

Causal Inference: A Tale of Three Frameworks

LINBO WANG^{1,*}, THOMAS S. RICHARDSON², AND JAMES M. ROBINS³

¹*Department of Statistical Sciences, University of Toronto, Toronto, ON, Canada*

²*Department of Statistics, University of Washington, Seattle, WA, U.S.A.*

³*Harvard T.H. Chan School of Public Health, Boston, MA, U.S.A.*

Abstract

Causal inference is a central goal across many scientific disciplines. Over the past several decades, three major frameworks have emerged to formalize causal questions and guide their analysis: the potential outcomes framework, structural equation models, and directed acyclic graphs. Although these frameworks differ in language, assumptions, and philosophical orientation, they often lead to compatible or complementary insights. This paper provides a comparative introduction to the three frameworks, clarifying their connections, highlighting their distinct strengths and limitations, and illustrating how they can be used together in practice. The discussion is aimed at researchers and graduate students with some background in statistics or causal inference who are seeking a conceptual foundation for applying causal methods across a range of substantive domains.

Keywords *directed acyclic graphs; identification; potential outcomes; structural equation models; SWIGs*

1 Introduction

Causal inference is the science of understanding the consequences of interventions, requiring assumptions that extend beyond those needed for purely associational analysis. Its importance has grown rapidly in the era of machine learning and artificial intelligence, where the ability to draw reliable causal conclusions is central to building systems that are not only predictive but also trustworthy, transparent, and robust to distributional shifts (Peters et al., 2016; Wachter et al., 2017; Pearl, 2019; Arjovsky et al., 2020; Bühlmann, 2020; Tjoo and Guan, 2021; Schölkopf, 2022; Jiao et al., 2024). Over the past decades, three foundational frameworks have emerged to formalize causal reasoning: the **potential outcomes** framework, nonparametric **structural equation models** (NPSEMs), and **directed acyclic graphs** (DAGs). Each framework carries its own formal machinery, conceptual underpinnings, and historical roots. Although they originated in distinct disciplinary traditions, they are now increasingly recognized as complementary, and in many cases translatable into one another.

A substantial literature surveys causal inference from within a single framework or with an emphasis on identification and estimation approaches (e.g. Imbens and Wooldridge, 2009; Spirtes, 2010; Pearl, 2010a; Kuang et al., 2020; Yao et al., 2021; Li et al., 2023; Jiao et al., 2024). However, there has been little concise, side-by-side treatment that translates assumptions and results across the three frameworks and clarifies when they agree and when they differ, despite earlier comparative discussions such as Greenland and Brumback (2002), Pearl (2009), and

*Corresponding author. Email: linbo.wang@utoronto.ca.

Robins and Richardson (2011). This review aims to fill that gap. We offer a selective review of the three frameworks, with the goal of elucidating their differences, similarities, and interconnections. Rather than offering a comprehensive survey, we aim to guide readers, particularly those new to the field or approaching causal inference from a single perspective, through the conceptual landscape of causal inference using simple, illustrative examples.

In this review, we focus on *interventional* questions, which ask what would happen under specific actions or treatments applied to the entire population, without conditioning on what actually occurred. These contrast with *counterfactual* questions, which ask what would have happened for the same individual, had things been different, given what actually occurred. For example, as noted by Peters et al. (2017, Remark 6.20(ii)), suppose someone offers you \$10,000 if you correctly predict the result of a coin flip. You guess “heads” and lose. The interventional effect of guessing “heads” versus “tails” is zero, whereas the counterfactual effect, given that you actually guessed “heads” and lost, is \$10,000. While counterfactuals play a central role in philosophical and legal reasoning, their identification typically requires stronger, often untestable, assumptions. See Remark 3 and Wu and Wang (2026) for further discussion of this distinction.

Although we do not delve into estimation methods or unmeasured confounding in depth, we briefly comment on these in later sections to clarify their connection to the foundational frameworks. Other approaches, such as Dawid’s decision-theoretic framework (Dawid, 2000, 2015; Richardson and Robins, 2023), are also omitted from our discussion. The decision-theoretic perspective defines causality through conditional independence relations among decisions, observations, and consequences, without invoking counterfactuals or structural equations. Conceptually, it shares common ground with all three frameworks reviewed here: like potential outcomes, it distinguishes between observed and hypothetical actions; like structural models, it emphasizes invariance of mechanisms under intervention; and like causal DAGs, it encodes causal assumptions through independence relations. However, it differs from these frameworks in that it formulates causal reasoning entirely within the observable world, avoiding reference to unobservable counterfactual quantities. Because our focus is on frameworks that explicitly model potential outcomes or interventions, we do not explore the decision-theoretic approach in detail. For these topics, and for comprehensive treatments of any individual framework, we refer readers to existing textbooks (e.g., Pearl, 2009; Peters et al., 2017; Imbens and Rubin, 2015; Hernán and Robins, 2025) and the broader specialized literature.

Another influential framework that we do not review in detail is the sufficient-component cause (SCC) framework; see Rothman (1976), Greenland and Poole (1988), Greenland and Brumback (2002), Rothman et al. (2008, Chap. 2), VanderWeele and Robins (2009), VanderWeele and Richardson (2012, § 7), and VanderWeele (2015, p. 10 & Chap. 10) for detailed discussions. SCC models provide a mechanistic representation of causation and are well suited for questions about explanatory causal structures and mechanistic interaction. These mechanistic questions differ somewhat from the interventional questions that motivate our review, and several SCC quantities cannot be expressed within the causal DAG formalism. Since our goal here is to compare three frameworks for interventional causal questions, we do not cover SCC models, although they remain an important component of the broader causal inference literature.

A further set of traditions, especially influential in health and legal applications, involves more informal approaches to causal inference alongside formal models. Greenland (2004), for example, distinguishes canonical approaches based on considerations such as the Hill’s list (Hill, 1965), in which causality is treated as a property of an observed association to be judged using qualitative “symptoms” like temporality, strength, consistency, and dose response, from methodologic modeling or bias analysis, which uses parametric models and sensitivity analyses to assess

the impact of selection, unmeasured confounding, and measurement error on effect estimates. These approaches are widespread in applied practice and can be combined with potential outcome, structural, and graphical models, but they fall outside our focus on three foundational frameworks for representing and identifying interventional effects.

The frameworks reviewed here primarily concern single-stage interventions. Many applied settings, however, involve sequential or adaptive decisions, such as treatment policies in medicine, dynamic pricing, or reinforcement learning. These settings have been extensively studied in the literature on dynamic treatment regimes and longitudinal causal inference, most notably through Robins’s development of g-methods and related approaches (Robins, 1986; Hernán and Robins, 2025). These frameworks generalize the ideas discussed here to multi-stage interventions and time-varying treatments. Because our focus is on the conceptual connections among the three foundational frameworks, we do not cover these extensions in detail.

Our review proceeds as follows. We begin by introducing the three frameworks in turn, highlighting how each connects observational data to hypothetical interventions and how each encodes key assumptions, such as the absence of unmeasured confounding. In Section 3, we explore how these frameworks can be formally related. We show how potential outcomes arise naturally within NPSEMs, and conversely, how systems of potential outcomes can be used to construct NPSEMs through canonical representations. We also examine how causal DAG models are implied by structural equation models under the assumptions of independent errors and autonomy, and how NPSEMs can, in turn, be generated from a given causal DAG using the functional representation lemma. Finally, we introduce Single World Intervention Graphs (SWIGs), which integrate the graphical and potential outcomes approaches by explicitly representing potential outcomes on a modified DAG, thereby bridging these frameworks without invoking all of the assumptions imposed by structural models.

Section 4 offers a comparative analysis of the three frameworks along several dimensions, including their philosophical orientation, expressive capacity, and identification power. We discuss when the additional assumptions imposed by NPSEMs with independent errors can lead to stronger identification results and when such assumptions may be overly restrictive. We conclude with practical guidance for applied researchers and reflect on recent methodological developments that seek to unify these perspectives into a coherent and transparent workflow for causal inference.

2 Three Frameworks for Causal Inference

The formalization of causal effects in mathematical terms is widely recognized as one of the major advances in the modern theory of causality. Unlike purely associational quantities, causal effects characterize how an outcome responds to different hypothetical interventions that set a treatment or exposure to alternative values. At a broader level, counterfactual reasoning considers what *would* have happened for the same individual under alternative scenarios that may be contrary to fact. This conceptual leap distinguishes causal inference from classical statistics, which has traditionally focused on describing correlations and conditional dependencies among *observed* data, or on making inferences regarding the distribution from which a sample was drawn.

Causal inference can be understood through at least three principal frameworks: (1) Potential Outcomes, (2) SEMs, and (3) DAGs. While each framework stems from a different disciplinary tradition, all aim to bridge the gap between observational data and interventional scenarios. This section provides a detailed account of each framework, with the key assumptions

Table 1: How the three causal inference frameworks relate the observational and interventional worlds, and how they encode assumptions about unmeasured confounding. Here, A denotes the treatment, Y denotes the outcome, and L denotes covariates.

Framework	How the observational world relates to the (hypothetical) interventional world	How to encode or check no unmeasured confounding
Potential Outcomes	Consistency: $Y = Y(A)$	Ignorability ^a : $A \perp\!\!\!\perp Y(a) \mid L$
NPSEM ^b	Autonomy ^c : intervention modeled by replacing the structural equation for A , while keeping all other equations unchanged	Independent errors ^d (NPSEM-IE)
Causal DAG	Modularity: intervention modeled by removing incoming edges	Causal sufficiency: no unmeasured common causes of variables in the graph; under this assumption, L can be chosen to satisfy the backdoor criterion ^e

^a $\perp\!\!\!\perp$ denotes independence.

^b An NPSEM typically assumes that interventions on all variables are well-defined.

^c In the NPSEM literature, this property is also called modularity, highlighting that the mechanism for A can be replaced without altering other equations.

^d A generic NPSEM does not impose independent errors; the NPSEM with independent errors (NPSEM-IE) is stronger in that it adds this assumption.

^e The scope differs across frameworks. In the potential outcomes framework, ignorability is typically imposed for a specific effect of A on Y , and thus enables identification of that individual effect alone. In contrast, an NPSEM-IE encodes a full system of structural equations and independent errors, which enables identification of all the interventional effects in the model. Causal DAG models with causal sufficiency rule out hidden common causes; under this assumption, graphical tools such as the backdoor criterion provide effect-specific tests for identification across multiple treatment–outcome pairs.

summarized in Table 1. We discuss their relationships and how they both unify and diverge in their approaches to causal reasoning in the following sections.

2.1 The Potential Outcomes Framework

The potential outcomes framework provides an intuitive and mathematically tractable foundation for causal inference. It was originally introduced by Neyman (1923) for randomized experiments and was later formalized and extended to observational studies by Rubin (1974). This was further extended by Robins (1986) to studies with time-varying treatments and confounders, as well as to direct and indirect effects and feedback of one cause on another. This framework is closely related to an independent line of work in econometrics that developed the so-called switching regression model (Roy, 1951; Manski, 1993). These developments allow researchers to pose well-defined causal questions, even in the absence of randomized experiments, by explicitly invoking assumptions that link observed data to potential outcomes. We refer interested readers to Rubin (2005); Morgan and Winship (2014); Imbens and Rubin (2015); Hernán and Robins (2025) for a more detailed introduction.

