valid studies of the association between opium use and lung cancer would be expected to include adjustment for tobacco use. Table 2.5 lists a selection of the backdoor paths from opium use to lung cancer (but note that there are more).

**Table 2.5.** Backdoor paths from opium use to lung cancer in Fig. 2.11 and their relevance to the control of confounding

| Number | Path | Status |
|---|---|---|
| 1 | opium use ← propensity to use substances ← SES → tobacco use → lung cancer | Open, unblocked |
| 2 | opium use ← propensity to use substances → tobacco use → lung cancer | Open, unblocked |
| 3 | opium use ← age → lung cancer | Open, unblocked |
| 4 | opium use ← age → tobacco use → lung cancer | Open, unblocked |
| 5 | opium use ← sex → tobacco use → lung cancer | Open, unblocked |
| 6 | opium use ← propensity to use substances ← age → tobacco use → lung cancer | Open, unblocked |
| 7 | opium use ← propensity to use substances ← age → lung cancer | Open, unblocked |
| 8 | opium use ← age → propensity to use substances → tobacco use → lung cancer | Open, unblocked |
| 9 | opium use ← sex → tobacco use ← age → lung cancer | Closed, blocked |
| 10 | opium use ← sex → tobacco use ← SES → propensity to use substances ← age → lung cancer | Closed, blocked |
| 11 | opium use ← propensity to use substances ← SES → tobacco use ← age → lung cancer | Closed, blocked |
| 12 | opium use ← propensity to use substances → tobacco use ← age → lung cancer | Closed, blocked |

SES, socioeconomic status.

**Example 2.7c.** Confounding by backdoor paths in the DAG for opium consumption and lung cancer (continued)

In the example of Table 2.5, 12 backdoor paths are noted in a partial list of all the backdoor paths between opium use and lung cancer, of which only eight are open (1–8). The four that are blocked (9–12) are blocked because they contain a collider (e.g. sex → tobacco use ← age, sex → tobacco use ← SES, propensity to use substances ← SES → tobacco use, or propensity to use substances → tobacco use ← age). The eight unblocked paths all contain sex, age, or propensity to use substances; all of these unblocked paths could be blocked (i.e. removing all the confounding created by these forking paths) by controlling for these three variables. Although some of these variables are colliders on other paths and could thus open new paths, the new paths would all be blocked by one of the three variables in the set used to remove the confounding (Fig. 2.12). For example, propensity to use substances is a collider on the path SES → propensity to use substances ← age; if it were somehow possible to measure and adjust for propensity to use substances through stratification, regression, or some other method, a path would be opened between SES and age such that there would now be an open backdoor path: opium use ← propensity to use substances ← SES --- age → lung cancer. However, this path is already blocked by adjusting for propensity to use substances and sex; thus, no new bias is created. The key point here is that there may be instances where it is in fact necessary to control for a collider on a backdoor path from the exposure to the outcome to remove all the bias. This is fine as long as any new paths opened up by controlling for the collider are blocked.

**Fig. 2.12.** Example of directed acyclic graph for a study of a possible causal effect of opium use on lung cancer, adjusted for age, sex, and propensity to use substances. SES, socioeconomic status. Dashed lines represent associations created by conditioning on a collider.

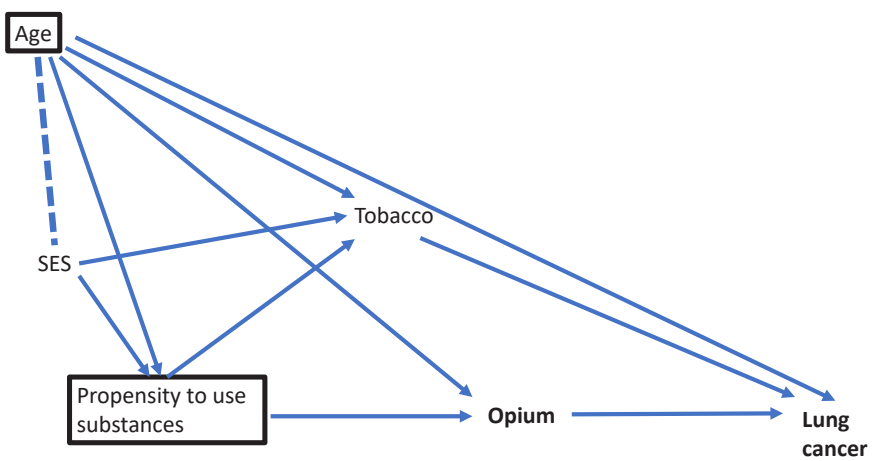


The eight unblocked paths also contain tobacco use or age (or both), so one could also adjust for both of these and block all the unblocked backdoor paths; however, note that tobacco use is a collider in each of the closed paths, so adjusting for tobacco use would open new pathways (Fig. 2.13). Nonetheless, although adjusting for tobacco use does open new pathways, none of them leads to a new unblocked backdoor path (confounding pathway) from opium use to lung cancer after adjusting for tobacco use and age, so this would also be an appropriate set. Because propensity to use substances may be difficult to measure in practice, adjustment for age and tobacco use may be an easier strategy.

**Example 2.7c.** Confounding by backdoor paths in the DAG for opium consumption and lung cancer (continued)

**Fig. 2.13.** Example of a directed acyclic graph for a study of a possible causal effect of opium use on lung cancer, adjusted for tobacco use and age. SES, socioeconomic status. Dashed lines represent associations created by conditioning on a collider.

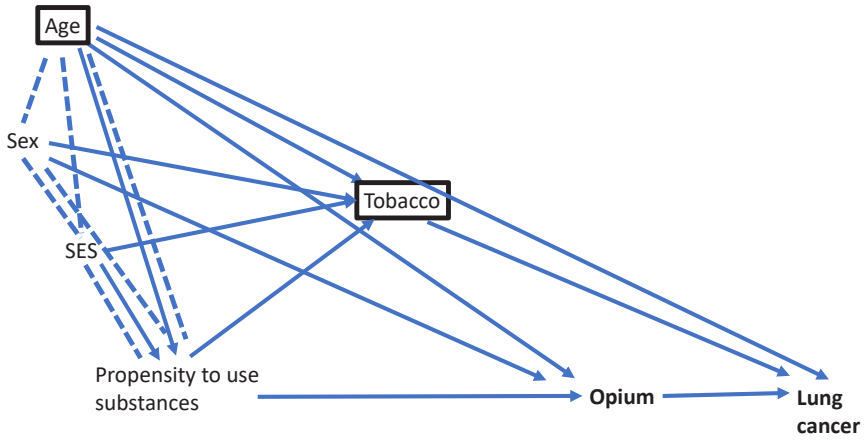


In Fig. 2.14, one variable is added: yellow fingers, which is an effect of tobacco use. The variable "yellow fingers" is a descendant of tobacco use (i.e. a variable affected by tobacco use), because it is possible to follow a directed path to get from tobacco use to yellow fingers. This demonstrates that adjustment for the descendant of a collider partially adjusts for the collider itself; in this instance, adjusting for yellow fingers (alone) would partially open a path between each of the parents. Although this may seem a silly example, it is meant to illustrate the approach used when investigators adjust for variables as a proxy solution in a situation where information on the variable they would have liked to adjust for (here, tobacco use) is missing. (text continues on page 42)

**Fig. 2.14.** Illustrative example of a directed acyclic graph for a study of a possible causal effect of opium use on lung cancer, as shown in Fig. 2.11, with an additional variable (yellow fingers). SES, socioeconomic status.

