

15

The Science and Ethics of Causal Modeling

Judea Pearl

University of California, Los Angeles

The research questions that motivate most quantitative studies in the health, social, and behavioral sciences are not statistical but causal in nature. For example, what is the efficacy of a given treatment or program in a given population? Whether data can prove an employer guilty of hiring discrimination? What fraction of past crimes could have been avoided by a given policy? What was the cause of death of a given individual in a specific incident? These are causal questions because they require some knowledge of the data-generating process; they cannot be computed from the data alone.

Solving causal problems mathematically requires certain extensions in the standard mathematical language of statistics, and these extensions are not generally emphasized in the mainstream literature and education. As a result, a profound tension exists between the scientific questions that a researcher wishes to ask and the type of questions traditional analysis can accommodate, let alone answer. Bluntly, scientists speak causation, and statistics delivers correlation. This tension has resulted in several ethical issues concerning the statement of a problem, the implementation of a study, and the reporting of finding. This chapter describes a simple causal extension to the language of statistics, shows how it leads to a coherent methodology that avoids the ethical problems mentioned, and permits researchers to benefit from the many results that causal analysis has produced in the past 2 decades.

Following an introductory section that defines the demarcation line between associational and causal analysis, the rest of the chapter will deal with the estimation of three types of causal queries: (a) queries about the effect of potential interventions, (b) queries about *counterfactuals* (e.g., whether event x would occur had event y been different), and (c) queries about the *direct* and *indirect effects*.

From Associational to Causal Analysis: Distinctions and Barriers

The Basic Distinction: Coping With Change

The aim of standard statistical analysis, typified by regression, estimation, and hypothesis testing techniques, is to assess parameters of a distribution from samples drawn of that distribution. With the help of such parameters, one can infer associations among variables, estimate probabilities of past and future events, and update probabilities of events in light of new evidence or new measurements. These tasks are managed well by standard statistical analysis so long as experimental conditions remain the same. Causal analysis goes one step further; its aim is to infer not only probabilities of events under static conditions but also the dynamics of events under *changing conditions*, for example, changes induced by treatments or external interventions.

This distinction implies that causal and associational concepts do not mix. There is nothing in the joint distribution of symptoms and diseases to tell us whether curing the former would or would not cure the latter. More generally, there is nothing in a distribution function to tell us how that distribution would differ if external conditions were to change—say from observational to experimental setup—because the laws of probability theory do not dictate how one property of a distribution ought to change when another property is modified. This information must be provided by causal assumptions that identify those relationships that remain invariant when external conditions change.

These considerations imply that the slogan “correlation does not imply causation” can be translated into a useful principle: One cannot substantiate causal claims from associations alone, even at the population level—behind every causal conclusion there must lie some causal assumption that is not testable in observational studies.¹

Formulating the Basic Distinction

A formal demarcation line that makes the distinction between associational and causal concepts crisp and easy to apply can be formulated as follows. An associational concept is any relationship that can be defined in terms of a joint distribution of observed variables, and a causal concept is any relationship that cannot be defined from the distribution alone. Examples of associational concepts are correlation, regression,

¹ The methodology of “causal discovery” (Pearl, 2000b, Chapter 2; Spirtes, Glymour, & Scheines, 2000) is likewise based on the causal assumption of “faithfulness” or “stability” but will not be discussed in this chapter.

dependence, conditional independence, likelihood, collapsibility, propensity score, risk ratio, odds ratio, marginalization, Granger causality, conditionalization, “controlling for,” and so on. Examples of causal concepts are randomization, influence, effect, confounding, “holding constant,” disturbance, spurious correlation, faithfulness/stability, instrumental variables, intervention, explanation, mediation, and attribution. The former can, whereas the latter cannot, be defined in terms of distribution functions.

This demarcation line is extremely useful in tracing the assumptions that are needed for substantiating various types of scientific claims. Every claim invoking causal concepts must rely on some premises that invoke such concepts; it cannot be inferred from, or even defined in, terms of statistical associations alone.

Ramifications of the Basic Distinction

This principle has far-reaching consequences that are not generally recognized in the standard statistical literature. Many researchers, for example, are still convinced that confounding is solidly founded in standard, frequentist statistics, and that it can be given an associational definition saying (roughly): “ U is a potential confounder for examining the effect of treatment X on outcome Y when both U and X and U and Y are not independent.” That this definition and all its many variants must fail (Pearl, 2009a, Section 6.2)² is obvious from the demarcation line above; if confounding were definable in terms of statistical associations, we would have been able to identify confounders from features of nonexperimental data, adjust for those confounders, and obtain unbiased estimates of causal effects. This would have violated our golden rule: Behind any causal conclusion there must be some causal assumption, untested in observational studies. Hence the definition must be false. Therefore, to the bitter disappointment of generations of epidemiologists and social science researchers, confounding bias cannot be detected or corrected by statistical methods alone; one must make some judgmental assumptions regarding causal relationships in the problem before an adjustment (e.g., by stratification) can safely correct for confounding bias.

This distinction implies that causal relations cannot be expressed in the language of probability and hence that any mathematical approach to causal analysis must acquire new notation for expressing causal relations—probability calculus is insufficient. To illustrate, the syntax of probability calculus does not permit us to express the simple fact that “symptoms do not cause diseases,” let alone draw mathematical conclusions from such facts. All

² Any intermediate variable U on a causal path from X to Y satisfies this definition, without confounding the effect of X on Y .

we can say is that two events are dependent—meaning that if we find one, we can expect to encounter the other, but we cannot distinguish statistical dependence, quantified by the conditional probability $P(\text{disease} \mid \text{symptom})$ from causal dependence, for which we have no expression in standard probability calculus. Therefore, scientists seeking to express causal relationships must supplement the language of probability with a vocabulary for causality, one in which the symbolic representation for the relation “symptoms cause disease” is distinct from the symbolic representation of “symptoms are associated with disease.”

Two Mental Barriers: Untested Assumptions and New Notation

The preceding two requirements: (a) to commence causal analysis with untested,³ theoretically or judgmentally based assumptions, and (b) to extend the syntax of probability calculus to articulate such assumptions, constitute the two main sources of confusion in the ethics of formulating, conducting, and reporting empirical studies.

Associational assumptions, even untested, are testable in principle, given a sufficiently large sample and sufficiently fine measurements. Causal assumptions, in contrast, cannot be verified even in principle, unless one resorts to experimental control. This difference stands out in Bayesian analysis. Although the priors that Bayesians commonly assign to statistical parameters are untested quantities, the sensitivity to these priors tends to diminish with increasing sample size. In contrast, sensitivity to prior causal assumptions, say that treatment does not change gender, remains substantial regardless of sample size.

This makes it doubly important that the notation we use for expressing causal assumptions be cognitively meaningful and unambiguous so that one can clearly judge the plausibility or inevitability of the assumptions articulated. Statisticians can no longer ignore the mental representation in which scientists store experiential knowledge because it is this representation and the language used to access that representation that determine the reliability of the judgments on which the analysis so crucially depends.

How does one recognize causal expressions in the statistical literature? Those versed in the potential–outcome notation (Holland, 1988; Neyman, 1923; Rubin, 1974) can recognize such expressions through the subscripts that are attached to counterfactual events and variables, for example, $Y_x(u)$ or Z_{xy} . Some authors use parenthetical expressions, for example, $Y(0)$, $Y(1)$, $Y(x, u)$, or $Z(x, y)$. The expression $Y_x(u)$, for example, may stand for the value that outcome Y would take in individual u , had treatment X been at level x . If u is chosen at random, Y_x is a random variable, and one can talk about the probability that Y_x would attain a value y in the population,

³ By “untested” I mean untested using frequency data in nonexperimental studies.

written $P(Y_x = y)$. Alternatively, Pearl (1995) used expressions of the form $P(Y = y \mid \text{set}(X = x))$ or $P(Y = y \mid \text{do}(X = x))$ to denote the probability (or frequency) that event $(Y = y)$ would occur if treatment condition $(X = x)$ were enforced uniformly over the population.⁴ Still a third notation that distinguishes causal expressions is provided by graphical models, where the arrows convey causal directionality.⁵

However, few have taken seriously the textbook requirement that any introduction of new notation must entail a systematic definition of the syntax and semantics that govern the notation. Moreover, in the bulk of the statistical literature before 2000, causal claims rarely appear in the mathematics. They surface only in the verbal interpretation that investigators occasionally attach to certain statistical parameters (e.g., regression coefficients), and in the verbal description with which investigators justify assumptions. For example, the assumption that a covariate not be affected by a treatment—a necessary assumption for the control of confounding (Cox, 1958, p. 48)—is expressed in plain English, not in a mathematical expression.

The next section provides a conceptualization that overcomes these mental barriers; it offers both a friendly mathematical machinery for cause–effect analysis and a formal foundation for counterfactual analysis.