Suppose the observed data consist of a triple (L, A, Y) , where L denotes (a vector of) covariates, $A \in \{0, 1\}$ is a binary treatment indicator (for example, placebo versus active treatment), and Y is the outcome of interest. The variables L capture characteristics that may confound the treatment and outcome relationship. These covariates play a key role in adjustment procedures that aim to mimic randomization in observational settings.

For example, consider a job-training program ($A = 1$) intended to improve employment outcomes (Y). Individuals with higher prior education or stronger motivation (L) may be both more likely to enroll in the program and more likely to achieve better employment outcomes regardless of enrollment. This example provides a simple setting for comparing how the three causal frameworks represent and identify the same underlying causal effect.

For a binary treatment variable A , we define two potential outcome variables:

- $Y(0)$: the value of Y that would be observed for a given unit if assigned $A = 0$ (placebo);
- $Y(1)$: the value of Y that would be observed for a given unit if assigned $A = 1$ (treatment).

The variables $Y(0)$ and $Y(1)$ are two distinct random variables representing hypothetical outcomes under different treatment assignments. They are *not* different realizations of the same variable. For each unit, only one of these two outcomes is realized and observed, depending on the actual treatment received. The other remains counterfactual. This leads to the *Fundamental Problem of Causal Inference* (Holland, 1986): we can never simultaneously observe both potential outcomes $Y(0)$ and $Y(1)$ for the same individual. From this perspective, the potential outcomes framework views causal inference as a missing data problem, since at least one of the two potential outcomes is missing for each unit.

To address this challenge, causal inference often focuses on population-level quantities that average over the distribution of potential outcomes. This leads to a distinction between two types of causal effects:

- **Individual causal effect:** Defined as $Y_i(1) - Y_i(0)$, this measures the causal effect for a specific unit i . Although conceptually straightforward, it is unobservable and generally not identifiable from data without strong assumptions or additional structure (such as deterministic effects or monotonicity).
- **Average causal effect (ACE):** Defined as $ACE = \mathbb{E}\{Y(1) - Y(0)\}$, the ACE captures the expected difference in outcomes had everyone in the population received treatment versus control. This is the primary estimand (i.e. parameter of interest) in many empirical studies, and it is identifiable under assumptions such as complete randomization or conditional ignorability (see eqn. (2.1)).

Remark 1. In addition to the individual and average causal effects, other causal estimands that offer more nuanced insights into treatment effects may also be of interest. One such quantity is the *Conditional Average Treatment Effect (CATE)*, defined as $\mathbb{E}\{Y(1) - Y(0) \mid X = x\}$, which represents the average treatment effect for individuals with covariates $X = x$ where $X \subset L$. CATE is central to personalized decision-making and helps characterize treatment effect heterogeneity across subpopulations. Another commonly used estimand is the *Effect of Treatment on the Treated (ETT)*, given by $\mathbb{E}\{Y(1) - Y(0) \mid A = 1\}$, which captures the causal effect among those who actually received the treatment; the ETT is useful, for example, when evaluating whether a program is effective for those who are participating in it. The *Quantile Treatment Effect (QTE)*, expressed as $Q_{Y(1)}(\tau) - Q_{Y(0)}(\tau)$, quantifies the difference in the τ -th quantiles of the potential outcome distributions under treatment and control. This allows researchers to assess how treatment affects the entire outcome distribution rather than just its mean. Importantly, in contrast to ACE, QTEs are not “individual-level” effects: the difference between the

treatment and control quantiles need not correspond to the effect experienced by any particular individual. Instead, QTEs characterize population-level distributional shifts, which may provide valuable insight into distributional heterogeneity even if they do not map to individual units.

Remark 2. Potential outcomes are sometimes referred to as *counterfactuals* in the literature. *Before* treatment assignment, both $Y(1)$ and $Y(0)$ are potential outcomes. *After* treatment assignment, the outcome corresponding to the received treatment is *factual*, while the other is *counterfactual*. (However, many papers use the terms potential outcomes and counterfactuals interchangeably.) To avoid ambiguity, we use the term *potential outcomes* to collectively refer to both $Y(1)$ and $Y(0)$.

Remark 3. The estimands we have introduced so far are sometimes referred to as *interventional* estimands, as they pertain to the distribution of outcomes under hypothetical interventions, such as setting the treatment to a fixed value. Another type of question that may be of interest involves *counterfactual* reasoning. For example, when Y is binary, the persuasion rate is defined as $\mathbb{P}\{Y(1) = 1 \mid Y(0) = 0\}$ (Jun and Lee, 2023), which quantifies the probability that an individual would take an action under treatment given that they would not have done so under control. Other examples include the probability of necessity, $\mathbb{P}\{Y(0) = 0 \mid A = 1, Y = 1\}$; the probability of sufficiency, $\mathbb{P}\{Y(1) = 1 \mid A = 0, Y = 0\}$; and the probability of necessity and sufficiency, $\mathbb{P}\{Y(1) = 1, Y(0) = 0\}$ (Pearl, 2009, § 9.2). Identification of counterfactual estimands is typically much more involved, as it relies on the dependence between potential outcomes, which cannot be recovered from observed data, even in randomized controlled experiments. We refer readers to Peters et al. (2017, Example 3.4) for an interesting example and discussions on whether counterfactuals or interventional quantities should serve as the basis for legal decisions; see also Robins and Greenland (1989), Robins and Greenland (1991), and Pearl (2009, § 1.4.4) and Wu and Wang (2026). Along with association questions, interventional and counterfactual questions are referred to as the three levels of the ladder of causation by Pearl and Mackenzie (2018, § 1).

A crucial component of the Potential Outcome framework is the *Stable Unit Treatment Value Assumption (SUTVA)* (Rubin, 1980), which posits that (i) there are no multiple versions of the treatment, and (ii) one unit’s outcome is unaffected by the treatment assigned to another unit, meaning that there is no interference between units.

The first part of SUTVA is closely related to the *consistency* assumption (Cole and Frangakis, 2009; VanderWeele, 2009; Pearl, 2010b), which is the fundamental assumption that links the observational world to the potential outcomes in a hypothetical interventional world. The consistency assumption states that if an individual receives treatment level $A = a$, then their observed outcome Y equals their potential outcome under that treatment:

$$Y = Y(a) \quad \text{if } A = a, \quad \text{or equivalently, } Y = Y(A).$$

This identity presumes that the treatment is well-defined and administered uniformly across units. In other words, there are no multiple or ambiguous versions of treatment corresponding to the same value a .

A classic example of an ill-defined treatment arises when the so-called “treatment” is the causal effect of obesity (defined as $\text{BMI} > 30$) on heart disease (Hernán and Robins, 2025).

Remark 4. The consistency assumption discussed here is not the same as the notion of consistency used to describe a statistical property of an estimator (e.g., Lehmann and Casella, 1998), which refers to convergence in probability to the true parameter value.

The second part of SUTVA, namely the no interference assumption, implies that each unit’s potential outcome depends only on its own treatment assignment and not on the treatment received by others. This assumption can be violated in many real-world contexts. For example, in vaccine studies, the likelihood of an individual contracting an infectious disease may depend not only on whether they are vaccinated but also on the vaccination status of others in their community. Such *interference* is common in infectious disease epidemiology and social network settings, where spillover, peer effects, or herd immunity can induce dependencies across units. A growing body of work has extended the potential outcomes framework to accommodate interference. We refer interested readers, for example, to Kuang et al. (2020, § 9) and Bajari et al. (2023) for surveys of this area.

To identify the ACE under the SUTVA, a key condition is the assumption of *no unmeasured confounding*, also known as *conditional (weak) ignorability* or *selection on observables*:

$$A \perp\!\!\!\perp Y(a) \mid L, \quad a = 0, 1. \quad (2.1)$$

This condition implies that, conditional on L , the treatment assignment A is as good as random and thus independent of the potential outcomes. Under this assumption and a positivity condition that ensures all levels of treatment are possible across values of L , the ACE is identified by the following formula:

$$\text{ACE} = \mathbb{E}_L \{ \mathbb{E}(Y \mid A = 1, L) - \mathbb{E}(Y \mid A = 0, L) \}, \quad (2.2)$$

where the outer expectation is over the marginal distribution of L . Equation (2.2) was termed “direct standardization” in the early epidemiological literature (Neison, 1844) and is also a special case of the g-formula of Robins (1986). This expression can be estimated using regression, stratification, matching, inverse probability weighting, or doubly robust methods, depending on the structure and quality of the observed data.

Remark 5. In contrast to (2.2), the association between A and Y can be written as $E(Y \mid A = 1) - E(Y \mid A = 0) = \mathbb{E}_{L|A=1} \mathbb{E}(Y \mid A = 1, L) - \mathbb{E}_{L|A=0} \mathbb{E}(Y \mid A = 0, L)$.

2.2 Structural Equation Models

An NSPEM (Haavelmo, 1943; Strotz and Wold, 1960; Pearl, 2009; Spirtes et al., 2000; Halpern and Pearl, 2000) consists of a finite set of random variables $V = \{V_1, \dots, V_p\}$, a collection of exogenous (unobserved) random variables $\varepsilon = (\varepsilon_1, \dots, \varepsilon_p)$, and a collection of measurable functions f_1, \dots, f_p , such that each endogenous variable V_j is determined by an equation of the form

$$V_j = f_j(U_j, \varepsilon_j), \quad (2.3)$$

where $U_j \subseteq V \setminus \{V_j\}$ denotes a subset of variables that serve as inputs to the structural assignment of V_j , known as the parents of V_j . In this context, the variables in V are referred to as *endogenous* variables because their values are determined within the system by the structural equations. In contrast, the components of ε are called *exogenous* variables, as they represent factors external to the system whose values are not influenced by other variables in the model. These exogenous variables typically encode latent background factors or sources of randomness that affect the system but are not explained by it.

Together, the functions f_j and exogenous variables ε_j specify a data-generating process and thereby induce a joint distribution over V . Importantly, NPSEMs do not need to be acyclic: cycles

or feedback relations can in principle be allowed. When the system is *acyclic*, however, there exists (at least) one ordering $V_{\pi(1)}, \dots, V_{\pi(p)}$ such that each $V_{\pi(j)}$ depends only on variables earlier in the ordering. In this case, the structure of dependence among variables can be represented by a directed acyclic graph (DAG) (see Section 3.2). The model is then fully defined by the tuple $(f_1, \dots, f_p, \varepsilon_1, \dots, \varepsilon_p)$ together with a valid causal ordering (e.g., as obtained from Simon’s causal ordering algorithm (Simon, 1953)). The nonparametric nature of NPSEMs arises from the fact that no parametric form is assumed for the functions f_j , allowing for considerable flexibility in modeling complex causal relationships. A key assumption in many NPSEM applications is that the exogenous variables in ε are jointly independent, yielding the special case known as the NPSEM with Independent Errors (NPSEM-IE).