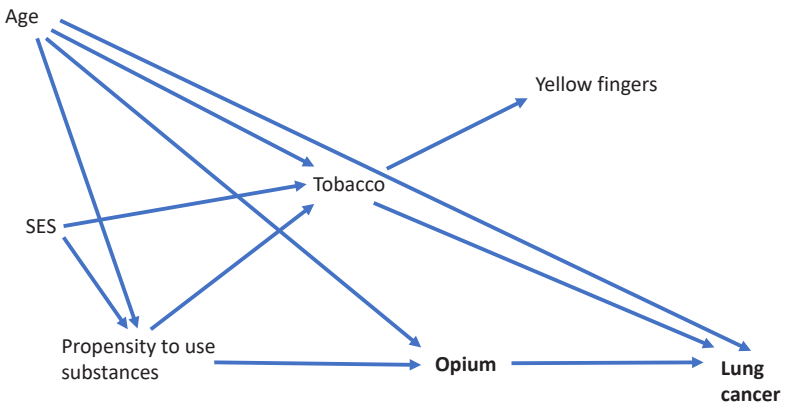

*(b) Implications for study results*

**Key message**

Given a DAG that describes the data-generation process under study, researchers and reviewers can determine whether a particular study sufficiently accounted for confounding bias. The rules given in this chapter can be used to check whether the analysis adjusted for a sufficient set of variables to remove all the confounding. If it becomes apparent that some confounding paths were not properly blocked, simple sensitivity analyses can be used to assess the magnitude of the residual confounding bias and whether this bias is likely to fully explain the observed association (see Chapters 3 and 6 for more details).

Sensitivity analyses can also be used to assess whether the size of the association between the exposure and the outcome would be larger than that observed had the bias been absent and, therefore, whether the data still suggest a cancer hazard. If the analysis involved a sufficient set of variables, confounding bias might still be present if some unobserved (or unobservable) variables were missing in the original DAG. Additional sensitivity analyses (Arah et al., 2008; VanderWeele and Arah, 2011; Arah, 2017), including analyses based on negative control exposures or outcomes (Flanders et al., 2022), can be carried out to explore this further (see Chapter 3).

### 2.4.2 Information bias

Another key source of bias that must be contended with in epidemiological research is information bias (Lash et al., 2021). As noted in the Preface, information bias results from the mismeasurement or misclassification of key variables. This section discusses ways of using DAGs to visualize different types of information bias when beginning to assess the possible impact that any mismeasurement of variables may have. This concept is discussed further in Chapter 4.

To obtain unbiased estimates of causal effects, accurate information is needed about the variables used in the study. Information bias occurs when the variables are not perfectly measured, and the mismeasured versions lead to a difference between the causal effect and the observed effect. For example, in the above-mentioned study of opium use and lung cancer, suppose that the study investigators assessed opium use with a questionnaire. Not all participants would provide accurate information about the amount of opium they typically used, for several reasons. Some may not accurately remember, and some may not want to tell the researchers, because opium use is usually illegal. Furthermore, if opium use was assessed after the lung cancer had already occurred, as may happen in a case–control study, it is possible that if the participants in the study thought there was a relation between opium use and lung cancer, the investigators may get more accurate information about those with lung cancer than those without; this could lead to a biased estimate of the true effect (often referred to as recall bias).

*(a) Types of variables affected*

All variables can be mismeasured to some degree. Although measurement error can be used as a catch-all term for mismeasured variables, mismeasurement that occurs in continuous variables is referred to as measurement error, whereas mismeasurement that occurs in categorical variables is referred to as misclassification. In both cases, it is possible to explore the impact of any potential bias created by the lack of perfect correspondence between the true value of a variable and its measured version. The next section first focuses on exposures and then discusses confounders.

*(b) Identification with DAGs*

Measurement error and misclassification can be depicted in DAGs, as demonstrated by Hernán and Cole (2009). With their approach, each factor in an analysis is represented with two variables: the true underlying variable and the measured version of that variable. Although the true version is almost never identified, a measurement approach is generally chosen that should be closely correlated with the actual values of the true variable to be measured.

Example 2.8 describes the heuristic (Lash, 2007) that many researchers rely on when they note in their discussion sections that non-differential measurement error was likely to have biased their results towards no effect, despite the fact that this can be incorrect in a number of circumstances (van Smeden et al., 2020; Yland et al., 2022). In actuality, the structures can become more complex, and the direction of the bias can become unpredictable (at least in aggregate), as will be demonstrated. However, where the bias is probably towards the null, those who are simply trying to identify a non-null causal link between an exposure and

**Example 2.8a.** Depiction of non-differential measurement error in a DAG for red meat consumption and colorectal cancer

The DAG in Fig. 2.15a shows the presumed data-generation process for a study of the relation between red meat consumption and CRC. For simplicity, assume that there are no shared causes of the two variables. A measured version of each variable is also represented with the same variable name but with an asterisk * (indicating that this is the measured version). For each variable, the measured version is affected by the true variable; this creates an association between them. This is essential, because if there were no relation, there would be no reason to use the measured version.

Fig. 2.15b shows the same DAG, but the associated error terms are added, denoted by *U* with a subscript label related to the variable of interest; these explain the difference between the measured and true versions of the variable. Adding these error terms allows for the description of different types of measurement error, which can have different impacts on the results, and therefore on the inferences to be drawn from the study. In a study in which the authors estimate the effect of red meat consumption on CRC, what can in fact be estimated is the association between red meat consumption* and CRC*. Assuming that the measurement error in each does not depend on any other variable, in a very simple scenario, the association between red meat consumption* and CRC* might be expected to be attenuated compared with the true effect of red meat consumption in causing CRC, because red meat consumption* and CRC* are imperfect proxies for the true versions. (text continues on page 46)

**Fig. 2.15.** Directed acyclic graphs for a study of a possible causal effect of red meat consumption (RM) on risk of colorectal cancer (CRC): (a) representing the data-generation process, as well as measured versions of each variable (each represented with an asterisk); (b) with the addition of an associated error term (*U*) for each variable.

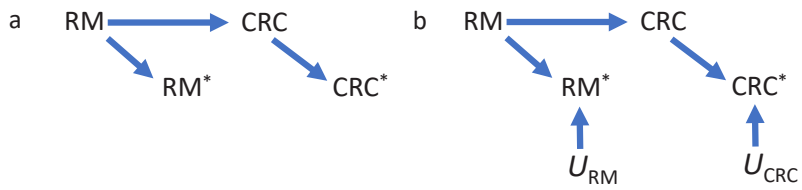


an outcome, and not the magnitude of the effect, might be able to focus less on this bias, because any observed association would probably have been stronger had the bias been absent.

This measurement error can be classified as independent, because the error terms are independent of each other, and non-differential, because the error terms do not depend on the actual value of any other variable. Here, we focus on the distinction between differential and non-differential error and leave a discussion of dependent and independent error to Side Box 2.6. For an exposure, non-differential measurement error (using measurement error as a catch-all term here) typically means that the amount of measurement error in the exposure does not depend on the actual value of the outcome (although non-differentiality could be defined with respect to another key variable in the study). In many situations, the existence of non-differential error leads to the expectation of a bias towards the null. However, because DAGs do not imply anything about the magnitude of the effect of the arrows, it is not possible to say how much bias there will be; therefore, some may not find adding nodes for measurement to be beneficial. Quantitative bias analyses (Fox et al., 2021a) can be quite helpful in this situation, as discussed in Chapter 4.