Structural Causal Models, Diagrams, Causal Effects, and Counterfactuals

Structural Equations as Oracles for Causes and Counterfactuals

How can one express mathematically the common understanding that symptoms do not cause diseases? The earliest attempt to formulate such a relationship mathematically was made in the 1920s by the geneticist Sewall Wright (1921), who used a combination of equations and graphs. For example, if X stands for a disease variable and Y stands for a certain symptom of the disease, Wright would write a linear equation:

$$y = \beta x + u \quad (15.1)$$

⁴ Clearly, $P(Y = y \mid \text{do}(X = x))$ is equivalent to $P(Y_x = y)$. This is what we normally assess in a controlled experiment, with X randomized, in which the distribution of Y is estimated for each level x of X .

⁵ These notational clues should be useful for detecting inadequate definitions of causal concepts; any definition of confounding, randomization, or instrumental variables that is cast in standard probability expressions, void of graphs, counterfactual subscripts, or $\text{do}(\ast)$ operators, can safely be discarded as inadequate.

where x stands for the level (or severity) of the disease, y stands for the level (or severity) of the symptom, and u stands for all factors, other than the disease in question, that could possibly affect Y . In interpreting this equation one should think of a physical process whereby Nature *examines* the values of x and u and, accordingly, *assigns* variable Y the value $y = \beta x + u$.

To express the directionality inherent in this assignment process, Wright augmented the equation with a diagram, later called "path diagram," in which arrows are drawn from (perceived) causes to their (perceived) effects, and, more importantly, the absence of an arrow makes the empirical claim that the value Nature assigns to one variable is indifferent to that taken by another (see Figure 15.1).

The variables V and U are called "exogenous"; they represent observed or unobserved background factors that the modeler decides to keep unexplained, that is, factors that influence but are not influenced by the other variables (called "endogenous") in the model.

If correlation is judged possible between two exogenous variables, U and V , it is customary to connect them by a dashed double arrow, as shown in Figure 15.1b.

To summarize, path diagrams encode causal assumptions via missing arrows, representing claims of zero influence, and missing double arrows (e.g., between V and U), representing the (causal) assumption $Cov(U, V) = 0$.

The generalization to nonlinear system of equations is straightforward. For example, the nonparametric interpretation of the diagram of Figure 15.2a corresponds to a set of three functions, each corresponding to one of the observed variables:

$$\begin{aligned} z &= f_z(w) \\ x &= f_x(z, v) \\ y &= f_y(x, u) \end{aligned} \quad (15.2)$$

where in this particular example, W , V , and U are assumed to be jointly independent but, otherwise, arbitrarily distributed.




FIGURE 15.1

A simple structural equation model and its associated diagrams. Unobserved exogenous variables are connected by dashed arrows.




FIGURE 15.2

(a) The diagram associated with the structural model of Equation 15.2. (b) The diagram associated with the modified model, M_{x_0} , of Equation 15.3, representing the intervention $do(X = x_0)$.

Remarkably, unknown to most economists and philosophers,⁶ *structural equation models* provide a formal interpretation and symbolic machinery for analyzing counterfactual relationships of the type: “Y would be y had X been x in situation $\mathbf{U} = \mathbf{u}$,” denoted $Y_x(\mathbf{u}) = y$. Here \mathbf{U} represents the vector of all exogenous variables.⁷

The key idea is to interpret the phrase “had X been x_0 ” as an instruction to modify the original model and replace the equation for X by a constant x_0 , yielding the submodel, M_{x_0} ,

$$\begin{aligned} z &= f_z(w) \\ x &= x_0 \\ y &= f_y(x, u) \end{aligned} \tag{15.3}$$

the graphical description of which is shown in Figure 15.2b.

This replacement permits the constant x_0 to differ from the actual value of X, namely, $f_x(z, v)$, without rendering the system of equations inconsistent, thus yielding a formal interpretation of counterfactuals in multistage models, where the dependent variable in one equation may be an independent variable in another (Balke & Pearl, 1994; Pearl, 2000). In general, we can formally define the postintervention distribution by the equation:

$$P_M(y \mid do(x)) \triangleq P_{M_x}(y). \tag{15.4}$$

In words: In the framework of model M , the postintervention distribution of outcome Y is defined as the probability that model M_x assigns to each outcome level $Y = y$.

⁶ Connections between structural equations and a restricted class of counterfactuals were recognized by Simon and Rescher (1966). These were later generalized by Balke and Pearl (1995), who used modified models to permit counterfactual conditioning on dependent variables.

⁷ Because $\mathbf{U} = \mathbf{u}$ may contain detailed information about a situation or an individual, $Y_x(\mathbf{u})$ is related to what philosophers called “token causation,” whereas $P(Y_x = y \mid Z = z)$ characterizes “Type causation,” that is, the tendency of X to influence Y in a subpopulation characterized by $Z = z$.

From this distribution, one is able to assess treatment efficacy by comparing aspects of this distribution at different levels of x_0 . A common measure of treatment efficacy is the difference

$$E(Y \mid do(x'_0)) - E(Y \mid do(x_0)) \quad (15.5)$$

where x'_0 and x_0 are two levels (or types) of treatment selected for comparison. For example, to compute $E(Y_{x_0})$, the expected effect of setting X to x_0 (also called the average causal effect of X on Y , denoted $E(Y \mid do(x_0))$ or, generically, $E(Y \mid do(x_0))$), we solve Equation 15.3 for Y in terms of the exogenous variables, yielding $Y_{x_0} = f_Y(x_0, u)$, and average over U and V . It is easy to show that in this simple system, the answer can be obtained without knowing the form of the function $f_Y(x, u)$ or the distribution $P(u)$. The answer is given by:

$$E(Y_{x_0}) = E(Y \mid do(X = x_0)) = E(Y \mid x_0)$$

which is estimable from the observed distribution $P(x, y, z)$. This result hinges on the assumption that W , V , and U are mutually independent and on the topology of the graph (e.g., that there is no direct arrow from Z to Y).

In general, it can be shown (Pearl, 2009a, Chapter 3) that, whenever the graph is Markovian (i.e., acyclic with independent exogenous variables) the postinterventional distribution $P(Y = y \mid do(X = x))$ is given by the following expression:

$$P(Y = y \mid do(X = x)) = \sum_t P(y \mid t, x)P(t) \quad (15.6)$$

where T is the set of direct causes of X (also called “parents”) in the graph. Again, we see that all factors on the right side are estimable from the distribution P of observed variables, and hence the counterfactual probability $P(Y_x = y)$ is estimable with mere partial knowledge of the generating process—the topology of the graph and independence of the exogenous variable are all that is needed.

When some variables in the graph (e.g., the parents of X) are unobserved, we may not be able to estimate (or “identify” as it is called) the postintervention distribution $P(y \mid do(x))$ by simple conditioning, and more sophisticated methods would be required. Likewise, when the query of interest involves several hypothetical worlds simultaneously, for example, $P(Y_x = y \mid Y_{x'} = y')$, the Markovian assumption may not suffice for identification and additional assumptions, touching on the form of the data-generating functions (e.g., monotonicity) may need to be invoked. These issues will be discussed in the “Confounding and Causal Effect Estimation” and “An Example: Mediation, Direct and Indirect Effects” sections.

This interpretation of counterfactuals, cast as solutions to modified systems of equations, provides the conceptual and formal link between structural equation models, used in economics and social science, and the Neyman–Rubin potential–outcome framework to be discussed in the “Relation to Potential Outcomes and the Demystification of ‘Ignorability’” section. But first we discuss two longstanding problems that have been completely resolved in purely graphical terms, without delving into algebraic techniques.

Confounding and Causal Effect Estimation

Although good statisticians have always known that the elucidation of causal relationships from observational studies must be shaped by assumptions about how the data were generated, the relative roles of assumptions and data, and ways of using those assumptions to eliminate confounding bias, have been a subject of much controversy.⁸ The structural framework of the “Structural Equations as Oracles for Causes and Counterfactual” section puts these controversies to rest.

Covariate Selection: The Back-Door Criterion

Consider an observational study where we wish to find the effect of X on Y , for example, treatment on response, and assume that the factors deemed relevant to the problem are structured as in Figure 15.3—some are affecting the response; some are affecting the treatment; and some are affecting both treatment and response.

Some of these factors may be unmeasurable, such as genetic trait or lifestyle; others are measurable, such as gender, age, and salary level. Our problem is to select a subset of these factors for measurement and adjustment, namely, that if we compare treated versus untreated subjects having the same values of the selected factors, we get the correct treatment effect in that subpopulation of subjects. Such a set of “deconfounding” factors is called a “sufficient set” or a set “admissible for adjustment.” The problem of defining a sufficient set, let alone finding one, has baffled epidemiologists and social scientists for decades (see Greenland, Pearl, & Robins, 1999; Pearl, 1998, 2003, for review).