As an example, in a causal system involving a covariate L , a treatment A , and an outcome Y , an NPSEM may specify the following equations:

$$\begin{aligned} L &= f_L(\varepsilon_L), \\ A &= f_A(L, \varepsilon_A), \\ Y &= f_Y(L, A, \varepsilon_Y), \end{aligned} \tag{2.4}$$

where ε_L , ε_A , and ε_Y are exogenous error terms representing unmeasured factors. In the job-training example, these equations can be interpreted as representing how individual characteristics (L) influence program participation (A) and, in turn, employment outcomes (Y). This formulation expresses the data-generating process in a way that makes causal relations explicit, distinguishing between background factors and the mechanisms that connect them.

Remark 6. The *Causally Interpretable Structured Tree Graph (CISTG)* as detailed as the data (Robins, 1986) is mathematically equivalent to an NPSEM model (with no assumption on the errors); see Equations (3.2) and (3.6) below.

On the surface, the NPSEM in (2.4) appears identical to a nonparametric regression model for (L, A, Y) . However, unlike statistical regression models that describe relationships in the observational world only, what makes NPSEMs *structural* is that they also define what would happen in an interventional world by removing or replacing any subset of structural equations.

The key assumption in the SEM framework that links the observational world to interventional worlds is the *autonomy assumption* (Aldrich, 1989), also known as the *invariance* or *modularity* assumption (Peters et al., 2017). It asserts that each structural equation represents an independent, self-contained causal mechanism that remains invariant under interventions on other variables. Formally, an *intervention* on a variable A corresponds to removing the structural equation for A and replacing it with a constant assignment. This is denoted by Pearl’s *do-operator* (Pearl, 2009), written as $\text{do}(A = a)$. Historically, Robins (1986) used the notation $g = a$ to express the same idea; Pearl’s *do* notation is now more widely used.

In the job-training example, this corresponds to evaluating a hypothetical scenario in which program participation is externally assigned rather than self-selected. The NPSEM in (2.4) implies that in an interventional world where A is fixed to a , the system becomes:

$$\begin{aligned} L &= f_L(\varepsilon_L), \\ A &= a, \\ Y &= f_Y(L, a, \varepsilon_Y). \end{aligned}$$

Under this formulation, the ACE can be represented as

$$\text{ACE} = \mathbb{E}\{Y \mid \text{do}(A = 1)\} - \mathbb{E}\{Y \mid \text{do}(A = 0)\} = \mathbb{E}\{f_Y(L, 1, \varepsilon_Y) - f_Y(L, 0, \varepsilon_Y)\}.$$

Autonomy ensures that interventions can be represented by replacing the structural equation for the intervened variable, without having to re-specify the rest of the system. This feature allows us to predict counterfactual outcomes under hypothetical interventions. In real-world systems, however, autonomy is not always guaranteed (e.g., Cartwright, 2007). Administering the treatment may change the patient’s behavior or alter how the outcome is measured. For instance, it may modify the reporting mechanism or induce side effects that feed back into the outcome-generation process in an unmodeled way. In such cases, the function f_Y may no longer be invariant to interventions on A . In other words, the mapping from (L, A, ε_Y) to Y may differ depending on whether A is set naturally or via intervention.

Remark 7. The term Structural Causal Models (SCMs) is often used in the literature. It may refer to the NPSEM (e.g., Pearl, 2009, p. 203, Definition 7.1.1) alone or specifically to the NPSEM with independent errors (NPSEM-IE) (e.g., Pearl, 2009, p. 44, Definition 2.2.2); see also Peters et al. (2017, p. 62, Definition 6.2). Following Shpitser et al. (2022), we use our terminology here to emphasize the distinction between NPSEM and NPSEM-IE.

Remark 8. The independent error and autonomy assumptions in the NPSEM-IE framework imply the *principle of independent mechanisms (PIM)* (Janzing and Schlölkopf, 2010; Peters et al., 2017). This principle states that the causal generative process can be decomposed into independent modules, one for each variable, that remain invariant to changes in the distributions of other variables. In the special case of two variables, PIM reduces to the so-called *independence of cause and mechanism*, which asserts that the conditional distribution $P(Y | X)$ is independent of the marginal distribution $P(X)$. In other words, under NPSEM-IE assumptions, the causal mechanism mapping $X \mapsto Y$ is invariant to changes in the distribution of X .

2.3 Directed Acyclic Graphs

The use of DAGs for causal inference stems from foundational work by Spirtes et al. (2000) and Pearl (1995), building on earlier ideas from Sewall Wright’s path diagrams (Wright, 1921). In this framework, a causal system is represented by a DAG, where nodes correspond to variables and directed edges encode possible direct causal effects. In particular, the absence of an edge asserts that no direct causal effect is present. In the following, we introduce the DAG model from three perspectives: purely graphical, statistical/probabilistic, and causal.

Directed Acyclic Graph

A *Directed Acyclic Graph* (DAG) is a finite set of nodes connected by directed edges such that no cycles are present. In other words, there is no way to start at a node and return to it by following a sequence of directed edges.

Formally, let $\mathcal{G} = (V, E)$ be a DAG, where

- $V = \{V_1, \dots, V_p\}$ is a set of vertices, and
- $E \subseteq V \times V$ is a set of directed edges.

Figure 1 shows two example DAGs. We next introduce some graph terminology that will be useful later. At this point, no statistical or causal interpretation is attached to the graph.

In a DAG \mathcal{G} , the *parents* of a node V_j , denoted $\text{Pa}_{\mathcal{G}}(V_j)$, or $\text{Pa}(V_j)$ in short, are all nodes with arrows pointing directly into V_j :

$$\text{Pa}(V_j) = \{V_k \in V : (V_k \rightarrow V_j) \in E\}.$$




Figure 1: Two example DAGs over variables L , A , and Y .

For example, in both Figure 1a and Figure 1b, $\text{Pa}(Y) = \{L, A\}$. If V_k is a parent of V_j , then V_j is a *child* of V_k .

A *path* between nodes V_k and V_j is a sequence of distinct nodes

$$V_k - V_{i_1} - V_{i_2} - \dots - V_j$$

such that each consecutive pair is connected by an edge, regardless of its orientation. A *directed path* or *causal path* from V_k to V_j is a path in which all edges point forward along the sequence, for example $V_k \rightarrow V_{i_1} \rightarrow \dots \rightarrow V_j$. Any other path from V_k to V_j is a *non-causal path*. A particularly important class of non-causal paths are the *backdoor paths*, which are paths from V_k to V_j that begin with an arrow into V_k .

The *ancestors* of V_j , denoted $\text{An}(V_j)$, are all nodes from which V_j is reachable by a directed path, with V_j itself included as an ancestor of V_j . Dually, the *descendants* of V_j , denoted $\text{De}(V_j)$, are all nodes reachable from V_j by a directed path, again including V_j itself. The set of *non-descendants* of V_j is

$$\text{ND}(V_j) := V \setminus \text{De}(V_j).$$

For example, in Figure 1a we have, for node L , $\text{An}(L) = \{L\}$, $\text{De}(L) = \{L, A, Y\}$, $\text{ND}(L) = \emptyset$. In the sparser DAG of Figure 1b, for the same node L , $\text{An}(L) = \{L\}$, $\text{De}(L) = \{L, Y\}$, $\text{ND}(L) = \{A\}$.

Lastly, let A , B , C be disjoint subsets of nodes in a DAG \mathcal{G} . A path between a node in A and a node in B is said to be *blocked* by C if it contains a subpath of one of the following forms:

1. A *chain* $X \rightarrow M \rightarrow Y$, $X \leftarrow M \leftarrow Y$ or a *fork* $X \leftarrow M \rightarrow Y$ such that $M \in C$;
2. A *collider* $X \rightarrow M \leftarrow Y$ such that neither M nor any of its descendants is in C .

Some authors refer to the subpaths in 1. as *non-colliders*. If all paths from nodes in A to nodes in B are blocked by C , then A and B are said to be *d-separated* by C (the d denotes *directional* (Pearl et al., 2016, p. 46)). Equivalently, d -separation can be defined using the notion of “moral graph” (Lauritzen et al., 1990). Specifically, consider the subgraph induced by the ancestors of $A \cup B \cup C$. For each pair of non-adjacent nodes that share a common child, add an undirected edge between them (thus “marrying” the parent variables). Then, remove the directions of all edges to obtain an undirected moral graph. We say A and B are d -separated by C in the original DAG if and only if every path between A and B is intercepted by a node in C in the resulting moral graph.

Unlike the concepts of parents and descendants, *d-separation* is a relatively modern development (Pearl and Paz, 1986; Pearl, 1988; Geiger et al., 1990; Pearl, 1995). It enables one to read off conditional independencies from directed graphs in probabilistic DAG models, to which we now turn.

Probabilistic DAG Model

In probabilistic or statistical modeling, a DAG is sometimes referred to as a Bayesian network (Pearl, 1985), and each node in a DAG represents a random variable. Here the DAG is used

to encode the conditional independence structure of a joint distribution over a set of random variables (V_1, \dots, V_p) . Specifically, let P be a joint distribution over (V_1, \dots, V_p) , and let \mathcal{G} be a DAG. We say that P is *Markovian* with respect to \mathcal{G} if it satisfies the following properties.

Definition 1 (Markov Properties).

1. **Markov Factorization Property:** The joint distribution factorizes as:

$$P(V_1, \dots, V_p) = \prod_{j=1}^p P(V_j \mid \text{Pa}_{\mathcal{G}}(V_j)). \quad (2.5)$$

2. **Global Markov Property:** For any disjoint sets $A, B, C \subseteq V$, if C d-separates A and B in \mathcal{G} , then $A \perp\!\!\!\perp B \mid C$ in P .
3. **Local Markov Property:** Each variable is conditionally independent of its non-descendants given its parents: $V_j \perp\!\!\!\perp \text{ND}(V_j) \mid \text{Pa}(V_j)$.

These Markov properties are equivalent if the joint distribution of V has a density with respect to a product measure (Lauritzen, 1996, Theorem 6.22).

Due to the global Markov property, in probabilistic DAG models, d-separation serves as a bridge between the structure of a DAG and the independencies encoded in the corresponding distribution. If the DAG is incomplete (i.e., contains missing edges), Pearl's *d-separation criterion* can be used to identify the conditional independence relationships implied by the factorization.

As an example, consider the DAG shown in Figure 1b. The DAG encodes the following factorization of the joint distribution as in (2.5): $P(L, A, Y) = P(L) \cdot P(A) \cdot P(Y \mid L, A)$. The local Markov property implies that A is independent of its non-descendant L . On the other hand, using d-separation, we see that $A \perp\!\!\!\perp L$, because the only path connecting them, $A \rightarrow Y \leftarrow L$, contains a collider Y .