When there is differential measurement error, there is no predictable direction of the bias; it can be

**Example 2.8b.** Depiction of differential measurement error in a DAG for red meat consumption and colorectal cancer

The DAG of Fig. 2.16 takes the DAG of Fig. 2.15b as a starting point but adds an arrow from CRC to $U_{RM}$. In this situation, the error in red meat consumption is affected by the true CRC status. This is an instance of differential measurement error, such as recall bias. For example, if the data on red meat consumption were collected before the cancer diagnosis, as would be depicted in the DAG of Fig. 2.15, it might be reasonable to assume that the error in information about red meat consumption is unrelated to whether a person develops CRC. However, if data on red meat consumption were collected by self-report after a diagnosis of CRC, the DAG in Fig. 2.16 might be more likely, because misreporting of red meat consumption might be different between those who did and did not have a diagnosis of CRC. This could occur in retrospective studies because those who have a diagnosis may spend more time trying to assess their exposures and may recall them more accurately than those who do not have a diagnosis. Alternatively, if people with a diagnosis believe that the cancer was caused by red meat consumption, they might overreport their red meat consumption compared with those who did not have a diagnosis (Lash et al., 2021). The key point with differential measurement error is that the error in one variable is related to the actual value of a second key variable (e.g. error in red meat consumption is related to actual CRC status). Thus, differential exposure measurement error typically means that the amount of measurement error in the exposure does depend on the actual value of the outcome (although it could be defined with respect to another key variable in the study). (text continues on page 46)

**Fig. 2.16.** Directed acyclic graph for a study of a possible causal effect of red meat consumption (RM) on risk of colorectal cancer (CRC), as well as measured versions of each variable (each represented with an asterisk) and the associated error term ($U$) for each variable, representing independent, differential measurement error.

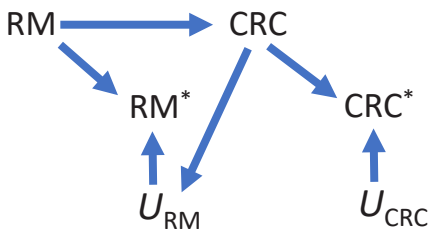


towards or away from the null, and even our intuitions on the direction can sometimes be wrong (Greenland and Robins, 1985). Chapter 4 shows that having good information about the amount of measurement error in a variable stratified by any variable the error might depend on is the key to assessing the likely direction and magnitude of the bias.

Note that DAGs cannot indicate the magnitude of bias created by any information bias. See Chapter 4 for sensitivity analyses that enable the expert reviewer to consider whether information bias could meaningfully change causal conclusions from individual studies.

### *2.4.3 Selection bias*

Although modern approaches to defining selection bias have focused on the use of causal diagrams, not all selection bias can be easily described using a DAG. Therefore, the focus here is on a mechanism by which selection bias occurs. As defined in the Preface, selection bias can occur when entry into or retention in a study is related to both the exposure and the outcome, although it should be noted that it can also occur because of the ways in which people are selected into analytical groups. In other words, selection biases create an association between exposure and outcome because of the way in which people are selected into or

**Side Box 2.6.** Dependent error

With respect to information bias, an informative scenario to consider is presented in Fig. 2.17; the error terms (unknown causes) are unrelated to the actual values of any other key analytical variables, but the error terms for both the exposure, $U_{RM}$, and the disease, $U_{CRC}$, have a shared cause, $U_{RM\text{-}CRC}$. This may occur if the same source (perhaps self-report) was used for both the exposure and the outcome. In such situations, the errors in the two variables are correlated, leading to non-differential but dependent measurement errors. This is also sometimes referred to as common-method or common-source bias (Podsakoff et al., 2003). In this situation, as was true in each of the previous DAGs (Fig. 2.15), because both red meat consumption and CRC were mismeasured, one would expect error in both variables. However, unlike in the previous two DAGs, the errors here are correlated with each other. This may be easiest to understand with a dichotomous exposure and a dichotomous outcome. If the red meat consumption (high or low) of 10% of study participants was incorrectly classified and the CRC status (yes or no) of 10% of study participants was incorrectly classified, then one would expect misclassification on both variables for 1% (the product of those two percentages) of study participants. However, in the DAG of Fig. 2.17, because the errors are correlated or dependent, one would expect misclassification on both variables for more than 1% of study participants. This is because if self-report was used for both the exposure and the outcome, people who are more likely to overreport their exposure might be more likely to overreport their outcome, and vice versa. In this scenario, as demonstrated in articles by Kristensen (1992) and Chavance and Dellatolas (1993), small amounts of non-differential but dependent measurement error can lead to strong bias away from the null for a truly null effect. Thus, to obtain valid estimates of the effect of an exposure on an outcome, it is critical to separate the sources for data on key variables in the study (Brennan et al., 2021). For example, if self-report was being used for red meat consumption, a medical record could be used to obtain information on the CRC diagnosis. Both could still be measured with error, but because the errors would not be correlated, the impact of the bias would often be smaller. Bias analyses are quite difficult to implement for dependent errors; therefore, *IARC Monographs* Working Groups should be cautious when reviewing studies that may contain dependent error.

**Fig. 2.17.** Directed acyclic graph representing the data-generation process for red meat consumption (RM) and colorectal cancer (CRC), as well as measured versions of each variable (each represented with an asterisk) and the associated error term ($U$) for each variable, representing dependent, non-differential measurement error.

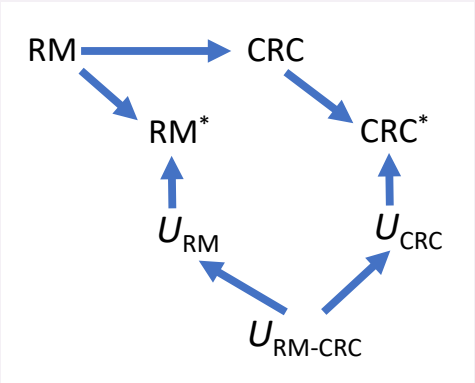


Although this is not shown here, errors in measurement can also be both dependent and differential, creating a very unpredictable and potentially strong bias. In such situations, it can be nearly impossible to assess the true underlying causal effect of an exposure on an outcome. (text continues on page 47)

out of a study. In nearly all studies, except those that use a census of the study population such that there is no selection into the study, participants are selected into the study either by the investigators or by their self-selection into the study (or a combination of the two). Selection alone does not always lead to selection bias; it is only when the forces that lead people to be selected into or out of a study or the ways in which researchers select people into or out of analytical groups distort the true causal effect for the target population, leaving a biased association (see Chapter 5 for more information). With causal diagrams, it is easier to demonstrate when selection bias occurs. For now, note that an example of selection bias would be selection on or adjustment of a shared effect of the exposure and the outcome (i.e. a collider).

## Key message

Selection bias can be produced at the time of study entry, at the time of sampling into a study, at the time of selection out of a study (e.g. loss to follow-up), or during analysis (analytical selection).

### *(a) Description and mechanisms*

As is shown in more detail in Chapter 5, selection bias is a common issue across all study designs (Lash and Rothman, 2021). In randomized trials, there can be selection bias due to loss to follow-up. In cohort studies, selection bias can arise because of how the cohort is selected. Case–control studies can have selection bias due to inappropriate choice of control participants. This is not an exhaustive list but underscores the ubiquity of the problem. This section connects the commonality of these biases via DAGs.

### *(b) Depiction and identification with DAGs*

Section 2.2.4 reviews how a closed backdoor path can be opened by conditioning on a collider (or a descendant of a collider) on that pathway. Such collider biases can occur from selection into the analytical dataset (Example 2.9).

Next, Example 2.10 elaborates on this simple causal DAG in the setting of a case–control study.

Loss to follow-up in any longitudinal study (e.g. randomized trials or cohort studies) can also create a selection bias, which can be depicted through conditioning on a collider in a causal diagram. In situations where loss to follow-up creates a bias, the time under observation in the study is related to the exposure and the outcome (Example 2.11).

Note that there are other ways in which loss to follow-up can be drawn in DAGs, but all of these structures reduce to the same issue: if we analyse only people who happened to continue to be observed in the study without further adjustment, we might be conditioning on a collider or a descendant of a collider in a path between exposure and outcome, as drawn in the DAG.