⁸ A recent flare-up of this controversy can be found in Pearl (2009c, 2009d, 2010c) and Rubin (2009), which demonstrate the difficulties statisticians encounter in articulating causal assumptions and typical mistakes that arise from pursuing causal analysis within the structure-less “missing data” paradigm.




FIGURE 15.3

Graphical model illustrating the back-door criterion. Error terms are not shown explicitly.

The following criterion, named “back-door” in Pearl (1993), settles this problem by providing a graphical method of selecting a sufficient set of factors for adjustment. It states that a set S is admissible for adjustment if two conditions hold:

1. No element of S is a descendant of X .
2. The elements of S “block” all “back-door” paths from X to Y , namely, all paths that end with an arrow pointing to X .⁹

Based on this criterion we see, for example, that in the sets $\{Z_1, Z_2, Z_3\}$, $\{Z_1, Z_3\}$, and $\{W_2, Z_3\}$, each is sufficient for adjustment because each blocks all back-door paths between X and Y . The set $\{Z_3\}$, however, is not sufficient for adjustment because, as explained in footnote 3, it does not block the path $X \leftarrow W_1 \leftarrow Z_1 \rightarrow Z_3 \leftarrow Z_2 \rightarrow W_2 \rightarrow Y$.

The implication of finding a sufficient set S is that stratifying on S is guaranteed to remove all confounding bias relative to the causal effect of X on Y . In other words, it renders the causal effect of X on Y estimable, via

$$\begin{aligned}
 &P(Y = y \mid do(X = x)) \\
 &= \sum_s P(Y = y \mid X = x, S = s)P(S = s).
 \end{aligned}
 \tag{15.7}$$

Because all factors on the right side of the equation are estimable (e.g., by regression) from the preinterventional data, the causal effect can likewise be estimated from such data without bias.

The back-door criterion allows us to write Equation 15.7 directly, after selecting a sufficient set S from the diagram, without resorting to any algebraic manipulation. The selection criterion can be applied systematically to diagrams of any size and shape, thus freeing analysts from

⁹ In this criterion, a set S of nodes is said to *block* a path p if either (a) p contains at least one arrow-emitting node that is in S , or (b) p contains at least one collision node that is outside S and has no descendant in S . (See Pearl, 2000b, pp. 16–17, 335–337.)

judging whether “ X is conditionally ignorable given S ,” a formidable mental task required in the potential–outcome framework (Rosenbaum & Rubin, 1983). The criterion also enables the analyst to search for an optimal set of covariates—namely, a set S that minimizes measurement cost or sampling variability (Tian, Paz, & Pearl, 1998). A complete identification condition, including models with no sufficient sets (e.g., Figure 15.3, assuming that X , Y , and W_3 are the only measured variables) is given in Shpitser and Pearl (2006).

Another problem that has a simple graphical solution is to determine whether adjustment for two sets of covariates would result in the same confounding bias (Pearl & Paz, 2010). This criterion allows one to assess, before taking any measurement, whether two candidate sets of covariates, differing substantially in dimensionality, measurement error, cost, or sample variability, are equally valuable in their bias-reduction potential.

Counterfactual Analysis in Structural Models

Not all questions of causal character can be encoded in $P(y \mid do(x))$ -type expressions, in much the same way that not all causal questions can be answered from experimental studies. For example, questions of attribution (e.g., I took an aspirin and my headache is gone; was it *due* to the aspirin?) or of susceptibility (e.g., I am a healthy nonsmoker; would I be as healthy had I been a smoker?) cannot be answered from experimental studies, and naturally, these kind of questions cannot be expressed in $P(y \mid do(x))$ notation.¹⁰ To answer such questions, a probabilistic analysis of counterfactuals is required, one dedicated to the relation “ Y would be y had X been x in situation $\mathbf{U} = \mathbf{u}$,” denoted $Y_x(\mathbf{u}) = y$.

As noted in the “Structural Equations as Oracles for Causes and Counterfactuals” section, the structural definition of counterfactuals involves modified models, like M_{x_0} of Equation 15.3, formed by the intervention $do(X = x_0)$ (Figure 15.2b). Denote the solution of Y in model M_x by the symbol $Y_{M_x}(\mathbf{u})$; the formal definition of the counterfactual $Y_x(\mathbf{u})$ in a structural causal model is given by (Blake & Pearl, 1994; Pearl, 2009a, p. 98):

$$Y_x(\mathbf{u}) \triangleq Y_{M_x}(\mathbf{u}). \quad (15.8)$$

¹⁰The reason for this fundamental limitation is that no death case can be tested twice, with and without treatment. For example, if we measure equal proportions of deaths in the treatment and control groups, we cannot tell how many death cases are attributable to the treatment itself; it is possible that many of those who died under treatment would be alive if untreated and, simultaneously, many of those who survived with treatment would have died if not treated.

The quantity $Y_x(\mathbf{u})$ can be given experimental interpretation; it stands for the way an individual with characteristics (\mathbf{u}) would respond had the treatment been x , rather than the treatment $x = f_x(\mathbf{u})$ received by that individual. In our example, because Y does not depend on v and w , we can write: $Y_{x_0}(\mathbf{u}) = f_Y(x_0, \mathbf{u})$. Clearly, the distribution $P(u, v, w)$ induces a well-defined probability on the counterfactual event $Y_{x_0} = y$, as well as on joint counterfactual events, such as " $Y_{x_0} = y$ AND $Y_{x_1} = y'$," which are, in principle, unobservable if $x_0 \neq x_1$. Thus, to answer attributional questions, such as whether Y would be y_1 if X were x_1 , given that in fact Y is y_0 and X is x_0 , we need to compute the conditional probability $P(Y_{x_1} = y_1 \mid Y = y_0, X = x_0)$, which is well defined once we know the forms of the structural equations and the distribution of the exogenous variables in the model. For example, assuming a linear equation for Y (as in Figure 15.1),

$$y = \beta x + u,$$

the conditions $Y = y_0$ and $X = x_0$ yield $U = y_0 - \beta x_0$, and we can conclude that, with probability one, Y_{x_1} must take on the value: $Y_{x_1} = \beta x_1 + U = \beta(x_1 - x_0) + y_0$. In other words, if X were x_1 instead of x_0 , Y would increase by β times the difference $(x_1 - x_0)$. In nonlinear systems, the result would also depend on the distribution of U , and, for that reason, attributional queries are generally not identifiable in nonparametric models (Pearl, 2009a, Chapter 9).

In general, if x and x' are incompatible, then Y_x and $Y_{x'}$ cannot be measured simultaneously, and it may seem meaningless to attribute probability to the joint statement " Y would be y if $X = x$ and Y would be y' if $X = x'$." Such concerns have been a source of objections to treating counterfactuals as jointly distributed random variables (Dawid, 2000). The definition of Y_x and $Y_{x'}$ in terms of two distinct submodels neutralizes these objections (Pearl, 2009a, p. 206) because the contradictory joint statement is mapped into an ordinary event (among the background variables) that satisfies both statements simultaneously, each in its own distinct submodel; such events have well-defined probabilities.

The structural interpretation of counterfactuals (Equation 15.8) also provides the conceptual and formal basis for the Neyman–Rubin potential–outcome framework, an approach that takes a controlled randomized trial as its starting paradigm, assuming that nothing is known to the experimenter about the science behind the data. This "black box" approach was developed by statisticians who found it difficult to cross the two mental barriers discussed in the "Two Mental Barriers" section. The next section establishes the precise relationship between the structural and potential–outcome paradigms, and outlines how the latter can benefit from the richer representational power of the former.

Relation to Potential Outcomes and the Demystification of “Ignorability”

The primitive object of analysis in the potential–outcome framework is the unit-based response variable, denoted $Y_x(u)$, read: “the value that outcome Y would obtain in experimental unit u , had treatment X been x ” (Neyman, 1923; Rubin, 1974). Here, *unit* may stand for an individual patient, an experimental subject, or an agricultural plot. In the “Counterfactual Analysis in Structural Models” section, we saw (Equation 15.8) that this counterfactual entity has a natural interpretation in structural equations as the solution for Y in a modified system of equation, where *unit* is interpreted as vector u of background factors that characterize an experimental unit. Thus, each structural equation model carries a collection of assumptions about the behavior of hypothetical units, and these assumptions permit us to derive the counterfactual quantities of interest. In the potential–outcome framework, however, no equations are available for guidance, and $Y_x(u)$ is taken as primitive, that is, an undefined quantity in terms of which other quantities are defined—not a quantity that can be derived from some model. In this sense, the structural interpretation of $Y_x(u)$ provides the formal basis for the *potential outcome* approach; the formation of the submodel M_x explicates mathematically how the hypothetical condition “had X been x ” could be realized and what the logical consequences are of such a condition.