Remark 9. Different graphs may encode the same set of conditional independence relationships. For example, the graphs $A \rightarrow B \rightarrow C$ and $A \leftarrow B \rightarrow C$ both encode the assumption that $A \perp\!\!\!\perp C \mid B$. In this case, these graphs are called *Markov equivalent*. Graphs that are Markov equivalent cannot be distinguished based solely on conditional independencies in observational data. However, they may be distinguished in the presence of additional assumptions about the joint distribution, such as constraints on the functional form of the conditional distributions. For detailed discussions, see Example 3 in Section 4.2 and Peters et al. (2017, § 4 & § 7).

We note that, in general, the converse of the (global) Markov property does not necessarily hold. That is, a probability distribution P that is Markovian with respect to a DAG may exhibit additional conditional independence relations that are not entailed by d-separation in the DAG. Such situations arise, for example, when special parameter values induce cancellations or deterministic relationships, leading to additional independencies not implied by the graph via d-separation. When the converse does hold, meaning that every conditional independence in P corresponds exactly to a d-separation in \mathcal{G} , we say that the distribution is *faithful* to the graph.

Definition 2 (Faithfulness). Given a DAG \mathcal{G} and a joint probability distribution P over its variables, P is *faithful* to \mathcal{G} if:

$$A \perp\!\!\!\perp_P B \mid C \quad \Leftrightarrow \quad A \text{ and } B \text{ are d-separated by } C \text{ in } \mathcal{G}.$$

Not all probability distributions can be faithfully represented by a DAG. In particular, the graph structure enforces certain logical closure properties for d-separation that need not hold

for conditional independence relations in general distributions. For example, if C d-separates A from B and also A from D , then C must d-separate A from (B, D) in the same DAG. However, this rule does not generally hold for probabilistic conditional independence. A distribution may satisfy both $A \perp\!\!\!\perp_P B \mid C$ and $A \perp\!\!\!\perp_P D \mid C$ without satisfying $A \perp\!\!\!\perp_P (B, D) \mid C$.

It can be shown that the set of distributions unfaithful to a graph \mathcal{G} has Lebesgue measure zero (Spirtes et al., 2000, Theorem 3.2); see also Uhler et al. (2013). This result, however, comes with an important caveat: there exist sequences of faithful distributions that converge arbitrarily close to unfaithful distributions. A more robust alternative is the concept of λ -strong faithfulness, which requires that the dependence strength (e.g., partial correlation) between d-connected variables exceeds some threshold $\lambda > 0$. Unlike exact faithfulness, the set of distributions that are not λ -strong faithful may occupy a substantial volume in the space of distributions (Uhler et al., 2013).

A careful reader might observe that the first paragraph after Definition 2 seems to suggest that, when sampling uniformly over all distributions, the faithfulness assumption fails with probability one. In contrast, the second paragraph implies that faithfulness holds almost surely. To reconcile this apparent contradiction, consider three binary variables A , B , and D , whose joint distribution is specified by probabilities $p(a, b, d) \in [0, 1]$ for $(a, b, d) \in \{0, 1\}^3$, subject to the constraint $\sum_{a,b,d} p(a, b, d) = 1$. The space of all such distributions is the 7-simplex

$$\Delta_7 := \left\{ p \in [0, 1]^8 : \sum_{a,b,d} p(a, b, d) = 1 \right\},$$

which is a 7-dimensional affine subset of \mathbb{R}^8 .

Now consider the following subsets of distributions:

$$\begin{aligned} \mathcal{S}_1 &= \{P_{A,B,D} : A \perp\!\!\!\perp B\}; \\ \mathcal{S}_2 &= \{P_{A,B,D} : A \perp\!\!\!\perp B, \quad A \perp\!\!\!\perp D\}; \\ \mathcal{S}_3 &= \{P_{A,B,D} : A \perp\!\!\!\perp (B, D)\}. \end{aligned}$$

Here \mathcal{S}_1 corresponds to the distributions obeying the Markov property for the DAG $(B \rightarrow D \leftarrow A)$, while \mathcal{S}_3 gives the set of distributions Markov with respect to the DAG $(B \rightarrow D \leftarrow A)$. The sets \mathcal{S}_1 , \mathcal{S}_2 , \mathcal{S}_3 have dimensions 6, 5, and 4, respectively, within the 7-simplex Δ_7 .

Now suppose we sample uniformly from the set \mathcal{S}_1 . Then the subset of such distributions that are unfaithful to the DAG $(B \rightarrow D \leftarrow A)$, such as those in \mathcal{S}_2 , has measure zero. Thus, faithfulness holds with probability one under this sampling scheme.

In contrast, if we sample uniformly from \mathcal{S}_2 , the set of distributions that satisfy $A \perp\!\!\!\perp B$ and $A \perp\!\!\!\perp D$, then with probability zero the resulting distribution will also satisfy $A \perp\!\!\!\perp (B, D)$, i.e., fall into the set \mathcal{S}_3 . Therefore, the distributions in \mathcal{S}_2 cannot be faithfully represented by any DAG almost surely.

Remark 10. Strictly speaking, the phrase “sample uniformly” is not well defined without first specifying a parameterization, a point related to *Borel’s paradox*. Conditional distributions on measure zero sets can depend on the chosen parameterization (Billingsley, 1995, p. 441, Problem 33.1); see also Wang (2022). In particular, subsets such as \mathcal{S}_1 and \mathcal{S}_2 are lower dimensional manifolds within the probability simplex, so there is no canonical uniform measure on them. For \mathcal{S}_1 and the associated DAG $B \rightarrow D \leftarrow A$, we make this precise by drawing each conditional probability table independently from a Dirichlet(1, . . . , 1) prior and combining them using the DAG factorization, under which the set of unfaithful distributions has Lebesgue measure zero.

For \mathcal{S}_2 , we describe in Appendix A.1 a sampling scheme under which a distribution satisfying $A \perp B$ and $A \perp D$ almost surely does not satisfy $A \perp (B, D)$.

Causal DAG Model

So far, we have described how DAGs encode *statistical* conditional independencies via the factorization of the observed data distribution. To endow a DAG with causal meaning, we must go beyond this statistical interpretation and commit to assumptions about how the system behaves under external interventions (Pearl, 2009, Definition 1.3.1); see also Peters et al. (2017, Definition 6.32).

Definition 3 (Causal DAG Model). A causal DAG model over random variables $V = \{V_1, \dots, V_p\}$ consists of a directed acyclic graph \mathcal{G} and a family of interventional distributions $\{P_{\text{do}(v_S)}\}$ such that:

1. (**Observational Markov condition**) The *observational distribution* P factorizes as (2.5).
2. (**Modularity**) For any joint intervention on a set $S \subseteq [p] \equiv \{1, \dots, p\}$ that sets V_S to a fixed value v_S :

$$P_{\text{do}(v_S)}(V_{[p] \setminus S}) = \prod_{k \notin S} P(V_k \mid V_{\text{Pa}_{\mathcal{G}}(V_k) \setminus S}, V_{\text{Pa}_{\mathcal{G}}(V_k) \cap S} = v_{\text{Pa}_{\mathcal{G}}(V_k) \cap S}). \quad (2.6)$$

To illustrate the modularity assumption in Definition 3, consider Figure 1a. When interpreted as a probabilistic DAG, this graph corresponds to a full model on the observed data distribution over (L, A, Y) , since no edges are missing. However, if we interpret the same DAG causally, then the effect of an intervention on A is captured by the truncated DAG in Figure 1b, where all incoming edges into A have been removed. Formally, (2.6) yields

$$P_{\text{do}(a)}(L, Y) = P(L) P(Y \mid A = a, L), \quad (2.7)$$

which can be obtained by removing the term $P(A \mid L)$ from the observational factorization and fixing A at the value a . Marginalizing over L gives

$$P_{\text{do}(a)}(Y) = \sum_{L=l} P(L=l) P(Y \mid A = a, L=l), \quad (2.8)$$

which coincides with the g-formula (2.2). This truncated factorization forms the basis of Pearl's *do-calculus*, introduced shortly before Remark 13.

Remark 11. Eqn. (2.6) was originally introduced by Robins (1986) and termed the g-computation algorithm formula or simply the g-formula. It was subsequently independently re-discovered and named the manipulation theorem (Spirtes et al., 2000) and the truncation formula (Pearl, 2009).

When a causal DAG is postulated, it is often assumed that the set of variables V is *causally sufficient*, in the sense that there are no unmeasured common causes $C \notin V$ that cause two or more variables in V (Spirtes, 2010). This assumption may be slightly relaxed to the notion of *interventional sufficiency*, which is sufficient to ensure that the truncation formula (2.6) holds; see Peters et al. (2017, § 9.1) for a detailed discussion. When unmeasured common causes do exist between variables, it is conventional to represent them by adding bidirected edges (\leftrightarrow) between the affected nodes, resulting in an Acyclic Directed Mixed Graph (ADMG) (Richardson, 2003). Formally, such graphs can be obtained as *latent projections* of DAGs with unobserved variables (Verma and Pearl, 1991; Richardson and Spirtes, 2002).

Definition and Identification of Causal Effects

Under the causal DAG model, the *average causal effect* (ACE) of a binary treatment $A \in \{0, 1\}$ on an outcome Y is defined as:

$$\text{ACE} = \mathbb{E}_{\text{do}(a=1)}(Y) - \mathbb{E}_{\text{do}(a=0)}(Y),$$

where the expectation is taken with respect to the interventional distributions $P_{\text{do}(a)}(Y)$ defined via the truncated factorization in eqn. (2.6).

In general, the interventional quantity $\mathbb{E}_{\text{do}(a)}(Y)$ is not equal to its observational counterpart $\mathbb{E}(Y | A = a)$ because of the presence of confounding. To recover causal effects, we may seek an *adjustment set* L of observed variables such that conditioning on L blocks all non-causal paths from A to Y .

When all variables in the DAG are observed, it is sufficient to include all the parents of A in the adjustment set L , in which case the average causal effect (ACE) is identified by eqn. (2.2) (Pearl, 2009, Theorem 3.2.2). More generally, even if some variables in the DAG are unobserved, the following *backdoor criterion* provides a graphical condition under which causal effects can still be identified from observational data.

A set of variables L satisfies the *backdoor criterion* relative to treatment A and outcome Y if:

1. No node in L is a descendant of A , and
2. L blocks all backdoor paths from A to Y in the DAG.

Under this criterion, the ACE is still identified by eqn. (2.2) (Pearl, 2009, Theorem 3.3.2).

In the job-training example, background factors such as education and motivation (L) may influence both participation in the program (A) and employment outcomes (Y). In a corresponding DAG, these relationships are represented by arrows $L \rightarrow A$ and $L \rightarrow Y$, together with a causal arrow $A \rightarrow Y$; see, for example, Figure 1a. If L blocks all backdoor paths from A to Y and contains no descendants of A , as is the case in Figure 1a, then adjustment for L satisfies the backdoor criterion and identifies the causal effect of program participation.

Remark 12. Although Pearl (2009, § 3.2.1, eqn. (3.2)) interprets the causal semantics of a DAG using the NPSEM-IE framework (see also Section 3.2), the proof of the backdoor criterion (Pearl, 2009, Theorem 3.2.2 and § 11.3.3) relies only on the g-formula (2.6), and therefore holds under the causal DAG model.