Examples 2.10 and 2.11 are only two ways in which causal graphs may depict selection bias; Chapter 5 describes others in detail. Let us now turn our attention to what can be done to avoid, address, or mitigate selection bias, and the role of DAGs in that process.

### *(c) Implications for study results*

What can be done about selection biases? Returning to the DAGs in Fig. 2.20, there would be no biasing pathway if there were no box around *S*. But this, of course, is not usually a realistic situation and is beyond the control of someone trying to analyse or review existing data. However, when studies have loss to follow-up greater than some de minimis value (e.g. 5%), the approaches to evaluating the sensitivity of results described in Chapter 5 could be helpful.

Selection bias can sometimes be minimized by design. For example, choosing control participants in a case–control study such that it is unlikely that a path exists between the proposed causal agent and the selection of the control group minimizes the bias created in the DAG in Fig. 2.19a, even if it cannot be guaranteed to prevent it completely. As another example, in studies where outcome assessments are obtained from routinely collected data rather than onerous study visits, loss to follow-up may be minimized through this reduced participant burden (although at a potential cost of information bias).

Even if selection bias has not been minimized by design, reviewers can use DAGs to ensure that in the studies being evaluated the analytical steps were taken to mitigate the bias to the best extent possible with the available data. The graph for loss to follow-up in Fig. 2.20b shows that the open pathway can be closed by adjusting for a variable on the newly opened path, namely SES. Chapter 5 describes in more detail the options for these adjustments, as well as ways to reason about the direction

**Example 2.9.** Depiction of collider bias (by hospitalization)

Suppose, for a study of the association between red meat consumption and CRC, that the DAG in Fig. 2.18a depicts the data-generation process but that in this situation the study being reviewed was conducted only among hospitalized patients. In other words, the study design conditioned on hospitalization, as shown in the DAG in Fig. 2.18b. Conditioning on hospitalization opens up the path red meat consumption → hospitalization ← CRC; therefore, there is an open path between red meat consumption and CRC other than the causal path of interest, and this new open path could explain any observed association between red meat consumption and hospitalization in the dataset. (text continues on page 50)

**Fig. 2.18.** Simple selection-bias diagrams showing (a) selection on hospitalization as a collider and (b) bias from conditioning on hospitalization. The dashed line represents an association created by conditioning on a collider.

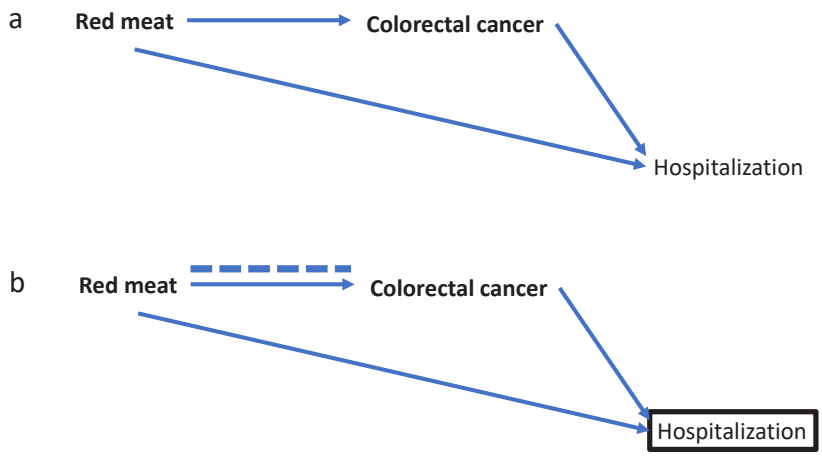


and magnitude of bias when adjustment is not possible.

## Key message

With respect to DAGs and review panels, perhaps the most useful implication of DAGs for selection bias is in identifying when selection bias is likely in a published study, and then using DAGs as a guide to inform a possible bias analysis (Fox et al., 2021b; Chapter 4) or sensitivity analysis for whatever remaining biases exist within the evidence at hand.

## 2.5 DAGs and multiple sources of bias

### *2.5.1 Identifying multiple sources of bias*

As noted previously, the full data-generation process can be represented in DAGs by including those sources of bias that occur in the population (e.g. confounding) and those that occur because of the study (e.g. selection bias and information bias); this will allow for a full picture of the ability of a study to identify causal effects from the observational data. When assessing the full impact of bias on study results, it may be necessary to think through how sources of bias interact with each other. It is not immediately clear from looking at a DAG whether two sources of bias will be additive in terms of their impact on study results or, if they act in opposite directions, whether they might cancel each other out (Greenland, 2005).

A simple scenario of how two sources of bias might interact with each other is information bias (in the measurement of exposures, outcomes, or other variables) and confounding. Suppose that a DAG representing the underlying data-generation process is used to identify a sufficient set of variables to control for all the confounding.

**Example 2.10.** Selection bias in a DAG for opium consumption and lung cancer

Consider an investigation of the effects of opium use on lung cancer that included a case–control study in which participants in the control group were selected from hospitalized patients. Let *S* denote an indicator of being included in the case–control study. In a case–control design, there is an arrow from the outcome to selection (lung cancer → selection [*S*]) by definition: having lung cancer (*Y* = 1) increases the probability of being selected into the study as a case participant (*S* = 1). Ideally, control participants are selected so that they represent the exposure distribution that gave rise to the case diseases; therefore, there should be no arrow from the exposure to selection (Fig. 2.19a). However, perhaps in this hypothetical study control participants were selected who had been hospitalized for other reasons, and people who were hospitalized were more likely to be older than the general population. In that situation, the DAG may look more like the DAG in Fig. 2.19b, where this choice of control participants creates a biasing pathway (lung cancer --- age → opium use). Side Box 2.7 describes selection bias in matched case–control studies. (text continues on page 50)

**Fig. 2.19.** Selection-bias diagram showing (a) selection (*S*) as a collider and (b) bias from conditioning on selection in a case–control study with control participants selected from among people with a condition related to the exposure. The dashed line represents an association created by conditioning on a collider.

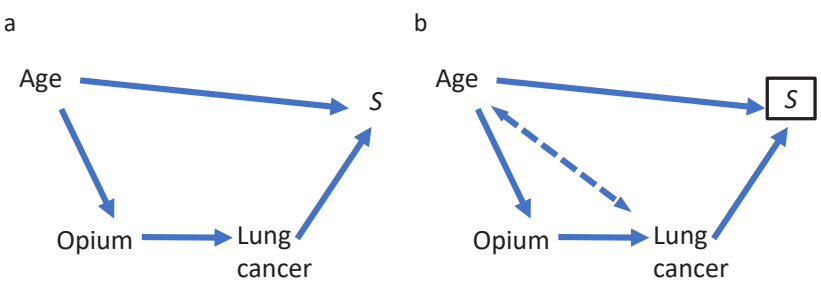


**Side Box 2.7.** Selection bias in matched case–control studies

Note that the type of biasing pathway shown in Example 2.10 also occurs in matched case–control studies in which confounders of exposure and outcome are chosen as matching factors. The matching creates a selection bias; this is why in such studies it is necessary to adjust for matched variables to remove the bias (i.e. block the backdoor path that is opened by the matching). Matching in a case–control study does not remove confounding, as is often thought; rather, it creates an efficient population within which to control for the confounding. This is sometimes referred to as selection bias by design. (text continues above)

**Example 2.11.** Depiction of collider bias (from loss to follow-up)

Consider a hypothetical cohort study of opium consumption, as shown in Fig. 2.20a, in which opium use is directly related to why participants continue to be observed in the study (i.e. were not lost to follow-up), and in which SES affects both cancer risk and the likelihood of remaining observed. This might occur if people who use opium are less likely to continue in a study than those who do not. Here again, there is a collider on the pathway opium use → selection ← SES → lung cancer. If the collider on this pathway were not conditioned on, it would not create any bias because this path is closed due to the collider. Because in this situation selection represents loss to follow-up, by definition selection has been conditioned on, because it is only possible to analyse people for whom there are data, and this gives the DAG in Fig. 2.20b, which now has an open unblocked backdoor path from opium use to lung cancer: opium use --- SES → lung cancer. (text continues on page 50)

**Fig. 2.20.** Simple selection-bias diagram showing (a) loss to follow-up as a collider (*S*) and (b) bias from conditioning on loss to follow-up. *S*, selection; SES, socioeconomic status. The dashed line represents an association created by conditioning on a collider.