The distinct characteristic of the potential outcome approach is that, although investigators must think and communicate in terms of undefined, hypothetical quantities such as $Y_x(u)$, the analysis itself is conducted almost entirely within the axiomatic framework of probability theory. This is accomplished by treating the new hypothetical entities Y_x as ordinary random variables; for example, they are assumed to obey the axioms of probability calculus, the laws of conditioning, and the axioms of conditional independence.

Naturally, these hypothetical entities are not entirely whimsy. They are assumed to be connected to observed variables via consistency constraints (Robins, 1986), such as

$$X = x \Rightarrow Y_x = Y, \quad (15.9)$$

which states that, for every u , if the actual value of X turns out to be x , then the value that Y would take on if “ X were x ” is equal to the actual value of Y . For example, a person who chose treatment x and recovered would also have recovered if given treatment x by design. Whether additional constraints should tie the observables to the unobservables is not a question that can be answered in the potential–outcome framework, which lacks an underlying model.

The main conceptual difference between the two approaches is that whereas the structural approach views the intervention $do(x)$ as an operation that changes the distribution but keeps the variables the same, the potential–outcome approach views the variable Y under $do(x)$ to be a different variable, Y_x , loosely connected to Y through relations such as Equation 15.9, but remaining unobserved whenever $X \neq x$. The problem of inferring probabilistic properties of Y_x then becomes one of “missing data,” for which estimation techniques have been developed in the statistical literature.

Pearl (2009a, Chapter 7) shows, using the structural interpretation of $Y_x(u)$ (Equation 15.8), that it is indeed legitimate to treat counterfactuals as jointly distributed random variables in all respects, that consistency constraints like Equation 15.9 are automatically satisfied in the structural interpretation, and, moreover, that investigators need not be concerned about any additional constraints except the following two:

$$Y_{yz} = y \text{ for all } y, \text{ subsets } Z, \text{ and values } z \text{ for } Z \quad (15.10)$$

$$X_z = x \Rightarrow Y_{xz} = Y_z \text{ for all } x, \text{ subsets } Z, \text{ and values } z \text{ for } Z. \quad (15.11)$$

Equation 15.10 ensures that the intervention $do(Y = y)$ results in the condition $Y = y$, regardless of concurrent interventions, say $do(Z = z)$, that may be applied to variables other than Y . Equation 15.11 generalizes Equation 15.9 to cases where Z is held fixed, at z .

Problem Formulation and the Demystification of “Ignorability”

The main drawback of this black box approach surfaces in the phase where a researcher begins to articulate the “science” or “causal assumptions” behind the problem at hand. Such knowledge, as we have seen in the “Two Mental Barriers” section, must be articulated at the onset of every problem in causal analysis—causal conclusions are only as valid as the causal assumptions on which they rest.

To communicate scientific knowledge, the potential–outcome analyst must express causal assumptions in the form of assertions involving counterfactual variables. For example, in our example of Figure 15.2a, to communicate the understanding that Z is randomized (hence independent of V and U), the potential–outcome analyst would use the independence constraint $Z \perp\!\!\!\perp \{X_z, Y_x\}$.¹¹ To further formulate the understanding that Z

¹¹ The notation $Y \perp\!\!\!\perp X \mid Z$ stands for the conditional independence relationship $P(Y = y, X = x \mid Z = z) = P(Y = y \mid Z = z) P(X = x \mid Z = z)$ (Dawid, 1979).

does not affect Y directly, except through X , the analyst would write a so-called, “exclusion restriction”: $Y_{xz} = Y_x$.

A collection of constraints of this type might sometimes be sufficient to permit a unique solution to the query of interest; in other cases, only bounds on the solution can be obtained. For example, if one can plausibly assume that a set Z of covariates satisfies the relation:

$$Y_x \perp\!\!\!\perp X \mid Z \quad (15.12)$$

(assumption that was termed *conditional ignorability* by Rosenbaum & Rubin, 1983), then the causal effect $P(Y_x = y)$ can readily be evaluated to yield:

$$\begin{aligned} P(Y_x = y) &= \sum_z P(Y_x = y \mid z)P(z) \\ &= \sum_z P(Y_x = y \mid x, y)P(z) \quad (\text{using (12)}) \\ &= \sum_z P(Y_x = y \mid x, z)P(z) \quad (\text{using (9)}) \\ &= \sum_z P(y \mid x, z)P(z). \end{aligned} \quad (15.13)$$

The last expression contains no counterfactual quantities and coincides precisely with the standard covariate-adjustment formula of Equation 15.7.

We see that the assumption of conditional ignorability (Equation 15.12) qualifies Z as a sufficient covariate for adjustment; indeed, one can show formally (Pearl, 2009a, pp. 98–102, 341–343) that Equation 15.12 is entailed by the “back-door” criterion of the “Confounding and Causal Effect Estimation” section.

The derivation above may explain why the potential outcome approach appeals to conservative statisticians; instead of constructing new vocabulary (e.g., arrows), new operators (*do(x)*), and new logic for causal analysis, almost all mathematical operations in this framework are conducted within the safe confines of probability calculus. Save for an occasional application of the consistency rule, Equation 15.11 or Equation 15.9, the analyst may forget that Y_x stands for a counterfactual quantity—it is treated as any other random variable, and the entire derivation follows the course of routine probability exercises.

However, this mathematical orthodoxy exacts a high cost at the critical stage where causal assumptions are formulated. The reader may appreciate this aspect by attempting to judge whether the assumption of conditional ignorability (Equation 15.12), the key to the derivation of Equation 15.15, holds in any familiar situation, say, in the experimental setup of

Figure 15.2a. This assumption reads: “the value that Y would obtain had X been x , is independent of X , given Z .” Even the most experienced potential–outcome expert would be unable to discern whether any subset Z of covariates in Figure 15.3 would satisfy this conditional independence condition.¹² Likewise, to convey the structure of the chain $X \rightarrow W_3 \rightarrow Y$ (Figure 15.3) in the language of potential–outcome, one would need to write the cryptic expression: $W_{3_x} \perp\!\!\!\perp \{Y_{w_3}, X\}$, read: “the value that W_3 would obtain had X been x is independent of the value that Y would obtain had W_3 been w_3 jointly with the value of X .” Such assumptions are cast in a language so far removed from ordinary understanding of cause and effect that, for all practical purposes, they cannot be comprehended or ascertained by ordinary mortals. As a result, researchers in the graphless potential–outcome camp rarely use “conditional ignorability” (Equation 15.12) to guide the choice of covariates; they view this condition as a hoped-for miracle of nature rather than a target to be achieved by reasoned design.¹³

Having translated “ignorability” into a simple condition (i.e., back door) in a graphical model permits researchers to understand what conditions covariates must fulfill before they eliminate bias, what to watch for and what to think about when covariates are selected, and what experiments we can do to test, at least partially, if we have the knowledge needed for covariate selection.

Aside from offering no guidance in covariate selection, formulating a problem in the potential–outcome language encounters three additional hurdles. When counterfactual variables are not viewed as byproducts of a deeper, process-based model, it is hard to ascertain whether *all* relevant counterfactual independence judgments have been articulated, whether the judgments articulated are *redundant*, or whether those judgments are *self-consistent*. The need to express, defend, and manage formidable counterfactual relationships of this type explains the slow acceptance of causal analysis among health scientists and statisticians, and why economists and social scientists continue to use structural equation models instead of the potential–outcome alternatives advocated in Angrist, Imbens, and Rubin (1996), Holland (1988), and Sobel (1998).

On the other hand, the algebraic machinery offered by the counterfactual notation, $Y_x(u)$, once a problem is properly formulated, can be

¹²Inquisitive readers are invited to guess whether $X_z \perp\!\!\!\perp Z \mid Y$ holds in Figure 15.2a.

¹³The opaqueness of counterfactual independencies explains why many researchers within the potential–outcome camp are unaware of the fact that adding a covariate to the analysis (e.g., Z_3 in Figure 15.3) may increase confounding bias. Paul Rosenbaum, for example, writes: “There is little or no reason to avoid adjustment for a variable describing subjects before treatment” (Rosenbaum, 2002, p. 76). Rubin (2009) goes as far as stating that refraining from conditioning on an available measurement is “nonscientific ad hocery” because it goes against the tenets of Bayesian philosophy. (See Pearl, 2009c, 2009d, 2010c, for a discussion of this fallacy.)

extremely powerful in refining assumptions (Angrist et al., 1996), deriving consistent estimands (Robins, 1986), bounding probabilities of necessary and sufficient causation (Tian & Pearl, 2000), and combining data from experimental and nonexperimental studies (Pearl, 2009a). Pearl (2009a, p. 232) presents a way of combining the best features of the two approaches. It is based on encoding causal assumptions in the language of diagrams, translating these assumptions into counterfactual notation, performing the mathematics in the algebraic language of counterfactuals (using Equations 15.9, 15.10, and 15.11), and, finally, interpreting the result in plain causal language. The “An Example: Mediation, Direct and Indirect Effects” section illustrates such symbiosis.