The backdoor criterion can be operationalized using the following simple rules (Pearl, 2009, § 11.3.1), such that after conditioning on L ,

Rule 1: All non-causal paths between A and Y are blocked, and no new ones are introduced.

Rule 2: All causal paths from A to Y remain open (unblocked).

The first rule of the backdoor criterion suggests that adjusting for post-treatment variables may introduce bias. To illustrate, consider the causal DAG in Figure 2(a), where the treatment A has two descendants, S and Z , in addition to the primary variables A and Y . Adjusting for S clearly violates Rule 2 above. Now consider adjusting for Z . On the surface, Z does not lie on any path from A to Y , and S is not a collider on the path $A \rightarrow S \rightarrow Y$. However, every observed variable has an associated error term, and in panel (a) we explicitly display ε_S , the error for S , since it plays a key role here. Conditioning on Z induces a dependence between A and ε_S , which in turn opens the non-causal path $A \leftrightarrow \varepsilon_S \rightarrow S \rightarrow Y$. Equivalently, this corresponds to the undirected path $A - \varepsilon_S - S - Y$ in the induced moral graph. Another way to see this is by marginalizing out S and examining the resulting causal DAG in Figure 2(b). In this graph, it



(a) A DAG with post-treatment variables S , Z , and Y .

(b) The same DAG after marginalizing over S .

Figure 2: Left: the original model including the mediator S . Right: the simplified model where S is removed and its influence is captured via direct arrows.

becomes clear that Z is a collider on the non-causal path $A \rightarrow Z \leftarrow \varepsilon_S \rightarrow S$, and hence adjusting for Z may introduce collider bias into the causal effect estimate (Pearl, 2009, § 11.3.1).

When the backdoor criterion fails to apply directly, identification may still be possible through Pearl’s *do-calculus* (Pearl, 1995), a set of three transformation rules that enable the conversion of interventional expressions (e.g., $P_{\text{do}(a)}(Y)$) into estimable quantities from the observational distribution. While the backdoor criterion suffices for simple DAGs, complete graphical criteria for covariate adjustment in ancestral and partially directed graphs have been developed (Perković et al., 2018). In fact, any identifiable interventional distribution under a causal DAG model that may involve unmeasured nodes can be derived through recursive application of the *do-calculus* rules, as formalized by the ID algorithm (Tian and Pearl, 2002; Huang and Valtorta, 2006; Shpitser and Pearl, 2006). See Tian and Shpitser (2010) for a review of these results.

It is important to note that this result pertains to identifiability under the causal DAG model alone. Given additional assumptions, identification may still be possible even when the ID algorithm fails to apply. For instance, one can show that the ACE is not identifiable under the instrumental variable model alone (Balke and Pearl, 1997). However, it can still be identified under additional effect homogeneity assumptions (Wang and Tchetgen Tchetgen, 2018, Theorem 1); see also Hartwig et al. (2023, Theorem 1) and Dong et al. (2025, Theorem 1).

Remark 13. There are more sophisticated algorithms that also identify counterfactual quantities, such as those discussed in Remark 3 (Shpitser and Pearl, 2007, 2008). These algorithms, unlike the ID algorithm mentioned above, rely on the NPSEM-IE framework.

3 Translating Between the Three Causal Frameworks

The question of which causal framework is best suited for practical applications has been the subject of extensive debate in the causal inference literature (e.g. Pearl, 1995; Rubin, 2004; Lauritzen, 2004; Robins and Richardson, 2011; Richardson and Robins, 2023); see also Pearl (2009, p. 106). While each approach, including potential outcomes, structural causal models, and graphical models, offers unique strengths, efforts have been made to reconcile and translate among them. Much like natural languages, certain causal concepts are more naturally expressed in one framework than another, making such translations both valuable when possible but, at

times, inherently limited. In this section, we review key attempts to bridge these perspectives. Detailed comparisons of these frameworks are deferred to Section 4.

3.1 The Potential Outcomes and NPSEM-IE Frameworks

From NPSEM-IE to Potential Outcomes Under an NPSEM-IE, each endogenous variable V_j is determined by a deterministic function of its parents and an exogenous noise variable ε_j , with the collection $\{\varepsilon_j\}$ assumed to be jointly independent. This structural system implicitly defines a rich set of potential outcomes. For example, consider the NPSEM in (2.4). In this case, the potential outcome $Y(l, a)$ is obtained by intervening to set $L = l$, $A = a$, and evaluating the outcome function at the realized value of the noise term ε_Y (Pearl, 2009, § 3.6.3, § 7.1):

$$Y(l, a) = f_Y(l, a, \varepsilon_Y); \quad (3.1)$$

see also Strotz and Wold (1960). In fact, researchers have proposed expressing NPSEMs using potential outcomes notation to clearly distinguish them from standard statistical regression models (e.g. Imbens, 2014; Richardson and Robins, 2014, 2023). For example, the NPSEM in (2.4) can be represented as

$$\begin{aligned} L &= f_L(\varepsilon_L), \\ A(l) &= f_A(l, \varepsilon_A), \\ Y(l, a) &= f_Y(l, a, \varepsilon_Y), \end{aligned} \quad (3.2)$$

where the potential outcomes notation explicitly emphasizes the causal interpretation of the structural equations.

More compactly, we can also define the marginal potential outcome by absorbing L into its realized value under the intervention:

$$Y(a) = Y(L, a) = f_Y(L, a, \varepsilon_Y). \quad (3.3)$$

One can also recursively define nested potential outcomes from NPSEMs. For example, consider a simple model with a treatment A , a mediator M , and an outcome Y , governed by:

$$\begin{aligned} A &= f_A(\varepsilon_A), \\ M &= f_M(A, \varepsilon_M), \\ Y &= f_Y(A, M, \varepsilon_Y). \end{aligned}$$

In this setting, the nested counterfactual $Y(a, M(a'))$ represents the value of Y if treatment were set to a , and the mediator were assigned the value it would attain under treatment a' . This can be expressed as:

$$Y(a, M(a')) = f_Y(a, f_M(a', \varepsilon_M), \varepsilon_Y). \quad (3.4)$$

Implications of Missing Input Variables in NPSEM-IE In the NPSEM defined in (2.3), the set of input variables U_j for a given node V_j often does not include all variables that precede V_j in the causal ordering. The omission of these variables reflects an implicit *exclusion restriction* on the corresponding potential outcomes (Pearl, 2009, p. 101, eqn. (3.55)):

$$V_j(u_j) = V_j(u_j, s),$$

where $S \subset V$ is disjoint from U_j . This equality states that setting variables in S does not affect the value of V_j once the variables in U_j are held fixed, thereby encoding a form of causal irrelevance.

As a concrete example, consider a simple instrumental variable (IV) model with an unmeasured confounder U that affects both the treatment A and the outcome Y . Let Z be an observed instrument that influences A but has no direct effect on Y . The structural equations may be specified as:

$$\begin{aligned} Z &= \varepsilon_Z, \\ U &= \varepsilon_U, \\ A &= f_A(Z, U, \varepsilon_A), \\ Y &= f_Y(A, U, \varepsilon_Y). \end{aligned}$$

In this setup, U introduces unmeasured confounding between A and Y . The exclusion restriction required for IV identification is encoded in the structural equation for Y : the variable Z is omitted from the input set in the SEM for Y . This implies that, for any a and z ,

$$Y(z, a, U) = Y(a, U),$$

or, equivalently,

$$Y(z, a) = Y(a). \quad (3.5)$$

Eqn. (3.5) is often known as the exclusion restriction in the IV literature (Angrist et al., 1996).

From Potential Outcomes to NPSEM-IE via One-Step-Ahead Counterfactuals

Robins and Richardson (2011); Richardson and Robins (2013) formalize the converse direction by showing how one can construct an NPSEM-IE from a given system of potential outcomes using what they term *one-step-ahead counterfactuals*, i.e., the potential outcome of a variable under interventions on all of its parents (Robins, 1986). This “counterfactual interpretation of the error terms in structural equation models” was also noted in Pearl (2009, § 7.4.5) and referred to as the “canonical representation of structural causal models” by Peters et al. (2017, § 3.4).

To illustrate, consider a binary covariate L and a binary treatment variable A . Suppose the available potential outcomes include $A(l)$ and $Y(l, a)$ for $l \in \{0, 1\}$ and $a \in \{0, 1\}$. One can reconstruct the NPSEM as follows:

$$\begin{aligned} \varepsilon_L &= L, & f_L(\varepsilon_L) &= \varepsilon_L, \\ \varepsilon_A &= (A(1), A(0)), & f_A(l, \varepsilon_A) &= \begin{cases} A(1), & \text{if } l = 1, \\ A(0), & \text{if } l = 0, \end{cases} \\ \varepsilon_Y &= (Y(1, 1), Y(1, 0), Y(0, 1), Y(0, 0)), & f_Y(a, l, \varepsilon_Y) &= \begin{cases} Y(1, 1), & \text{if } l = 1, a = 1, \\ Y(1, 0), & \text{if } l = 1, a = 0, \\ Y(0, 1), & \text{if } l = 0, a = 1, \\ Y(0, 0), & \text{if } l = 0, a = 0. \end{cases} \end{aligned} \quad (3.6)$$

In this construction, each exogenous variable encapsulates the relevant potential outcomes for the corresponding endogenous variable, while the structural functions act as selectors based on the realized values of upstream variables.

3.2 NPSEM-IE and Causal DAG

From an NPSEM-IE to a Causal DAG

Consider the NPSEM (2.3). To construct a graph from this NPSEM, one proceeds as follows: (1) Draw a node for each endogenous variable V_j ; (2) For each variable V_j , draw directed edges from every variable that appears as an argument in f_j to V_j . These variables are the parents of V_j in the graph. If the resulting graph is acyclic, i.e., it forms a DAG, then the corresponding NPSEM is referred to as an *acyclic NPSEM*. In other words, any acyclic NPSEM induces a DAG in a purely graphical sense. The following proposition establishes the connection between the independent error and autonomy assumptions in an NPSEM-IE and the causal DAG model.

Proposition 1. *Any NPSEM-IE implies a unique causal DAG model. In particular, it induces (1) an observational distribution that factorizes according to the DAG induced by the NPSEM; (2) interventional distributions that satisfy the truncated factorization formula (2.6).*

The first claim in Proposition 1 follows from Pearl (2009, Theorem 1.4.1); see also Peters et al. (2017, Proposition 6.31). The second claim holds because, under the autonomy assumption, the structural equation for the intervened variable is replaced while all others remain unchanged; see also Strotz and Wold (1960, § 2), Pearl (2009, p. 72 eqn. (3.10)) and Peters et al. (2017, p. 109, eqn. (6.7)).