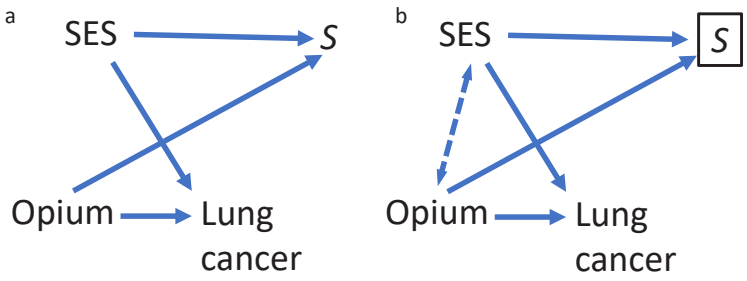


However, suppose in addition that there is an imperfect measure of the key confounders. Although the DAG can show which variables are necessary to remove the impact of confounding, the mismeasurement of those variables can lead to imperfect control. If the mismeasurement is severe enough, the residual confounding will be quite strong. Approaches to evaluating the direction and magnitude of such residual confounding are described in Chapter 3.

### *2.5.2 Representing and identifying multiple sources of bias in a DAG*

Representing the data-generation process to identify confounding, adding selection nodes to represent the selection of the study population and possibly nodes to represent selection out of the study (i.e. loss to follow-up), and adding nodes to represent the measured version of each variable and any biasing structures related to the error terms can create a very complex DAG. Although this process would ideally be followed for all variables, it may be helpful to focus on the variables that represent the largest sources of bias. However, this is challenging, because without knowing the impact that a particular source of bias has (say through a bias analysis method, described in later chapters), we are left with our intuition and our expert experience as to which biases are most important. It is recommended to start with as complete a DAG as possible for a particular study or set of studies and then remove biasing pathways that are thought to have minimal impact on the study results. See Chapter 6 for methods on this topic for triangulation.

## 2.6 Signed DAGs

The DAGs introduced thus far do not directly indicate the direction of a bias, but signed DAGs offer an approach that aids in identifying the direction of a bias. While signed DAGs can clarify many forms of bias (VanderWeele and Hernán, 2012), the focus here is on their use in understanding confounding. Suppose that an *IARC Monographs* Working Group is considering one uncontrolled

(dichotomous) confounder that does not modify the effect of the exposure on the outcome on the chosen effect measure scale (e.g. relative risk or risk difference) and wishes to understand whether this source of uncontrolled confounding is likely to explain all or some of the observed non-null association.

Signed DAGs are augmented to contain + or − symbols along the arrows to indicate the net or average direction of the effect (VanderWeele et al., 2008). A positive sign (+) indicates that an increase in (or the presence of) the variable at the tail of the arrow leads to an average increase (or no change) in the variable at the arrowhead, while a negative sign (−) indicates that an increase in (or the presence of) the variable at the tail of the arrow leads to an average decrease (or no change) in the variable at the arrowhead (see Side Box 2.8).

When one thinks about paths that can run between several variables and therefore have several arrows, rather than a path that is simply between two variables and has a single arrow, the sign of a path in a signed DAG is given by the product of the signs of its component arrows. An *IARC Monographs* Working Group that is interested in assessing the likely direction of confounding can begin by augmenting an existing DAG (as described in the previous sections) with these + or − symbols to represent the well-informed hypothesized direction of the relations.

To demonstrate how signed DAGs work, Example 2.12 extends the simple DAG shown in Fig. 2.4b. With two arrows and two possible signs that could be applied to the arrows, there are four possible scenarios and two possible results; a positive sign in the result describes the direction of the confounding as representing positive or upward bias (i.e. the bias leads to an observed estimate that is higher than the true effect), and a negative sign represents negative or downward bias (i.e. the bias leads to an observed estimate that is lower than the true effect). The net direction of the confounding created by each scenario follows the multiplication rules of positive and negative numbers, as shown in Table 2.6 and Fig. 2.23.

**Side Box 2.8.** Interpreting lack of change in signed DAGs

When interpreting signed DAGs, it may seem odd that "or no change" is included; it might be assumed that having no arrow would imply no change. This would be a reasonable assumption, but a lack of an arrow specifically implies no effect of the exposure on the outcome for any individual in the population (i.e. the sharp null). In contrast, there could be no average effect in the presence of an arrow if the number of people who experienced harmful effects was the same as the number of people who experienced preventive effects, such that the observed association averaged to the null. This might occur if the exposure prevented the outcome for some people in the population and caused it for other people, as might occur for seat belt use and death in an automobile accident. Although in this example it would be unlikely that the number of people for whom the exposure causes the outcome would be the same as the number of people for whom it prevents the outcome, in some exposure–outcome pairs such a result may be possible. In such a situation, on average, the exposure would be inferred to have no effect, even though for some people the exposure caused the outcome and for other people it prevented the outcome.

In this chapter, signs are only used in signed DAGs under the weak monotonicity assumption of non-decreasing (i.e. positive) or non-increasing (i.e. negative) average causal effects to assess the sign or direction of uncontrolled confounding due to an unmeasured confounder. Under this monotonicity assumption, a positive average monotonic effect, depicted as a positive sign on an arrow, means that increasing the value of the variable at the tail of the arrow always increases or leaves unchanged the average value of the variable at the arrowhead, for all values of the other covariates adjusted for in the analysis, in the entire population. Similarly, a negative average monotonic effect, depicted as a negative sign on an arrow in the DAG, means that increasing the value of the variable at the tail of the arrow always decreases or leaves unchanged the average value of the variable at the arrowhead, for all values of the other covariates adjusted for in the analysis, in the entire population. (text continues above)

**Example 2.12a.** Depiction of signed DAGs

In this example, there is a concern that a family history of CRC might confound an estimate of the effect of red meat consumption on CRC, as in Fig. 2.21. Furthermore, suppose that the study results indicated an increased risk of CRC associated with red meat consumption. As explained earlier in this chapter, the unblocked backdoor path red meat consumption ← family history of CRC → CRC, which represents a source of confounding bias, would need to be addressed to determine the causal effect of red meat consumption on CRC. If the study being assessed did not control for family history of CRC, then before dismissing the study, the *IARC Monographs* Working Group would want to decide whether the uncontrolled confounding might explain the finding. In other words, the reviewers would want to know: if family history of CRC had been controlled for in the analysis of the study, is it at least possible that the true effect would have been null? Here, signed DAGs can help.

**Fig. 2.21.** Fork structure denoting confounding by family history of colorectal cancer in a study of a possible causal effect of red meat consumption on risk of colorectal cancer.