Methodological Dictates and Ethical Considerations

The structural theory described in the previous sections dictates a principled methodology that eliminates the confusion between causal and statistical interpretations of study results, as well as the ethical dilemmas that this confusion tends to spawn. The methodology dictates that every investigation involving causal relationships (and this entails the vast majority of empirical studies in the social and behavioral sciences) should be structured along the following four-step process:¹⁴

1. **Define:** Express the target quantity Q as a function $Q(M)$ that can be computed from any model M , regardless of how realistic it is.
2. **Assume:** Formulate causal assumptions using ordinary scientific language, and represent their structural part in graphical form.
3. **Identify:** Determine whether the target quantity is identifiable (i.e., expressible as distributions).
4. **Estimate:** Estimate the target quantity if it is identifiable, or approximate it if it is not.

Defining the Target Quantity

The definitional phase is the most neglected step in current practice of quantitative analysis. The structural modeling approach insists on defining the target quantity, be it “causal effect,” “program effectiveness,” “mediated effect,” “effect on the treated,” or “probability of causation” before specifying any aspect of the model, without making functional or

¹⁴Pearl (2010a) identifies five steps, which include model testing.

distributional assumptions, before choosing a method of estimation, and before seeing any data.

The investigator should view this definition as an *algorithm* that receives a model M as an input and delivers the desired quantity $Q(M)$ as the output. Surely, such an algorithm should not be tailored to any aspect of the input M ; it should be general and ready to accommodate any conceivable model M whatsoever. Moreover, the investigator should imagine that the input M is a completely specified model, with all the functions f_X, f_Y, \dots and all the U variables (or their associated probabilities) given precisely. This is the hardest step for statistically trained investigators to make; knowing in advance that such model details will never be estimable from the data, the definition of $Q(M)$ appears like a futile exercise in fantasyland—it is not.

For example, the formal definition of the causal effect $P(y \mid do(x))$, as given in Equation 15.4, is universally applicable to all models, parametric and nonparametric, through the formation of a submodel M_x . By defining causal effect procedurally, thus divorcing it from its traditional parametric representation, the structural theory avoids the many pitfalls and confusions that have plagued the interpretation of structural and regressional parameters for the past half century.¹⁵

Explicating Causal Assumptions

This is the second most neglected step in causal analysis. In the past, the difficulty has been the lack of language suitable for articulating causal assumptions, which, aside from impeding investigators from explicating assumptions, also inhibited them from giving causal interpretations to their findings.

Structural equation models, in their counterfactual reading, have settled this difficulty. Today we understand that the versatility and natural appeal of structural equations stem from the fact that they permit investigators to communicate causal assumptions formally and in the very same vocabulary that scientific knowledge is stored.

Unfortunately, however, this understanding is not shared by all causal analysts; some analysts vehemently resist the resurrection of structural models and insist instead on articulating causal assumptions exclusively

¹⁵Note that b in Equation 15.1, the incremental causal effect of X on Y , is defined procedurally by

$$\beta = E(Y \mid do(x_0 + 1)) - E(Y \mid do(x_0)) = \frac{\partial}{\partial x} E(Y_x).$$

Naturally, all attempts to give b statistical interpretation have ended in frustrations (Holland, 1988; Wermuth, 1992; Wermuth & Cox, 1993; Whittaker, 1990), some persisting well into the 21st century (Sobel, 2008).

in the unnatural (although formally equivalent) language of potential outcomes, ignorability, treatment assignment, and other metaphors borrowed from clinical trials. This assault on structural modeling is perhaps more dangerous than the causal-associational confusion because it is riding on a halo of exclusive ownership to scientific principles and, while welcoming causation, uproots it away from its natural habitat.

Early birds of this exclusivist attitude have already infiltrated the American Psychological Association's (APA) guidelines (Wilkinson & the Task Force on Statistical Inference, 1999), where we can read passages such as: "The crucial idea is to set up the causal inference problem as one of missing data," (item 72) or "If a problem of causal inference cannot be formulated in this manner (as the comparison of potential outcomes under different treatment assignments), it is not a problem of inference for causal effects, and the use of 'causal' should be avoided," (item 73) or, even more bluntly, "The underlying assumptions needed to justify any causal conclusions should be carefully and explicitly argued, not in terms of technical properties like 'uncorrelated error terms,' but in terms of real world properties, such as how the units received the different treatments" (item 74).

The methodology expounded in this article testifies against such restrictions. It demonstrates a viable and principled formalism based on traditional structural equations paradigm, which stands diametrically opposed to the "missing data" paradigm. It renders the vocabulary of "treatment assignment" stifling and irrelevant (e.g., there is no "treatment assignment" in sex discrimination cases). Most importantly, it strongly prefers the use of "uncorrelated error terms" (or "omitted factors") over its "strong ignorability" alternative, which even experts admit cannot be used (and has not been used) to reason about underlying assumptions.

In short, the APA's guidelines should be vastly more inclusive and borrow strength from multiple approaches. The next section demonstrates the benefit of a symbiotic, graphical-structural-counterfactual approach to deal with the problem of mediation, or effect decomposition.

An Example: Mediation, Direct and Indirect Effects

Direct Versus Total Effects

The causal effect we have analyzed so far, $P(y \mid do(x))$, measures the *total* effect of a variable (or a set of variables) X on a response variable Y . In many cases, this quantity does not adequately represent the target of

investigation, and attention is focused instead on the direct effect of X on Y . The term *direct effect* is meant to quantify an effect that is not mediated by other variables in the model or, more accurately, the sensitivity of Y to changes in X while all other factors in the analysis are held fixed. Naturally, holding those factors fixed would sever all causal paths from X to Y with the exception of the direct link $X \rightarrow Y$, which is not intercepted by any intermediaries.

A classical example of the ubiquity of direct effects involves legal disputes over race or sex discrimination in hiring. Here, neither the effect of sex or race on applicants' qualification nor the effect of qualification on hiring is a target of litigation. Rather, defendants must prove that sex and race do not *directly* influence hiring decisions, whatever indirect effects they might have on hiring by way of applicant qualification.

From a policy-making viewpoint, an investigator may be interested in decomposing effects to quantify the extent to which racial salary disparity is the result of educational disparity, or, taking a health care example, the extent to which sensitivity to a given exposure can be reduced by eliminating sensitivity to an intermediate factor, standing between exposure and outcome. Another example concerns the identification of neural pathways in the brain or the structural features of protein-signaling networks in molecular biology (Brent & Lok, 2005). Here, the decomposition of effects into their direct and indirect components carries theoretical scientific importance because it tells us "how nature works" and therefore enables us to predict behavior under a rich variety of conditions.

Yet despite its ubiquity, the analysis of mediation has long been a thorny issue in the social and behavioral sciences (Baron & Kenny, 1986; Judd & Kenny, 1981; MacKinnon, Fairchild, & Fritz, 2007a; Muller, Judd, & Yzerbyt, 2005; Shrout & Bolger, 2002) primarily because structural equation modeling in those sciences was deeply entrenched in linear analysis, where the distinction between causal parameters and their regressional interpretations can easily be conflated. As demands grew to tackle problems involving binary and categorical variables, researchers could no longer define direct and indirect effects in terms of structural or regressional coefficients, and all attempts to extend the linear paradigms of effect decomposition to nonlinear systems produced distorted results (MacKinnon, Lockwood, Brown, Wang, & Hoffman, 2007b). These difficulties have accentuated the need to redefine and derive *causal effects* from first principles, uncommitted to distributional assumptions, or a particular parametric form of the equations. The structural methodology presented in this chapter adheres to this philosophy, and it has produced indeed a principled solution to the mediation problem, based on the counterfactual reading of structural equations (Equation 15.8). The following subsections summarize the method and its solution.

Controlled Direct Effects

A major impediment to progress in mediation analysis has been the lack of notational facility for expressing the key notion of “holding the mediating variables fixed” in the definition of direct effect. Clearly, this notion must be interpreted as (hypothetically) setting the intermediate variables to constants by physical intervention, not by analytical means such as selection, regression conditioning, matching, or adjustment. For example, consider the simple mediation models of Figure 15.4, where the error terms (not shown explicitly) are assumed to be independent. It will not be sufficient to measure the association between gender (X) and hiring (Y) for a given level of qualification (Z) (see Figure 15.4b) because, by conditioning on the mediator Z , we create spurious associations between X and Y through W_2 , even when there is no direct effect of X on Y (Pearl, 1998).

Using the $do(x)$ notation enables us to correctly express the notion of “holding Z fixed” and formulate a simple definition of the *controlled direct effect* (CDE) of the transition from $X = x$ to $X = x'$:

$$CDE \triangleq E(Y \mid do(x'), do(z)) - E(Y \mid do(x), do(z)).$$

Or, equivalently, using counterfactual notation:

$$CDE \triangleq E(Y_{x'z}) - E(Y_{xz})$$

where Z is the set of all mediating variables. The readers can easily verify that, in linear systems, the controlled direct effect reduces to the path coefficient of the link $X \rightarrow Y$ (see footnote 14) regardless of whether confounders are present (as in Figure 15.4b) and regardless of whether the error terms are correlated.