Remark 14. In fact, a stronger claim holds: the Single-World Intervention Graph (SWIG) model (see § 3.3 below) with interventions on all variables also implies a causal DAG model (Richardson and Robins, 2023, Theorem 11). This is stronger because the NPSEM-IE is a strict submodel of the SWIG model associated with the same DAG. The converse is not necessarily true. For example, the SWIG in Figure 3 implies the so-called extended g-formula (Robins et al., 2004) that

$$P(A = a^*, L = l, Y(a) = y) = P(L = l)P(A = a^* | L = l)P(Y = y | A = a, L = l) \quad (3.7)$$

for any values of a^* , a , l , and y . Although this implies the truncated factorization formula (2.7), the converse does not hold.

Remark 15. The extended g-formula (3.7) can be used to identify counterfactual quantities such as the ETT. As noted later in Section 4.2, the ETT is not expressible in the causal DAG framework.

From Causal DAGs to NPSEM-IE

Let \mathcal{G} be a causal DAG over variables $V = \{V_1, \dots, V_p\}$, and let $P(V_1, \dots, V_p)$ be a joint distribution that satisfies the Markov factorization with respect to \mathcal{G} :

$$P(V_1, \dots, V_p) = \prod_{j=1}^p P(V_j | \text{Pa}_{\mathcal{G}}(V_j)).$$




Figure 3: The single world intervention graph (template) corresponding to intervening on A in the causal DAG in Figure 1a.

The Functional Representation Lemma (El Gamal and Kim, 2011, p. 626, Appendix B) ensures that any random variable with a known conditional distribution can be expressed as a measurable function of its parents and an independent noise variable. A concrete construction is provided in Peters et al. (2017, Appendix C.9).

Applying the functional representation lemma recursively to each conditional distribution $P(V_j \mid \text{Pa}_{\mathcal{G}}(V_j))$, we obtain an NPSEM:

$$V_j = f_j(\text{Pa}_{\mathcal{G}}(V_j), \varepsilon_j), \quad j = 1, \dots, p, \quad (3.8)$$

where each f_j is a measurable function, each $\varepsilon_j \sim \text{Uniform}[0, 1]$, and the noise variables $\varepsilon_1, \dots, \varepsilon_p$ are jointly independent. The NPSEM (3.8) with independent errors is compatible with both the causal DAG \mathcal{G} and the joint distribution P . A similar construction for the discrete case can be found in Druzdzel and Simon (1993).

Proposition 2. *Consider a causal DAG model over random variables $V = \{V_1, \dots, V_p\}$ as in Definition 3. Then there exists an NPSEM-IE that implies the causal DAG model.*

Remark 16. Although this construction is always possible under the assumptions of the functional representation lemma, the resulting NPSEM-IE is generally not unique, as different (f_j, ε_j) pairs may induce identical observational distributions and factorizations.

3.3 Potential Outcomes and DAGs: Single World Intervention Graphs (SWIGs)

The potential outcomes and DAG frameworks can be unified through *single world intervention graphs* (SWIGs) (Richardson and Robins, 2013). SWIGs explicitly encode potential outcomes on the graph by splitting nodes into fixed (intervened) and random components. This representation enables two key features: (1) conditional independence statements involving both observed and potential outcomes can be read using a modified d-separation rule, and (2) identification results equivalent to do-calculus can be derived from these conditional independence conditions via a simple reformulation known as the potential outcome (po) calculus (Malinsky et al., 2019; Shpitser et al., 2022). This reformulation shows that from the perspective of potential outcomes, the po-calculus (and hence the do-calculus) is reducible to conditional independence, recursive substitution and causes only influence descendants in the SWIG. For completeness, we note that SWIGs can also be defined via a local Markov property with respect to a total ordering of the variables, without explicit reference to a graphical representation (see Richardson and Robins, 2023). The same potential outcome model was originally termed the Finest Fully Randomized Causally Interpretable Structured Tree Graph (FFRCISTG); see Robins (1986) and Richardson and Robins (2013, Appendix C); note that the “tree graphs” referred to here are “event trees,” not DAGs.

For example, consider intervening on the variable A in the DAG presented in Figure 1a. As illustrated in Figure 3, a SWIG can be constructed in two steps. In the first step, one splits each intervention node, in this case, A , into two nodes: a random node A , which inherits all incoming arrows from the original DAG, and a fixed node a , which inherits all outgoing arrows from the original DAG. In the second step, all descendants of the fixed node a are relabeled as potential outcomes, expressed as functions of the fixed node’s value. Formally, the graph in Figure 3 is a *template* (or a graph-valued function), since the fixed node a can take multiple values, such as $a = 0$ or $a = 1$.




Figure 4: Does the causal DAG model imply $Y(a, b) \perp\!\!\!\perp B \mid Z, A = a$?

One can read off conditional independence relationships from a SWIG using the d-separation rule, with the extension that a d-connecting path must *not* contain fixed (red) nodes. For example, from the SWIG in Figure 3, one can directly read off the weak ignorability condition (2.1) that $A \perp\!\!\!\perp Y(a) \mid L$, for $a = 0, 1$ using the rule of d-separation.

Graphs like the one in Figure 3 are called *single-world* because potential outcomes corresponding to *different* hypothetical interventions on the same variable never appear in the same SWIG. Consequently, a SWIG does not encode assumptions about relationships between potential outcomes across different interventions on the same variable. For example, the SWIG in Figure 3 does *not* imply the strong ignorability condition $A \perp\!\!\!\perp (Y(1), Y(0)) \mid L$. We discuss the distinction between single-world and cross-world assumptions in more detail in Section 4.2.

SWIGs are useful not only for validating well-known conditional independence assumptions that hold under the causal DAG model, but also as graphical tools for assessing conditional independence statements that are not obvious from a causal DAG model. For example, consider the conditional independence statement

$$Y(a, b) \perp\!\!\!\perp B \mid Z, A = a. \quad (3.9)$$

Pearl (2009, Ex. 11.3.3) claimed that the causal DAG model in Figure 4a does not imply (3.9). However, as can be seen from the SWIG in Figure 4b, the relationship $Y(a, b) \perp\!\!\!\perp B(a) \mid Z(a), A = a$ holds, which implies (3.9) through the consistency assumption.

4 Comparison Between the Three Frameworks

4.1 Minimalist vs. Mechanistic Perspectives on Causality

Although, as we have shown, these frameworks often lead to overlapping mathematical results, they arise from different philosophical traditions and emphasize different aspects of the causal enterprise (Pearl, 2009, § 1.4). The potential outcomes framework adopts a *minimalist* or design-based view of causality, focusing on comparisons between hypothetical outcomes under different treatment conditions, such as $Y(1) - Y(0)$. This approach does not require specifying how all outcomes are generated from other variables in the system. It is particularly attractive in settings where causal identification relies on features of the study design, such as randomization or the use of instrumental variables, rather than on a full model of the data-generating process. Because of its minimal assumptions and broad applicability, this framework has been described by some authors as the most general approach to causal inference (Imbens and Rubin, 2015; Rosenbaum, 2010).

In contrast, NPSEM-IEs adopt a more *mechanistic* perspective, rooted in classical physics and engineering, where causality is represented through deterministic functional relationships between variables and exogenous noise terms. In this tradition, causal systems are fully specified via structural equations, and randomness reflects ignorance about latent factors rather than intrinsic stochasticity. This view aligns with Laplace’s deterministic vision of natural laws (Laplace, 1825) and provides a rich language for reasoning about complex causal systems.

Causal DAGs provide a bridge between this minimalist perspective and mechanistic modeling. By encoding assumptions about independencies, DAGs help justify potential-outcomes-based identifiability conditions (e.g., ignorability) while avoiding over-specification of functional forms. For instance, a DAG makes it transparent why randomization implies that no arrows point into A . This, in turn, ensures the independence $A \perp\!\!\!\perp Y(a)$, $a \in \{0, 1\}$, which can be read off from a SWIG intervening on A . Moreover, DAGs clarify backdoor adjustment sets for observational studies (Pearl, 2009). In contrast to NPSEMs, all relationships in causal DAG models are formulated as inherently stochastic. As Pearl (2009, p. 26) observed, this reflects the modern conception of physics, sometimes described as quantum mechanical, in which probabilistic laws replace strict determinism as a more realistic description of physical processes. At the same time, DAGs retain the ability to articulate complex structural relationships, making them a powerful tool for both identification and hypothesis generation.

As a caveat, we note that although Pearl’s analogy draws on the idea that modern physics views natural laws as fundamentally probabilistic rather than strictly deterministic, it should be interpreted only metaphorically: causal DAGs are classical probabilistic models that admit hidden variable representations, whereas quantum systems in general do not. Quantum phenomena require richer nested probabilistic frameworks, sometimes viewed as supermodels of the classical theory, that extend beyond the scope of standard DAGs (Richardson et al., 2023).

In what follows, we compare these three frameworks in greater depth, highlighting their respective assumptions, strengths, and areas of applicability.

4.2 Single-World vs. Cross-World Causal Inference

Difference in Expressive Power for Counterfactual Quantities The key difference between the causal DAG framework with its associated “do” notation and the other two frameworks (NPSEM-IE and potential outcomes) is that the former does not provide a way to directly express counterfactual quantities that condition on observed variables rather than interventions, such as $P(Y(a) \mid A = a')$ (cf. Pearl, 2009, § 1.4.4). An important example is the ETT, $\mathbb{E}[Y(1) - Y(0) \mid A = 1]$, which is not expressible in the “do” notation. Compared with the causal DAG model, the NPSEM-IE model also imposes many more assumptions due to the additional cross-world assumptions. We now elaborate this point.

Implications of the Independent Error Assumptions in NPSEM-IE Because the independent error assumptions in NPSEM, when taken at face value, resemble the independent error assumptions commonly made in statistical modeling, they are often perceived as mild by statisticians who are not familiar with the causal modeling framework. Interestingly, many of the same statisticians would view the assumption of no unmeasured confounding as an extremely strong assumption. Before the recent evolution in causal inference, it was standard practice in statistics to avoid causal claims altogether, precisely because of the concern that confounding renders correlation insufficient for establishing causation (Pearl et al., 2016, Preface).

However, as we demonstrate below, the independent error assumptions in an NPSEM-IE are, in fact, much stronger than the assumption of no unmeasured confounding as in eqn. (2.1).

Consider the NPSEM-IE in eqn. (2.4). The independent error assumption states that $\varepsilon_L \perp\!\!\!\perp \varepsilon_A \perp\!\!\!\perp \varepsilon_Y$. Now consider the canonical representation in eqn. (3.6). Under this construction, the joint independence of the exogenous variables implies

$$L \perp\!\!\!\perp (A(l) : l \in \mathcal{L}) \perp\!\!\!\perp (Y(l, a) : a \in \{0, 1\}, l \in \mathcal{L}), \quad (4.1)$$

where \mathcal{L} denotes the support of L (Pearl, 2009, p. 101 footnote 14). Basic algebra then yields the strong ignorability condition,

$$A \perp\!\!\!\perp (Y(1), Y(0)) \mid L, \quad (4.2)$$

which is strictly stronger than the weaker no-unmeasured-confounding assumption in eqn. (2.1).

In fact, we have the following general result due to Pearl (2009, p. 101, eqn. (3.56) & footnote 14).