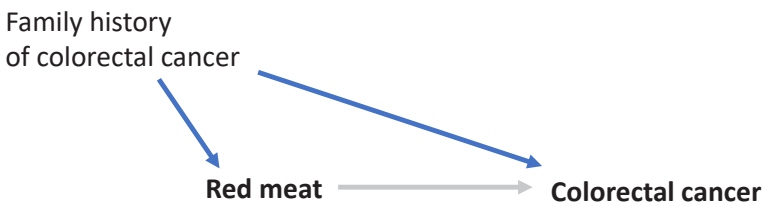


The first step in using a signed DAG is to hypothesize about the direction of the effect of the blue arrows. The arrow from red meat consumption to CRC has been left grey to indicate that this is the causal relation that is to be assessed. Fig. 2.22 shows the hypotheses about the blue arrows. On average, a family history of CRC is expected to increase the risk of developing CRC, perhaps due to a genetic predisposition; this is depicted with a positive sign to indicate a positive association. Furthermore, a family history of CRC is hypothesized, on average, to decrease red meat consumption, given the awareness of a potential link between the two and a desire of people with a family history of CRC to avoid developing the disease. This hypothesis of a negative association is depicted with a negative sign in the DAG. In this scenario, it is possible to identify the likely expected direction of this bias. (text continues on page 56)

**Fig. 2.22.** Signed DAG for assessing the direction of confounding by family history of colorectal cancer in a study of a possible causal effect of red meat consumption on risk of colorectal cancer.

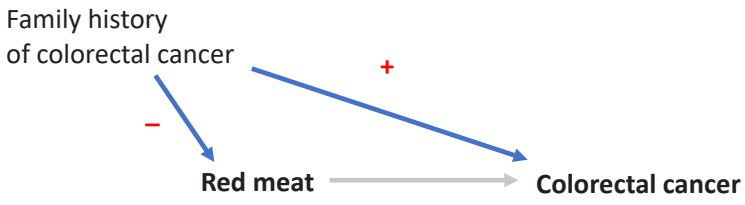

**Example 2.12b.** Using signed DAGs to determine the possible impact of biases

Returning to the signed DAG in Fig. 2.22, because the arrow from family history of CRC to CRC is positive and the arrow from family history of CRC to red meat consumption is negative, the probable net bias in the association between red meat consumption and CRC in a study in which family history of CRC was not adjusted for would be downwards or negative. This means that if a positive association (e.g. relative risk [RR] = 1.6) was observed between red meat consumption and CRC, because the bias was likely to be downwards (towards the null), if there had been data on family history of CRC and it was adjusted for, the estimate of the effect would be expected to be even larger than what was observed (in this example, RR > 1.6). In other words, because the negative uncontrolled confounding from the signed DAG and the estimated positive association from the study have opposite signs, the observed association probably underestimated the unobserved effect adjusted for the unmeasured confounder. Accordingly, such a study could not be dismissed, given that the goal was to determine whether consumption of red meat is carcinogenic and not the magnitude of the effect (which would indeed be biased).

Suppose, however, that the *IARC Monographs* Working Group encounters a study in which the observed association was that red meat consumption was associated with a reduced risk of CRC (e.g. RR = 0.8, indicating a negative association), but the study also did not adjust for a family history of CRC. In this situation, the Working Group would make all the same assumptions as before, that family history of CRC increases risk of CRC but decreases red meat consumption, yielding negative uncontrolled confounding. However, because the observed association was negative (i.e. protective against cancer), the expected bias, which is also negative, could have been part of the observed association, and adjusting for the unmeasured family history of CRC could have removed some or all of the observed association between red meat consumption and CRC. Thus, the result would probably have been less negative (closer to the null, or even positive) than what was observed (in this example, RR > 0.8). In this scenario, strong conclusions cannot be drawn. The true unbiased result could have been a less protective, a null, or a harmful effect of red meat consumption, in which the negative uncontrolled confounding was strong enough to induce some or all of the negative association or to mask a weaker positive (thus, harmful) effect, leading to the observation of a protective association. With only a signed DAG, it is not possible to tell which is correct, and the sensitivity analysis approaches described in later chapters would become essential. (text continues below)

If both arrows are positive (represented by + signs) or negative (represented by − signs), the likely direction of the net bias will be positive or upwards (represented by a + sign), because multiplying two numbers with the same sign will result in a positive number. If the two arrows have opposite signs, the likely direction of the net bias will be negative or downwards (i.e. towards the null for a positive association), represented by a − sign.

If the signs of the uncontrolled confounding and the observed (biased) study estimate are opposite, it could be concluded that the true bias-adjusted effect would have been in the same direction as observed in the biased study estimate. Such cases can still allow imperfect evidence to contribute informative information to support, rather than detract from, a given evaluation.

## 2.7 Use of DAGs in evidence synthesis

In the synthesis of the evidence across a number of studies with different study designs and different study populations, there is unlikely to be a single DAG that can describe the data-generation process in full. However, it can be helpful in evidence synthesis to begin with a working DAG that can be adapted to study-specific assessments to identify the potential limitations of each study and identify a set of variables that are likely to be necessary to control for confounding. It is also helpful for a group conducting evidence synthesis to work through the working DAG to ensure that assumptions are clearly understood between the group members and to identify areas of disagreement.

**Table 2.6.** Likely direction of confounding bias in the simplified scenario of a single uncontrolled confounder (*C*) for the directed acyclic graph (DAG) in Fig. 2.22, if the monotonicity assumptions for signed DAGs are met

| Sign of arrow 1 from family history of colorectal cancer to red meat consumption ($C \rightarrow X$) | Sign of arrow 2 from family history of colorectal cancer to colorectal cancer ($C \rightarrow Y$) | Likely direction of confounding |
|---|---|---|
| + (*C* increases risk of *X*) | + (*C* increases risk of *Y*) | + (positive[a]) |
| − (*C* decreases risk of *X*) | − (*C* decreases risk of *Y*) | + (positive[a]) |
| + (*C* increases risk of *X*) | − (*C* decreases risk of *Y*) | − (negative[b]) |
| − (*C* decreases risk of *X*) | + (*C* increases risk of *Y*) | − (negative[b]) |

*C*, uncontrolled confounder (family history of colorectal cancer); *X*, exposure (red meat consumption); *Y*, outcome (colorectal cancer).
[a] Positive uncontrolled confounding: not adjusting for *C* induces a positive association between *X* and *Y*, even when *X* does not affect *Y*.
[b] Negative uncontrolled confounding: not adjusting for *C* induces a negative association between *X* and *Y*, even when *X* does not affect *Y*.

For case–control studies, arrows from the outcome to the selection node will need to be included. For studies in which healthy worker biases are common, it may be essential to add nodes that describe the selection and confounding biases created as a result. Moreover, different measures used for different variables, or the timing of those measures, may lead to different information bias structures.

Researchers may find signed DAGs less useful for complex scenarios, for example in situations when they are trying to use signed DAGs and anticipate selection bias, non-monotonic effects, complex confounding structures, effect heterogeneity, and so on. Readers will find it helpful to refer to the more detailed discussions of signed DAGs in the literature (VanderWeele et al., 2008; Lipsky and Greenland, 2022). The following chapters provide information on other tools for understanding the direction of bias.

Finally, something that has not been mentioned yet is that DAGs can also be useful for non-traditional analyses of data from cohort or case–control studies, including the use of instrumental variable methods, Mendelian randomization approaches, and other quasi-experimental designs that may be used in triangulation processes for evidence synthesis (Swanson, 2015). In fact, some of the principles described in this chapter can help in reasoning about bias in those studies, too. For example, loss to follow-up can create a selection or collider stratification bias in such studies, and drawing a DAG can help to understand why (Swanson, 2019).

As noted previously, it is always challenging to draw DAGs that truly represent the underlying data-generation process. It may be helpful to consult a review of published DAGs for examples (Tennant et al., 2021). Because disagreements about the structure of the DAG can occur, it can be helpful to draw more than one DAG, to tease out the different assumptions that members of a group conducting evidence synthesis may have about a particular study. This can guide critical sensitivity analyses in evidence synthesis (Mathur and VanderWeele, 2020a, b, 2022).