This separates the task of definition from that of identification, as demanded by the “Defining the Target Quantity” section. The identification




FIGURE 15.4 (a) A generic model depicting mediation through Z with no confounders. (b) A mediation model with two confounders, W_1 and W_2 .

of CDE would depend, of course, on whether confounders are present and whether they can be neutralized by adjustment, but these do not alter its definition. Graphical identification conditions for expressions of the type $E(Y \mid do(x), do(z_1), do(z_2), \dots, do(z_k))$ in the presence of unmeasured confounders were derived by Pearl and Robins (1995) (see Pearl, 2009a, Chapter 4) and invoke sequential application of the back-door conditions discussed in the “Confounding and Causal Effect Estimation” section.

Natural Direct Effects

In linear systems, the direct effect is fully specified by the path coefficient attached to the link from X to Y ; therefore, the direct effect is independent of the values at which we hold Z . In nonlinear systems, those values would, in general, modify the effect of X on Y and thus should be chosen carefully to represent the target policy under analysis. For example, it is not uncommon to find employers who prefer males for the high-paying jobs (i.e., high z) and females for low-paying jobs (low z).

When the direct effect is sensitive to the levels at which we hold Z , it is often more meaningful to define the direct effect relative to some “natural” baseline level that may vary from individual to individual and represents the level of Z just before the change in X . Conceptually, we can define the natural direct effect $DE_{x,x'}(Y)$ as the expected change in Y induced by changing X from x to x' while keeping all mediating factors constant at whatever value they *would have obtained* under $do(x)$. This hypothetical change, which Robins and Greenland (1992) conceived and called “pure” and Pearl (2001) formalized and analyzed under the rubric “natural,” mirrors what lawmakers instruct us to consider in race or sex discrimination cases: “The central question in any employment-discrimination case is whether the employer would have taken the same action had the employee been of a different race (age, sex, religion, national origin, etc.) and everything else had been the same” (*Carson v. Bethlehem Steel Corp.*, 1996).

Extending the subscript notation to express nested counterfactuals, Pearl (2001) gave a formal definition for the “natural direct effect”:

$$DE_{x,x'}(Y) = E(Y_{x',z_x}) - E(Y_x). \quad (15.14)$$

Here Y_{x',z_x} represents the value that Y would attain under the operation of setting X to x' and simultaneously setting Z to whatever value it would have obtained under the setting $X = x$. We see that $DE_{x,x'}(Y)$, the natural direct effect of the transition from x to x' , involves probabilities of *nested counterfactuals* and cannot be written in terms of the $do(x)$ operator. Therefore, the natural direct effect cannot in general be identified or estimated, even with the help of ideal, controlled experiments (see footnote 10)—a point emphasized in Robins and Greenland (1992). However, aided by Equation

15.8 and the notational power of nested counterfactuals, Pearl (2001) was nevertheless able to show that if certain assumptions of “no confounding” are deemed valid the natural direct effect can be reduced to

$$DE_{x,x'}(Y) = \sum_z [E(Y | do(x', z)) - E(Y | do(x, z))]P(z | do(x)). \quad (15.15)$$

The intuition is simple; the natural direct effect is the weighted average of the controlled direct effect, using the causal effect $P(z | do(x))$ as a weighting function.

One condition for the validity of Equation 15.17 is that $Z_x \perp\!\!\!\perp Y_{x',z} | W$ holds for some set W of measured covariates. This technical condition in itself, like the ignorability condition of Equation 15.12, is close to meaningless for most investigators because it is not phrased in terms of realized variables. The structural interpretation of counterfactuals (Equation 15.8) can be invoked at this point to unveil the graphical interpretation of this condition. It states that W should be admissible (i.e., satisfy the back-door condition) relative the path(s) from Z to Y . This condition, satisfied by W_2 in Figure 15.4b, is readily comprehended by empirical researchers, and the task of selecting such measurements, W , can then be guided by the available scientific knowledge. Additional graphical and counterfactual conditions for identification are derived in Pearl (2001), Petersen, Sinisi, and van der Laan (2006), and Imai, Keele and Yamamoto (2008).

In particular, it was shown (Pearl, 2001) that Equation 15.15 is both valid and identifiable in Markovian models (i.e., no unobserved confounders) where each term on the right can be reduced to a “do-free” expression using Equation 15.6 or Equation 15.7 and then estimated by regression.

For example, for the model in Figure 15.4b, Equation 15.15 reads:

$$DE_{x,x'}(Y) = \sum_z \sum_{w_2} P(w_2) [E(Y | x', z, w_2) - E(Y | x, z, w_2)] \sum_{w_1} \frac{P(z | x, w_2, w_1) P(w_1)}{P(z | x, w_1) P(w_1)}. \quad (15.16)$$

However, for the confounding-free model of Figure 15.4a, we have:

$$DE_{x,x'}(Y) = \sum_z [E(Y | x', z) - E(Y | x, z)]P(z | x). \quad (15.17)$$

Both Equations 15.16 and 15.17 can easily be estimated by a two-step regression.

Natural Indirect Effects

Remarkably, the definition of the natural direct effect (Equation 15.14) can be turned around and provides an operational definition for the *indirect*

effect—a concept shrouded in mystery and controversy because it is impossible, using standard intervention, to disable the direct link from X to Y so as to let X influence Y solely via indirect paths.

The *natural indirect effect* (IE) of the transition from x to x' is defined as the expected change in Y affected by holding X constant, at $X = x$, and changing Z to whatever value it would have attained had X been set to $X = x'$. Formally, this reads (Pearl, 2001):

$$IE_{x,x'}(Y) = E[(Y_{x,Z_{x'}}) - E(Y_x)], \quad (15.18)$$

which is almost identical to the direct effect (Equation 15.14) save for exchanging x and x' in the first term.

Indeed, it can be shown that, in general, the total effect (TE) of a transition is equal to the *difference* between the direct effect of that transition and the indirect effect of the reverse transition. Formally,

$$TE_{x,x'}(Y) = E(Y_{x'} - Y_x) = DE_{x,x'}(Y) - IE_{x',x}(Y). \quad (15.19)$$

In linear systems, where reversal of transitions amounts to negating the signs of their effects, we have the standard additive formula:

$$TE_{x,x'}(Y) = DE_{x,x'}(Y) + IE_{x,x'}(Y). \quad (15.20)$$

Because each term above is based on an independent operational definition, this equality constitutes a formal justification for the additive formula used routinely in linear systems.

Note that, although it cannot be expressed in *do*-notation, the indirect effect has clear policy-making implications. For example, in the hiring discrimination context, a policy maker may be interested in predicting the gender mix in the work force if gender bias is eliminated and all applicants are treated equally—say, the same way that males are currently treated. This quantity will be given by the indirect effect of gender on hiring, mediated by factors such as education and aptitude, which may be gender dependent.

More generally, a policy maker may be interested in the effect of issuing a directive to a select set of subordinate employees, or in carefully controlling the routing of messages in a network of interacting agents. Such applications motivate the analysis of *path-specific effects*, that is, the effect of X on Y through a selected set of paths (Avin, Shpitser, & Pearl, 2005).

In all these cases, the policy intervention invokes the selection of signals to be sensed, rather than variables to be fixed. Therefore, Pearl (2001) has suggested that *signal sensing* is more fundamental to the notion of causation than *manipulation*, the latter being but a crude way of stimulating the

former in experimental setup. The mantra “No causation without manipulation” must be rejected (see Pearl, 2009a, Section 11.4.5).

It is remarkable that counterfactual quantities like DE and IE that could not be expressed in terms of $do(x)$ operators, and therefore appear void of empirical content, can, under certain conditions, be estimated from empirical studies and serve to guide policies. Awareness of this potential should embolden researchers to go through the definitional step of the study and freely articulate the target quantity $Q(M)$ in the language of science, that is, counterfactuals, despite the seemingly speculative nature of each assumption in the model (Pearl, 2000).

The Mediation Formula: A Simple Solution to a Thorny Problem

This subsection demonstrates how the solution provided in Equations 15.17 and 15.20 can be applied to practical problems of assessing mediation effects in nonlinear models. We will use the simple mediation model of Figure 15.4a, where all error terms (not shown explicitly) are assumed to be mutually independent, with the understanding that adjustment for appropriate sets of covariates W may be necessary to achieve this independence and that integrals should replace summations when dealing with continuous variables (Imai et al., 2008).