Proposition 3. *In NPSEM (2.3), joint independence of the p exogenous errors ε_j , $j = 1, \dots, p$ is equivalent to joint independence of the following p collections of potential outcomes $\{V_j(\text{pa}(V_j) = a) : a \in \mathcal{A}_{\text{pa}(V_j)}\}$, $j = 1, \dots, p$, where $\mathcal{A}_{\text{pa}(V_j)}$ denotes the set of all possible assignments to the parents of V_j .*

Note that joint independence of the exogenous errors does *not* imply that the variables $V_j(\text{pa}(V_j) = a)$, $a \in \mathcal{A}_{\text{pa}(V_j)}$ are independent of one another; thus for example this does not imply $Y(a = 0) \perp\!\!\!\perp Y(a = 1)$.

To understand the number of extra assumptions imposed by the NPSEM-IE, consider again the NPSEM-IE in eqn. (2.4) and a simple case where L , A , and Y are all binary. A full model for the joint distribution of $(L, A(l), Y(l, a))$, for $a \in \{0, 1\}$ and $l \in \{0, 1\}$, has dimension $2^{1+2+4} - 1 = 127$. The no-unmeasured-confounding assumption in eqn. (2.1) reduces this to a 123-dimensional model. In contrast, the strong ignorability assumption in eqn. (4.2) further reduces the model to 121 dimensions, and the independent error assumption of the NPSEM-IE implies a model of only 19 dimensions! In other words, NPSEM-IE imposes an additional $123 - 19 = 104$ constraints, all of which are not necessary for the identification of the ACE (Robins and Richardson, 2011). Most of these constraints, such as $A(l = 1) \perp\!\!\!\perp Y(l = 0, a)$, are examples of *cross-world assumptions*, which remain untestable even under randomized treatment assignment. These assumptions violate the principle of “no causation without manipulation” (Holland, 1986) and undermine the connection between experimentation and causal inference (Robins, 2003; Robins and Richardson, 2011).

When the NPSEM-IE Framework Leads to Stronger Identification Results In some cases, however, the formulation and additional assumptions in NPSEM-IE indeed lead to stronger identification results. Here, we provide several examples, each representing a different scenario in which these stronger results hold under the NPSEM-IE but not under alternatives such as the causal DAG model.

Example 1: Mediation Formula In some cases, the additional independence assumptions in NPSEM-IE lead to new identification results. For example, Figure 5 shows the mediation setting with A the treatment, M the mediator, and Y the outcome of interest. An estimand of interest in




Figure 5: Mediation model used to derive the mediation formula.

this setting is known as the natural (or pure) direct effect (Robins and Greenland, 1992), defined as $NDE = E[Y(1, M(0)) - Y(0)]$. The NDE can be identified under the NPSEM-IE using the mediation formula (Pearl, 2001)

$$\mathbb{E}[Y(1, M(0))] = \sum_m \mathbb{E}[Y \mid A = 1, M = m] \cdot P(M = m \mid A = 0).$$

Specifically, it relies on the cross-world assumption that $M(0) \perp\!\!\!\perp Y(1, m)$, which holds under the NPSEM-IE but not the causal DAG or SWIG assumptions.

Example 2: Monotonicity Assumption Some assumptions, such as the monotonicity assumption that $Y(a = 0) \leq Y(a = 1)$ almost surely, cannot be expressed directly using a stochastic model like a causal DAG. Assumptions of this kind often lead to stronger identification results. For example, under the no unmeasured confounding assumption (2.1) and the monotonicity assumption, the probability of necessity and sufficiency is identified as

$$PNS = P(Y = 1 \mid A = 1) - P(Y = 1 \mid A = 0).$$

Other examples include the identification of the local average treatment effect in instrumental variable models (Angrist et al., 1996), tighter bounds on natural indirect effects in mediation analysis without assuming cross-world independencies (Imai et al., 2010), and the identification of, or tighter bounds on, the survivor average treatment effect (Ding et al., 2011; Wang et al., 2017a,b).

Example 3: Additive Noise Models In addition to strengthening identification of causal effects, NPSEM-IE assumptions can also aid in identifying the causal graph structure itself. This connects to our earlier discussion of Markov equivalence in Remark 9, where multiple DAGs may encode the same set of conditional independencies. In many cases, however, restricting the function class in NPSEM-IE resolves this ambiguity and leads to identifiability of the underlying causal structure. These assumptions are more naturally expressed in the language of NPSEMs. One such assumption is that the error terms are additive. Under this condition, many results have been established that lead to identification of the causal system in different scenarios. Below, we present several examples assuming an additive model for (2.3):

$$V_j = f_j(U_j) + \varepsilon_j, \quad (4.3)$$

where U_j are the parents of V_j , and ε_j are jointly independent.

1. **(Linear, non-Gaussian)** Shimizu et al. (2006) showed that if the functions f_j in the NPSEM-IE (4.3) are linear in U_j and ε_j with nonzero coefficients, and if the error terms ε_j are non-Gaussian with strictly positive density, then the underlying graph structure is identifiable from the observed data distribution. This model is known as the *linear non-Gaussian acyclic model (LiNGAM)*.

2. **(Nonlinear)** Consider the bivariate causal model

$$Y = f_Y(X) + \varepsilon_Y, \quad X \perp\!\!\!\perp \varepsilon_Y. \quad (4.4)$$

Hoyer et al. (2008) showed that, except in “rare” cases, a distribution does not admit an additive noise model structure in both directions simultaneously; see also Peters et al. (2017, Theorem 4.5). In fact, if the noise ε_Y is Gaussian, then only *linear* functions allow for additive noise models in both directions (Hoyer et al., 2008, Corollary 1). More broadly, Peters et al. (2014) showed that if the functions f_j in the NPSEM-IE (2.3) are three times differentiable and nonlinear in any component of U_j (see Peters et al., 2014, Theorem 7.7 for a precise definition), then the corresponding graph structure is identifiable from the observed data distribution. These results extend to the so-called *post-nonlinear models* of the form $V_j = g_j(f_j(U_j) + \varepsilon_j)$. These models also generally exist in at most one direction, except in rare cases (Zhang and Hyvärinen, 2009).

Identification of the causal structure in linear models with Gaussian additive noise is more challenging. For example, consider again the bivariate causal model in (4.4). If $f_Y(X)$ is linear in X , then there exist β and ε_X such that $X = \beta Y + \varepsilon_X$ and $Y \perp\!\!\!\perp \varepsilon_X$ if and only if both ε_Y and X are Gaussian (Peters et al., 2017, Theorem 4.2). More generally, identification of the causal graph under linear models with Gaussian additive noise is possible if the error terms ε_j are i.i.d., that is, the noise variances are constant across j (Peters and Bühlmann, 2014). However, the equal variance assumption is not robust to rescaling of the variables.

We conclude this example with a caveat: in the nonlinear case, additive noise models are generally not closed under marginalization, even when marginalizing over unobserved variables. For instance, if $Y = f_Y(X) + \varepsilon_Y$ and $Z = f_Z(Y) + \varepsilon_Z$, then the marginal distribution $P_{X,Z}$ does not necessarily admit an additive noise model from X to Z . This limitation restricts the applicability of additive noise models, as in most, if not all, scenarios there are intermediate variables on a causal path that are unobserved (Peters et al., 2017, § 7.1.2). Therefore, additive models should perhaps be viewed as approximations to the true data-generating process rather than exact representations.

5 Summary and Practical Recommendations

This paper has reviewed three major frameworks for causal inference: the potential outcomes framework, NPSEMs, and DAGs. Each offers a distinct perspective on how to represent, interpret, and identify causal relationships, and each has its own strengths and limitations.

The potential outcomes framework provides a conceptually simple and widely applicable language for defining causal effects, such as average treatment effects and conditional effects. It is especially well suited for settings where study design plays a central role, such as randomized trials and quasi-experiments. Its mathematical simplicity makes it attractive for formulating identification strategies. However, practitioners sometimes posit counterfactual independence assumptions without providing a substantive justification or circumstances when such assumptions might be likely to hold. These shortcomings can often be addressed using the SWIG/FFRCISTG framework, central to which is the notion of one-step ahead counterfactuals and recursive substitution, which provide a stepwise data-generating process. In the case where there are well-defined interventions on every variable, this leads to a potential outcome model that is isomorphic to an NPSEM.

The NPSEM framework formalizes causal systems using deterministic functional relationships and exogenous error terms. This allows for a generative view of how outcomes arise from

their causes, enabling researchers to articulate and evaluate identification assumptions in terms of underlying data-generating mechanisms. As emphasized by Pearl (2009, p. 244), structural models and their associated graphs are particularly useful as tools for expressing assumptions about cause-effect relationships, as they provide a natural and intuitive representation of mechanistic knowledge. Counterfactual independence assumptions that may appear opaque or difficult to interpret when stated in the potential outcomes notation, such as the ignorability assumption, can often be given immediate and process-based interpretations when expressed structurally. The NPSEM framework also facilitates encoding assumptions on the functional form or even parameter values of the structural functions. However, when paired with the assumption of independent errors (NPSEM-IE), this framework often imposes strong and untestable cross-world assumptions that go beyond what is needed for identification and remain untestable even under randomized treatment assignment. These added assumptions can constrain the space of admissible models and lead to overly rigid inferences.

The causal DAG framework provides a graphical language for encoding conditional independence assumptions and visualizing the structure of confounding, mediation, and other causal pathways. It facilitates identification via graphical criteria such as the backdoor condition, and its connection to probabilistic DAGs supports a formal link between graphical structure and observed data. Unlike NPSEM-IE, causal DAGs (like SWIGs) avoid strong cross-world assumptions and are therefore more flexible in some contexts. On the other hand, DAGs do not inherently express counterfactuals, and certain causal quantities, especially those involving nested or cross-world counterfactuals, must be defined outside the graphical formalism or with added assumptions.

Taken together, these frameworks offer complementary tools for causal inference. A practical and robust workflow often benefits from combining them:

1. Define the causal estimand in potential outcomes notation to clarify the scientific question. This step requires basic assumptions such as the SUTVA, but does not otherwise rely on parametric modeling assumptions.
2. Express and evaluate assumptions using DAGs, SWIGs or NPSEMs to articulate structural knowledge and identify valid adjustment strategies.
3. Establish identification using graphical criteria such as d-separation, and algebraic tools like the g-formula, do-calculus, po-calculus or ID algorithm.
4. Choose an estimation method appropriate to the structure and available data, drawing on statistical tools such as regression, weighting, or machine learning methods.

In this workflow, structural assumptions are naturally expressed in graphical or NPSEM terms, then translated into potential outcomes notation, and finally subjected to algebraic derivation (Pearl, 2009, p. 245) to isolate those that are strictly necessary for identification. In this way, NPSEMs and their associated graphs (see Section 3.2 for details) serve as tools for evaluating the plausibility or implausibility of conditional independence assumptions stated in terms of potential outcomes. Contemporary implementations such as the target trial emulation framework (Hernán and Robins, 2016; Hernán et al., 2022) illustrate how structured causal roadmaps can help prevent common inferential pitfalls in observational studies. We refer readers to Wang et al. (2017b); Wang and Tchetgen Tchetgen (2018); Yang et al. (2019); Zhou et al. (2024); Dong et al. (2025) for examples that follow this workflow.