## 2.8 Summary

DAGs make different assumptions about the data-generation process explicit, enable the identification of areas of disagreement in those assumptions between members of a group conducting evidence synthesis, and help to identify important sources of bias in the individual studies and the collective body of evidence being used to identify hazards. Working through DAGs collectively can create a motivation for additional bias analyses or sensitivity analyses that can be used to identify which sources of bias are most likely to matter in drawing conclusions about a particular hazard. DAGs also provide a systematic way of identifying critical variables for valid estimation of the effect of an exposure on an outcome. Thus, they provide a useful tool for hazard identification, as a place to communicate the working model used to make judgements about the quality of the underlying studies, and serve as a model for using the evidence presented in the most efficient way possible.

## Key message

Table 2.7 presents the likely conclusions that can be drawn from the results of a study and the results of a simple signed DAG with a single confounder, about whether the exposure is likely to have an effect on the outcome.

**Table 2.7.** Likely conclusions that can be drawn from a study about the existence of a non-null effect of an exposure on an outcome, based on the direction of confounding diagnosed with a signed directed acyclic graph (DAG) and the direction of the observed association

| Observed association of exposure with cancer indicates | Signed DAG indicates that confounding is likely to be | Conclusion that can be drawn about the existence of a true (bias-adjusted) non-null effect (bias-adjusted RR, RD) |
|---|---|---|
| Elevated risk[a] (observed RR > 1, RD > 0) | Positive | *Unclear*<br>Adjusting for the confounder would probably remove some or all of the estimate of the effect. Thus, it is not possible to say whether the estimate of the effect adjusted for the confounder would indicate increased, null, or decreased risk. |
| | Negative[a] | *Elevated cancer risk from exposure is likely*<br>The observed estimate probably underestimates the true effect. Thus, adjusting for the unmeasured confounder would probably increase the estimate of the effect (bias-adjusted RR > observed RR; bias-adjusted RD > observed RD). Adjustment for the confounder would not bring the result back to the null or flip its direction. |
| No change in risk (observed RR = 1, RD = 0) | Positive | *Masked reduced cancer risk is likely*<br>Adjusting for the unmeasured confounder would probably reveal a negative effect estimate (bias-adjusted RR < 1; RD < 0), indicating a probable reduced risk associated with the exposure. |
| | Negative | *Masked elevated cancer risk is likely*<br>Adjusting for the unmeasured confounder would probably reveal a positive effect estimate (bias-adjusted RR > 1; RD > 0), indicating a probable elevated risk associated with the exposure. |
| Reduced risk (observed RR < 1, RD < 0) | Positive | *Reduced cancer risk from exposure is likely*<br>Adjusting for the confounder would probably decrease the estimate of the effect. Uncontrolled confounding by this factor is unlikely to explain the observed result (i.e. adjustment for the confounder would not bring the result back to the null). |
| | Negative | *Unclear*<br>Adjusting for the confounder would probably remove some or all of the estimate of the effect. Thus, it is not possible to say whether the estimate of the effect adjusted for the confounder would indicate decreased, null, or increased risk. |

RD, risk difference; RR, relative risk.

[a] Indicates a scenario that would be most applicable to an *IARC Monographs* Working Group assessing whether an exposure could be carcinogenic (assuming positively coded exposure and cancer outcome variables, such that a positive exposure–outcome association with RR > 1 or RD > 0 would indicate harm).

**Fig. 2.23.** Possible results for the direction of bias as diagnosed with a signed DAG. The left side of each scenario shows the hypothesized direction (positive or negative) of the arrow, and the right side of each scenario depicts the likely direction (positive or negative) of the net bias in the *X*–*Y* relation.

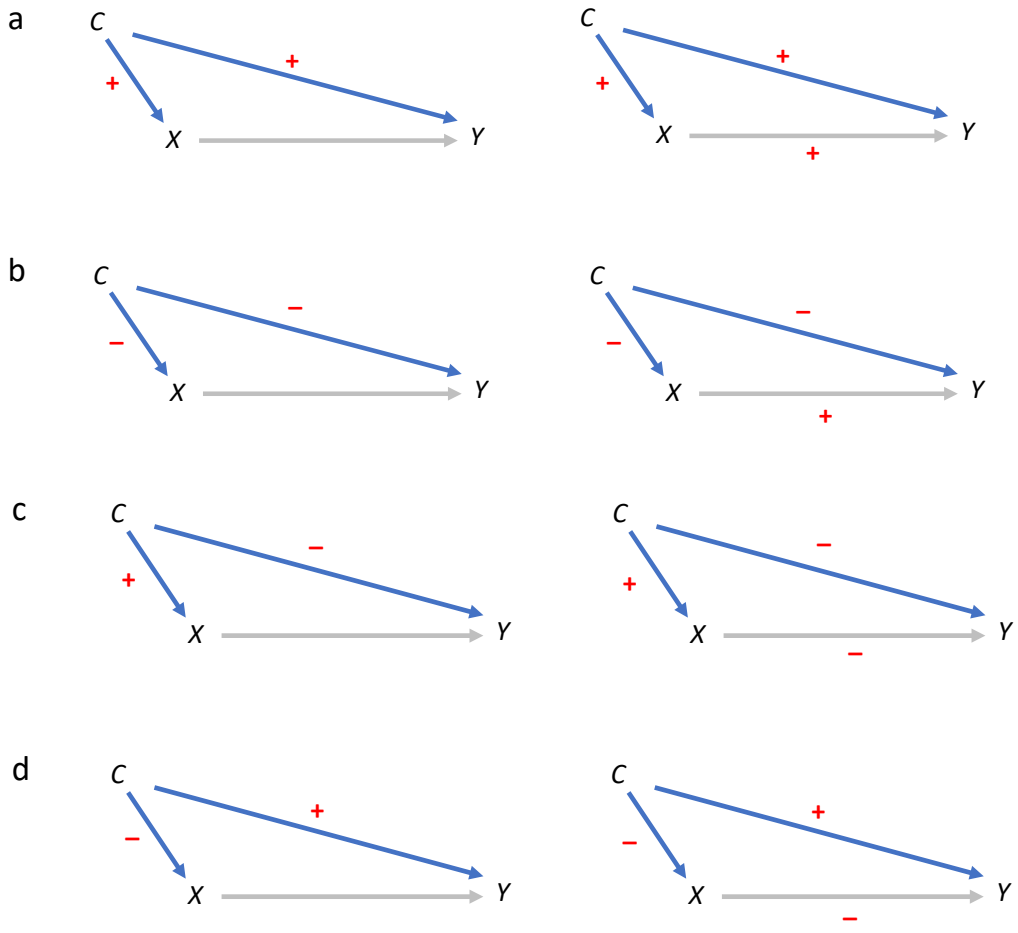


It should be cautioned that DAGs that depict the full data-generation process, capturing information bias and selection bias, can make it seem impossible to approximate the causal effect. In some circumstances, this will indeed be true, but because the magnitude of the bias cannot be demonstrated in DAGs, it can be easy to think that all potential sources of bias are equal, are additive, and are severe, when in fact this may not be true. The following chapters discuss ways to identify the possible magnitude of the impact, so that sources of bias that have minimal impact can be ignored. Because DAGs do not represent the amount of bias created, they can lead to excessive concerns about some sources of bias. In such situations, bias analyses can help to sort out which sources are most likely to matter; thus, the DAG is only a first step.