Combining Equations 15.17, 15.19, and 15.20, the expression for the indirect effect, IE , becomes:

$$IE_{xx'}(Y) = \sum_z E(Y | x, z)[P(z | x') - P(z | x)], \quad (15.21)$$

which provides a general and easy-to-use formula for mediation effects, applicable to any nonlinear system, any distribution (of U), and any type of variables. Moreover, the formula is readily estimable by regression, making no assumption whatsoever about the parametric form of the underlying process. Owing to its generality and ubiquity, I have referred to this expression as the “Mediation Formula” (Pearl, 2009b).

The Mediation Formula represents the average increase in the outcome Y that the transition from $X = x$ to $X = x'$ is expected to produce absent any direct effect of X on Y . Although based on solid causal principles, it embodies no causal assumption other than the generic mediation structure of Figure 15.4a. When the outcome Y is binary (e.g., recovery, or hiring) the ratio $(1 - IE/TE)$ represents the fraction of responding individuals who owe their response to direct paths, whereas $(1 - IE/TE)$ represents the fraction who owe their response to Z -mediated paths.

The Mediation Formula tells us that IE depends only on the expectation of the counterfactual $Y_{xz'}$, not on its functional form $f_Y(x, z, u_Y)$ or its distribution $P(Y_{xz} = y)$. Therefore, it calls for a two-step regression that, in

principle, can be performed nonparametrically. In the first step, we regress Y on X and Z and obtain the estimate:

$$g(x, z) = E(Y | x, z)$$

for every (x, z) cell. In the second step, we estimate the expectation of $g(x, z)$ conditional on $X = x'$ and $X = x$, respectively, and take the difference:

$$IE_{x,x'}(Y) = E_{Z|x'}(g(x, z)) - E_{Z|x}(g(x, z)).$$

Nonparametric estimation is not always practical. When Z consists of a vector of several mediators, the dimensionality of the problem would prohibit the estimation of $E(Y | x, z)$ for every (x, z) cell, and the need arises to use parametric approximation. We can then choose any convenient parametric form for $E(Y | x, z)$ (e.g., linear, logit, probit), estimate the parameters separately (e.g., by regression or maximum likelihood methods), insert the parametric approximation into Equation 15.21, and estimate its two conditional expectations (over z) to get the mediated effect (Pearl, 2010b; VanderWeele, 2009).

When applied to linear models, the Mediation Formula yields, of course, the standard product of coefficients. For example, the linear version of Figure 15.4a reads:

$$\begin{aligned} x &= u_x \\ z &= b_x x + u_z \\ y &= c_x x + c_z z + u_y. \end{aligned} \tag{15.22}$$

Computing the conditional expectation in Equation 15.21 gives:

$$E(Y | x, z) = E(c_x x + c_z z + u_y) = c_x x + c_z z$$

and yields:

$$\begin{aligned} IE_{x,x'}(Y) &= \sum_z (c_x x + c_z z) [P(z | x') - P(z | x)] \\ &= c_z [E(Z | x') - E(Z | x)] \end{aligned} \tag{15.23}$$

$$= (x' - x)(c_z b_x) \tag{15.24}$$

$$= (x' - x)(b - c_x) \tag{15.25}$$

where b is the total effect coefficient, $b = (E(Y | x') - E(Y | x)) / (x' - x) = c_x + c_z b_x$.

Thus, we obtained the standard expressions for indirect effects in linear systems, which can be estimated either as a difference in two regression coefficients (Equation 15.25) or a product of two regression coefficients (Equation 15.24), with Y regressed on both X and Z . However, when extended to nonlinear systems, these two strategies yield conflicting results (MacKinnon & Dwyer, 1993; MacKinnon et al., 2007b), and the question arose as to which strategy should be used in assessing the size of mediated effects (Freedman, Graubard, & Schatzkin, 1992; MacKinnon & Dwyer, 1993; MacKinnon et al., 2007b; Molenberghs et al., 2002). Pearl (2010b) shows that both strategies yield highly distorted results in nonlinear models, even when correct parametric forms are assumed. The reason lies in a violation of step 1 (defining the target quantity) of the “Methodological Dictates and Ethical Considerations” section. Researchers failed to define the causal quantity of interest and were postulating, estimating, and comparing parameters that were related to, yet hardly resembling, DE and IE . The Mediation Formula captures the correct target quantity and helps researchers cross the nonlinear barrier that has held back the mediation literature for more than half a century. Simple examples using Bernoulli/binary noise, logistic, and probit models are illustrated in Pearl (2010a, 2010b).

In addition to providing causally sound estimates for mediation effects, the Mediation Formula also enables researchers to evaluate analytically the effectiveness of various parametric specifications relative to any assumed model. This type of analytical “sensitivity analysis” has been used extensively in statistics for parameter estimation but could not be applied to mediation analysis because of the absence of an objective target quantity that captures the notion of indirect effect in both linear and nonlinear systems, free of parametric assumptions. The Mediation Formula has removed this barrier (Imai, Keele, & Tingley, 2010; Li, Schneider, & Bennett, 2007).

The derivation of the Mediation Formula (Pearl, 2001) was facilitated by taking seriously the four steps of the structural methodology (“Methodological Dictates and Ethical Considerations” section) together with the graph-counterfactual-structural symbiosis spawned by the structural interpretation of counterfactuals (Equation 15.8). In contrast, when the mediation problem is approached from an exclusivist potential-outcome viewpoint, void of the structural guidance of Equation 15.8, counterintuitive definitions ensue, carrying the label “principal stratification” (Rubin, 2004, 2005), which are at variance with common understanding of direct and indirect effects. For example, the direct effect is definable only in units absent of indirect effects. This means that a grandfather would be deemed to have no direct effect on his grandson’s behavior in families where he has had some effect on the father. This precludes from the analysis all typical families, in which a father and a grandfather have simultaneous,

References

- Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables (with comments). *Journal of the American Statistical Association*, *91*, 444–472.
- Avin, C., Shpitser, I., & Pearl, J. (2005). Identifiability of path-specific effects. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence IJCAI-05* (pp. 357–363). Edinburgh, UK: Morgan-Kaufmann Publishers.
- Balke, A., & Pearl, J. (1994). Probabilistic evaluation of counterfactual queries. In *Proceedings of the Twelfth National Conference on Artificial Intelligence* (Vol. I, pp. 230–237). Menlo Park, CA: MIT Press.
- Balke, A., & Pearl, J. (1995). Counterfactuals and policy analysis in structural models. In P. Besnard & S. Hanks (Eds.), *Uncertainty in artificial intelligence 11* (pp. 11–18). San Francisco: Morgan Kaufmann.
- Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, *51*, 1173–1182.
- Brent, R., & Lok, L. (2005). A fishing buddy for hypothesis generators. *Science*, *308*, 523–529.
- Cox, D. R. (1958). *The planning of experiments*. New York: John Wiley and Sons.
- Dawid, A. P. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society, Series B*, *41*, 1–31.
- Dawid, A. P. (2000). Causal inference without counterfactuals (with comments and rejoinder). *Journal of the American Statistical Association*, *95*, 407–448.
- Elwert, F., & Winship, C. (2010). Effect heterogeneity and bias in main-effects-only regression models. In R. Dechter, H. Geffner, & J. Y. Halpern (Eds.), *Heuristics, probability and causality: A tribute to Judea Pearl* (pp. 327–336). London: College Publications.
- Freedman, L. S., Graubard, B. I., & Schatzkin, A. (1992). Statistical validation of intermediate endpoints for chronic diseases. *Statistics in Medicine*, *8*, 167–178.
- Greenland, S., Pearl, J., & Robins, J. M. (1999). Causal diagrams for epidemiologic research. *Epidemiology*, *10*, 37–48.
- Hafeman, D. M., & Schwartz, S. (2009). Opening the black box: A motivation for the assessment of mediation. *International Journal of Epidemiology*, *38*, 838–845.
- Holland, P. W. (1988). Causal inference, path analysis, and recursive structural equations models. In C. Clogg (Ed.), *Sociological Methodology* (pp. 449–484). Washington, DC: American Sociological Association.
- Imai, K., Keele, L., & Tingley, D. (2010). *A general approach to causal mediation analysis*. Technical report. Princeton, NJ: Princeton University.
- Imai, K., Keele, L., & Yamamoto, T. (2008). *Identification, inference, and sensitivity analysis for causal mediation effects*. Technical report. Princeton, NJ: Princeton University.
- Judd, C. M., & Kenny, D. A. (1981). Process analysis: Estimating mediation in treatment evaluations. *Evaluation Review*, *5*, 602–619.

complementary influences on children's upbringing. In linear systems, to take an even sharper example, the direct effect would be undefined whenever indirect paths exist from the cause to its effect. The emergence of such paradoxical conclusions underscores the wisdom, if not necessity, of a symbiotic analysis, in which the counterfactual notation $Y_x(u)$ is governed by its structural definition, Equation 15.8.¹⁶ It also brings into focus the ethical issue of inclusiveness and its role in scientific research and education.

Conclusion

Statistics is strong in inferring distributional parameters from sample data. Causal inference requires two additional ingredients: a science-friendly language for articulating causal knowledge, and a mathematical machinery for processing that knowledge, combining it with data and drawing new causal conclusions about a phenomenon. This chapter presents nonparametric *structural causal models* as a formal and meaningful language for meeting these challenges, thus easing the ethical tensions that follow the disparity between causal quantities sought by scientists and associational quantities inferred from observational studies. The algebraic component of the structural language coincides with the potential–outcome framework, and its graphical component embraces Wright's method of path diagrams (in its nonparametric version). When unified and synthesized, the two components offer empirical investigators a powerful methodology for causal inference that resolves longstanding problems in the empirical sciences. These include the control of confounding, the evaluation of policies, the analysis of mediation, and the algorithmization of counterfactuals.

In particular, the analysis of mediation demonstrates the benefit of adhering to the methodological principles described. The development of the Mediation Formula (Equations 15.17 and 15.20) has liberated researchers from the blindfolds of parametric thinking and allows them to assess direct and indirect effects for any type of variable, with minimum assumptions regarding the underlying process.¹⁷

¹⁶ Such symbiosis is now standard in epidemiology research (Hafeman & Schwartz, 2009; Petersen et al., 2006; Robins, 2001; VanderWeele, 2009; VanderWeele & Robins, 2007) and is making its way slowly toward the social and behavioral sciences (e.g., Elwert & Winship, 2010; Morgan & Winship, 2007).

¹⁷ Author note: Portions of this chapter are adapted from Pearl (2009a, 2009b, 2010a). I am grateful to A. T. Panter and Sonya K. Serba for their encouragement and flexibility in the writing of this chapter. This research was supported in part by grants from National Science Foundation (IIS-0535223) and Office of Naval Research (N000-14-09-1-0665).

- Li, Y., Schneider, J. A., & Bennett, D. A. (2007). Estimation of the mediation effect with a binary mediator. *Statistics in Medicine*, 26, 3398–3414.
- MacKinnon, D. P., & Dwyer, J. H. (1993). Estimating mediated effects in prevention studies. *Evaluation Review*, 4, 144–158.
- MacKinnon, D. P., Fairchild, A. J., & Fritz, M. S. (2007a). Mediation analysis. *Annual Review of Psychology*, 58, 593–614.
- MacKinnon, D. P., Lockwood, C. M., Brown, C. H., Wang, W., & Hoffman, J. M. (2007b). The intermediate endpoint effect in logistic and probit regression. *Clinical Trials*, 4, 499–513.
- Molenberghs, G., Buyse, M., Geys, H., Renard, D., Burzykowski, T., & Alonso, A. (2002). Statistical challenges in the evaluation of surrogate endpoints in randomized trials. *Controlled Clinical Trials*, 23, 607–625.
- Morgan, S. L., & Winship, C. (2007). *Counterfactuals and causal inference: Methods and principles for social research (analytical methods for social research)*. New York: Cambridge University Press.
- Muller, D., Judd, C. M., & Yzerbyt, V. Y. (2005). When moderation is mediated and mediation is moderated. *Journal of Personality and Social Psychology*, 89, 852–863.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. Essay on principles. Section 9. *Statistical Science*, 5, 465–480.
- Pearl, J. (1993). Comment: Graphical models, causality, and intervention. *Statistical Science*, 8, 266–269.
- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82, 669–710.
- Pearl, J. (1998). Graphs, causality, and structural equation models. *Sociological Methods and Research*, 27, 226–284.
- Pearl, J. (2000). Comment on A. P. Dawid's causal inference without counterfactuals. *Journal of the American Statistical Association*, 95, 428–431.
- Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence* (pp. 411–420). San Francisco: Morgan Kaufmann.
- Pearl, J. (2003). Statistics and causal inference: A review. *Test Journal*, 12, 281–345.
- Pearl, J. (2009a). *Causality: Models, reasoning, and inference* (2nd ed.). New York: Cambridge University Press.
- Pearl, J. (2009b). Causal inference in statistics: An overview. *Statistics Surveys*, 3, 96–146. Retrieved from http://ftp.cs.ucla.edu/pub/stat_ser/r350.pdf
- Pearl, J. (2009c). Letter to the editor: Remarks on the method of propensity scores. *Statistics in Medicine*, 28, 1415–1416. Retrieved from http://ftp.cs.ucla.edu/pub/stat_ser/r345-sim.pdf
- Pearl, J. (2009d). *Myth, confusion, and science in causal analysis*. Technical report R-348. Los Angeles: University of California, Los Angeles. Retrieved from http://ftp.cs.ucla.edu/pub/stat_ser/r348.pdf
- Pearl, J. (2010a). An introduction to causal inference. *The International Journal of Biostatistics*, 6. doi: 10.2202/1557-4679.1203. Retrieved from <http://www.bepress.com/ijb/vol6/iss2/7>
- Pearl, J. (2010b). *The mediation formula: A guide to learning causal pathways*. Technical report TR-363. Los Angeles: University of California, Los Angeles. Retrieved from http://ftp.cs.ucla.edu/pub/stat_ser/r363.pdf

- Pearl, J. (2010c). On a class of bias-amplifying covariates that endanger effect estimates. In P. Grunwald & P. Spirtes (Eds.), *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence* (pp. 417–432). Corralis, OR: AUAI Press.
- Pearl, J., & Paz, A. (2010). Confounding equivalence in observational studies. In P. Grunwald & P. Spirtes (Eds.), *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence* (pp. 433–441). Corralis, OR: AUAI Press. Retrieved from http://ftp.cs.ucla.edu/pub/stat_ser/r343.pdf
- Pearl, J., & Robins, J. M. (1995). Probabilistic evaluation of sequential plans from causal models with hidden variables. In P. Besnard & S. Hanks (Eds.), *Uncertainty in artificial intelligence 11* (pp. 444–453). San Francisco: Morgan Kaufmann.
- Petersen, M. L., Sinisi, S. E. & van der Laan, M. J. (2006). Estimation of direct causal effects. *Epidemiology*, 17, 276–284.
- Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—applications to control of the healthy workers survivor effect. *Mathematical Modeling*, 7, 1393–1512.
- Robins, J. M. (2001). Data, design, and background knowledge in etiologic inference. *Epidemiology*, 12, 313–320.
- Robins, J. M., & Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3, 143–155.
- Rosenbaum, P., & Rubin, D. (1983). The central role of propensity score in observational studies for causal effects. *Biometrika*, 70, 41–55.
- Rosenbaum, P. R. (2002). *Observational studies* (2nd ed.). New York: Springer-Verlag.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66, 688–701.
- Rubin, D. B. (2004). Direct and indirect causal effects via potential outcomes. *Scandinavian Journal of Statistics*, 31, 161–170.
- Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100, 322–331.
- Rubin, D. B. (2009). Author's reply: Should observational studies be designed to allow lack of balance in covariate distributions across treatment group? *Statistics in Medicine*, 28, 1420–1423.
- Shpitser, I., & Pearl, J. (2006). Identification of conditional interventional distributions. In R. Dechter & T. S. Richardson (Eds.), *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence* (pp. 437–444). Corvallis, OR: AUAI Press.
- Shrout, P. E., & Bolger, N. (2002). Mediation in experimental and nonexperimental studies: New procedures and recommendations. *Psychological Methods*, 7, 422–445.
- Simon, H. A., & Rescher, N. (1966). Cause and counterfactual. *Philosophy and Science*, 33, 323–340.
- Sobel, M. E. (1998). Causal inference in statistical models of the process of socioeconomic achievement. *Sociological Methods & Research*, 27, 318–348.
- Sobel, M. E. (2008). Identification of causal parameters in randomized studies with mediating variables. *Journal of Educational and Behavioral Statistics*, 33, 230–231.

- Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, prediction, and search* (2nd ed.). Cambridge, MA: MIT Press.
- Tian, J., Paz, A., & Pearl, J. (1998). *Finding minimal separating sets*. Technical report R-254. Los Angeles: University of California, Los Angeles.
- Tian, J., & Pearl, J. (2000). Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28, 287–313.
- VanderWeele, T. J. (2009). Marginal structural models for the estimation of direct and indirect effects. *Epidemiology*, 20, 18–26.
- VanderWeele, T. J., & Robins, J. M. (2007). Four types of effect modification: A classification based on directed acyclic graphs. *Epidemiology*, 18, 561–568.
- Wermuth, N. (1992). On block-recursive regression equations. *Brazilian Journal of Probability and Statistics* (with discussion), 6, 1–56.
- Wermuth, N., & Cox, D. (1993). Linear dependencies represented by chain graphs. *Statistical Science*, 8, 204–218.
- Whittaker, J. (1990). *Graphical models in applied multivariate statistics*. Chichester, UK: John Wiley.
- Wilkinson, L., & the Task Force on Statistical Inference. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54, 594–604.
- Wright, S. (1921). Correlation and causation. *Journal of Agricultural Research*, 20, 557–585.