# References


Arah OA (2017). Bias analysis for uncontrolled confounding in the health sciences. Annu Rev Public Health. 38(1):23–38. doi:10.1146/annurev-publhealth-032315-021644 PMID:28125388

Arah OA, Chiba Y, Greenland S (2008). Bias formulas for external adjustment and sensitivity analysis of unmeasured confounders. Ann Epidemiol. 18(8):637–46. doi:10.1016/j.annepidem.2008.04.003 PMID:18652982

Berkson J (1946). Limitations of the application of fourfold table analysis to hospital data. Biometrics. 2(3):47–53. doi:10.2307/3002000 PMID:21001024

Brennan AT, Getz KD, Brooks DR, Fox MP (2021). An underappreciated misclassification mechanism: implications of nondifferential dependent misclassification of covariate and exposure. Ann Epidemiol. 58:104–23. doi:10.1016/j.annepidem.2021.02.007 PMID:33621629

Chavance M, Dellatolas G (1993). Bias from dependent misclassification. Epidemiology. 4(2):180–2. doi:10.1097/00001648-199303000-00015 PMID:8318094

Flanders WD, Waller LA, Zhang Q, Getahun D, Silverberg M, Goodman M (2022). Negative control exposures: causal effect identifiability and use in probabilistic-bias and Bayesian analyses with unmeasured confounders. Epidemiology. 33(6):832–9. doi:10.1097/EDE.0000000000001528 PMID:35895515

Fox MP, MacLehose RF, Lash TL (2021a). Misclassification. In: Applying quantitative bias analysis to epidemiologic data. Statistics for biology and health. 2nd ed. Cham, Switzerland: Springer; pp. 141–195. doi:10.1007/978-3-030-82673-4_6

Fox MP, MacLehose RF, Lash TL (2021b). Selection bias. In: Applying quantitative bias analysis to epidemiologic data. Statistics for biology and health. 2nd ed. Cham, Switzerland: Springer; pp. 75–103. doi:10.1007/978-3-030-82673-4_4

Glymour C, Scheines R (1986). Causal modeling with the TETRAD program. Synthese. 68(1):37–63. doi:10.1007/BF00413966

Greenland S (2005). Multiple-bias modelling for analysis of observational data. J R Stat Soc Ser A Stat Soc. 168(2):1–25. doi:10.1111/j.1467-985X.2004.00349.x

Greenland S, Pearl J, Robins JM (1999). Causal diagrams for epidemiologic research. Epidemiology. 10(1):37–48. doi:10.1097/00001648-199901000-00008 PMID:9888278

Greenland S, Robins JMM (1985). Confounding and misclassification. Am J Epidemiol. 122(3):495–506. doi:10.1093/oxfordjournals.aje.a114131 PMID:4025298

Hernán MA, Cole SR (2009). Invited Commentary: Causal diagrams and measurement bias. Am J Epidemiol. 170(8):959–62, discussion 963–4. doi:10.1093/aje/kwp293 PMID:19755635

Hernán MA, Hernández-Díaz S, Werler MM, Mitchell AA (2002). Causal knowledge as a prerequisite for confounding evaluation: an application to birth defects epidemiology. Am J Epidemiol. 155(2):176–84. doi:10.1093/aje/155.2.176 PMID:11790682

Hill AB (1965). The environment and disease: association or causation? Proc R Soc Med. 58(5):295–300. doi:10.1177/003591576505800503 PMID:14283879

Kristensen P (1992). Bias from nondifferential but dependent misclassification of exposure and outcome. Epidemiology. 3(3):210–5. doi:10.1097/00001648-199205000-00005 PMID:1591319

Lash TL (2007). Heuristic thinking and inference from observational epidemiology. Epidemiology. 18(1):67–72. doi:10.1097/01.ede.0000249522.75868.16 PMID:17149141

Lash TL, Rothman KJ (2021). Selection bias and generalizability. In: Lash TL, VanderWeele TJ, Haneuse S, Rothman KJ, editors. Modern epidemiology. 4th ed. Philadelphia (PA), USA: Wolters Kluwer; pp. 315–332.

Lash TL, VanderWeele TJ, Rothman KJ (2021). Measurement and measurement error. In: Lash TL, VanderWeele TJ, Haneuse S, Rothman KJ, editors. Modern epidemiology. 4th ed. Philadelphia (PA), USA: Wolters Kluwer; pp. 287–314.

Lipsky AM, Greenland S (2022). Causal directed acyclic graphs. JAMA. 327(11):1083–4. doi:10.1001/jama.2022.1816 PMID:35226050

Mathur MB, VanderWeele TJ (2020a). Robust metrics and sensitivity analyses for meta-analyses of heterogeneous effects. Epidemiology. 31(3):356–8. doi:10.1097/EDE.0000000000001180 PMID:32141922

Mathur MB, VanderWeele TJ (2020b). Sensitivity analysis for unmeasured confounding in meta-analyses. J Am Stat Assoc. 115(529):163–72. doi:10.1080/01621459.2018.1529598 PMID:32981992

Mathur MB, VanderWeele TJ (2022). Methods to address confounding and other biases in meta-analyses: review and recommendations. Annu Rev Public Health. 43(1):19–35. doi:10.1146/annurev-publhealth-051920-114020 PMID:34535060

Pearl J (2009). Causality: models, reasoning, and inference. 2nd ed. Cambridge, UK: Cambridge University Press. doi:10.1017/CBO9780511803161

Pearl J, Mackenzie D (2018). The book of why: the new science of cause and effect. New York (NY), USA: Basic Books.

Podsakoff PM, MacKenzie SB, Lee J-Y, Podsakoff NP (2003). Common method biases in behavioral research: a critical review of the literature and recommended remedies. J Appl Psychol. 88(5):879–903. doi:10.1037/0021-9010.88.5.879 PMID:14516251

Sheikh M, Shakeri R, Poustchi H, Pourshams A, Etemadi A, Islami F, et al. (2020). Opium use and subsequent incidence of cancer: results from the Golestan Cohort Study. Lancet Glob Health. 8(5):e649–60. doi:10.1016/S2214-109X(20)30059-0 PMID:32353313

Swanson SA (2015). Communicating causality. Eur J Epidemiol. 30(10):1073–5. doi:10.1007/s10654-015-0086-6 PMID:26445997

Swanson SA (2019). A practical guide to selection bias in instrumental variable analyses. Epidemiology. 30(3):345–9. doi:10.1097/EDE.0000000000000973 PMID:30896458

Tennant PWG, Murray EJ, Arnold KF, Berrie L, Fox MP, Gadd SC, et al. (2021). Use of directed acyclic graphs (DAGs) to identify confounders in applied health research: review and recommendations. Int J Epidemiol. 50(2):620–32. doi:10.1093/ije/dyaa213 PMID:33330936

van Smeden M, Lash TL, Groenwold RHH (2020). Reflection on modern methods: five myths about measurement error in epidemiological research. Int J Epidemiol. 49(1):338–47. doi:10.1093/ije/dyz251 PMID:31821469

Vanderweele TJ, Arah OA (2011). Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. Epidemiology. 22(1):42–52. doi:10.1097/EDE.0b013e3181f74493 PMID:21052008

VanderWeele TJ, Hernán MA (2012). Results on differential and dependent measurement error of the exposure and the outcome using signed directed acyclic graphs. Am J Epidemiol. 175(12):1303–10. doi:10.1093/aje/kwr458 PMID:22569106

VanderWeele TJ, Hernán MA, Robins JM (2008). Causal directed acyclic graphs and the direction of unmeasured confounding bias. Epidemiology. 19(5):720–8. doi:10.1097/EDE.0b013e3181810e29 PMID:18633331

VanderWeele TJ, Shpitser I (2013). On the definition of a confounder. Ann Stat. 41(1):196–220. doi:10.1214/12-AOS1058 PMID:25544784

Wright S (1960). The treatment of reciprocal interaction, with or without lag, in path analysis. Biometrics. 16(3):423. doi:10.2307/2527693

Yland JJ, Wesselink AK, Lash TL, Fox MP (2022). Misconceptions about misclassification: non-differential misclassification does not always bias results toward the null. Am J Epidemiol. doi:10.1093/aje/kwac035 PMID:35231925