We now illustrate these steps in detail using Wang and Tchetgen Tchetgen (2018). Their study begins by defining the causal estimand, the ACE, in potential outcomes notation, thereby clarifying the scientific target of inference (Step 1). They then express structural assumptions through a DAG and a corresponding SWIG with bidirected arrows, explicitly distinguishing

identification assumptions such as the exclusion restriction and independence from modeling assumptions (Step 2). Using algebraic derivations, they establish identification of the ACE under new no-interaction assumptions (A5.a or A5.b), showing that the estimand can be represented as the average Wald estimand (Step 3). Finally, they develop a series of bounded, efficient, and multiply robust estimators that remain consistent if any one of several working models is correctly specified (Step 4). This example demonstrates how potential outcomes, graphical models, and semiparametric theory can be systematically combined into a coherent causal workflow.

Recent advances continue to expand the reach of causal inference beyond traditional domains. Active research areas include causal discovery algorithms for learning causal structure from data (e.g., Spirtes et al., 2000; Peters et al., 2017), fairness in artificial intelligence (e.g., Kusner et al., 2017), and methods for robust learning and generalization across heterogeneous environments (e.g., Peters et al., 2016). Related ideas have also influenced reinforcement learning and decision-making in artificial intelligence, where causal reasoning is increasingly recognized as essential for interpretability and transferability. A full treatment of these topics is beyond the scope of this review, but they exemplify how causal thinking continues to shape contemporary research at the intersection of statistics, computer science, and the social and health sciences.

A Appendix

A.1 Constructing Priors and “Sampling Uniformly” on \mathcal{S}_2

We now provide a principled way to “sample uniformly” from the space

$$\mathcal{S}_2 = \{P : A \perp\!\!\!\perp B, A \perp\!\!\!\perp D\}.$$

The strategy is to express \mathcal{S}_2 using a set of variation independent parameters. This creates a simple product space on which a prior can be placed, and the resulting measure pushes forward to a well-defined prior on \mathcal{S}_2 .

Throughout we write

$$p_a^A = P(A = a), \quad p_b^B = P(B = b), \quad p_d^D = P(D = d), \quad a, b, d \in \{0, 1\}.$$

The assumptions $A \perp\!\!\!\perp B$ and $A \perp\!\!\!\perp D$ imply that the conditional marginals of B and D do not depend on A , so

$$P(B = b \mid A = a) = p_b^B, \quad P(D = d \mid A = a) = p_d^D$$

for all a . What remains is to characterize the association between B and D within each stratum of A , while preserving these common one way marginals.

For binary variables, it is well known that a 2×2 table with fixed marginals has exactly one degree of freedom. Any strictly positive such table can therefore be uniquely represented by a single scalar measuring the strength of association. A standard choice for this scalar is the conditional odds ratio (Chen, 2007; Tchetgen Tchetgen et al., 2018):

$$\text{OR}(B, D \mid A = a) = \frac{P(B = 1, D = 1 \mid A = a) P(B = 0, D = 0 \mid A = a)}{P(B = 1, D = 0 \mid A = a) P(B = 0, D = 1 \mid A = a)}.$$

The odds ratio is variation independent of the one way marginals of B and D , and under $A \perp\!\!\!\perp B$ and $A \perp\!\!\!\perp D$, these conditional marginals coincide with the unconditional marginals. Thus, specifying $\text{OR}(B, D \mid A = a)$ together with (p_1^B, p_1^D) fully determines the entire conditional 2×2 table in stratum $A = a$.

These observations motivate parameterizing \mathcal{S}_2 by the three free marginal probabilities $(p_1^A, p_1^B, p_1^D) \in (0, 1)^3$ together with the two stratum specific odds ratios

$$\psi_a = \text{OR}(B, D \mid A = a) \in (0, \infty), \quad a = 0, 1.$$

The five parameters $(p_1^A, p_1^B, p_1^D, \psi_0, \psi_1)$ vary freely over the Cartesian product $(0, 1)^3 \times (0, \infty)^2$ and uniquely determine a strictly positive distribution in \mathcal{S}_2 .

To construct a prior on \mathcal{S}_2 , one may place independent Dirichlet(1, 1) priors on (p_0^A, p_1^A) , (p_0^B, p_1^B) and (p_0^D, p_1^D) , together with independent priors on the log odds ratios, for example

$$\log \psi_0, \log \psi_1 \sim \text{Normal}(0, 1).$$

This induces a smooth prior with full support on the strictly positive interior of \mathcal{S}_2 and provides a concrete interpretation of “sampling uniformly” over \mathcal{S}_2 .

Acknowledgement

We are grateful to the Editor for encouraging the development of this work, and to Forrest Crawford and Sander Greenland for insightful discussions that helped shape several key aspects. We acknowledge the use of OpenAI’s ChatGPT for proofreading and refining the writing.

References

- Aldrich J (1989). Autonomy. *Oxford Economic Papers*, 41(1): 15–34. <https://doi.org/10.1093/oxfordjournals.oep.a041889>
- Angrist JD, Imbens GW, Rubin DB (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434): 444–455. <https://doi.org/10.1080/01621459.1996.10476902>
- Arjovsky M, Bottou L, Gulrajani I, Lopez-Paz D (2020). Invariant risk minimization. arXiv preprint: <https://arxiv.org/abs/1907.02893>, Original version 2019.
- Bajari P, Burdick B, Imbens GW, Masoero L, McQueen J, Richardson TS, et al. (2023). Experimental design in marketplaces. *Statistical Science*, 38(3): 458–476. <https://doi.org/10.1214/23-STS883>
- Balke A, Pearl J (1997). Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439): 1171–1176. <https://doi.org/10.1080/01621459.1997.10474074>
- Billingsley P (1995). *Probability and Measure*. Wiley, New York, 3rd edition.
- Bühlmann P (2020). Invariance, causality and robustness. *Statistical Science*, 35(3): 404–426.
- Cartwright N (2007). *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. Cambridge University Press, Cambridge, UK.
- Chen HY (2007). A semiparametric odds ratio model for measuring association. *Biometrics*, 63(2): 413–421. <https://doi.org/10.1111/j.1541-0420.2006.00701.x>
- Cole SR, Frangakis CE (2009). The consistency statement in causal inference: A definition or an assumption? *Epidemiology*, 20(1): 3–5. <https://doi.org/10.1097/EDE.0b013e3181818ef366>
- Dawid AP (2000). Causal inference without counterfactuals. *Journal of the American Statistical Association*, 95(450): 407–424. <https://doi.org/10.1080/01621459.2000.10474210>

- Dawid AP (2015). Statistical causality from a decision-theoretic perspective. *Annual Review of Statistics and Its Application*, 2(1): 273–303. <https://doi.org/10.1146/annurev-statistics-010814-020105>
- Ding P, Geng Z, Yan W, Zhou X-H (2011). Identifiability and estimation of causal effects by principal stratification with outcomes truncated by death. *Journal of the American Statistical Association*, 106(496): 1578–1591. <https://doi.org/10.1198/jasa.2011.tm10265>
- Dong M, Liu L, Tang D, Liu G, Xu W, Wang L (2025). Marginal causal effect estimation with continuous instrumental variables. arXiv preprint: <https://arxiv.org/abs/2510.14368>.
- Druzdzel MJ, Simon HA (1993). Causality in Bayesian belief networks. In: *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence (UAI-93)* (D Heckerman, E Mamdani, eds.), 3–11. Morgan Kaufmann, San Francisco, CA.
- El Gamal A, Kim Y-H (2011). *Network Information Theory*. Cambridge University Press.
- Geiger D, Verma T, Pearl J (1990). Identifying independence in Bayesian networks. *Networks*, 20(5): 507–534. <https://doi.org/10.1002/net.3230200504>
- Greenland S (2004). An overview of methods for causal inference from observational studies. In: *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives: An Essential Journey with Donald Rubin's Statistical Family* (A Gelman, X-L Meng, eds.), pages 1–13. John Wiley & Sons, Chichester, UK.
- Greenland S, Brumback B (2002). An overview of relations among causal modelling methods. *International Journal of Epidemiology*, 31(5): 1030–1037. <https://doi.org/10.1093/ije/31.5.1030>
- Greenland S, Poole C (1988). Invariants and noninvariants in the concept of interdependent effects. *Scandinavian Journal of Work, Environment & Health*, 14(2): 125–129. <https://doi.org/10.5271/sjweh.1945>
- Haavelmo T (1943). The statistical implications of a system of simultaneous equations. *Econometrica*, 11(1): 1–12. <https://doi.org/10.2307/1905714>
- Halpern JY, Pearl J (2000). Axiomatizing causal reasoning. *Journal of Artificial Intelligence Research*, 12: 317–337. <https://doi.org/10.1613/jair.648>
- Hartwig FP, Wang L, Smith GD, Davies NM (2023). Average causal effect estimation via instrumental variables: The no simultaneous heterogeneity assumption. *Epidemiology*, 34(3): 325–332. <https://doi.org/10.1097/EDE.0000000000001596>
- Hernán MA, Robins JM (2016). Using big data to emulate a target trial when a randomized trial is not available. *American Journal of Epidemiology*, 183(8): 758–764. <https://doi.org/10.1093/aje/kwv254>
- Hernán MA, Robins JM (2025). *Causal Inference: What If*. Chapman & Hall/CRC, Boca Raton, FL, 1st edition.
- Hernán MA, Wang W, Leaf DE (2022). Target trial emulation: a framework for causal inference from observational data. *Journal of the American Medical Association*, 328(24): 2446–2447. <https://doi.org/10.1001/jama.2022.21383>
- Hill AB (1965). The environment and disease: Association or causation? *Proceedings of the Royal Society of Medicine*, 58: 295–300. <https://doi.org/10.1177/003591576505800503>
- Holland PW (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81(396): 945–960. <https://doi.org/10.1080/01621459.1986.10478354>
- Hoyer P, Janzing D, Mooij JM, Peters J, Schölkopf B (2008). Nonlinear causal discovery with additive noise models. In: *Advances in Neural Information Processing Systems* (D Koller, D Schuurmans, Y Bengio, L Bottou, eds.), volume 21. Curran Associates, Inc.
- Huang Y, Valtorta M (2006). Pearl's calculus of intervention is complete. In: *Proceedings of the*

- Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI-2006)* (R Dechter, T Richardson, eds.), 217–224. AUAI Press, Arlington, VA.
- Imai K, Keele L, Yamamoto T (2010). Identification, inference and sensitivity analysis for causal mediation effects. *Statistical Science*, 25(1): 51–71. <https://doi.org/10.1214/10-STS321>
- Imbens GW (2014). Instrumental variables: An econometrician’s perspective. *Statistical Science*, 29(3): 323–358.
- Imbens GW, Rubin DB (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, New York.
- Imbens GW, Wooldridge JM (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1): 5–86. <https://doi.org/10.1257/jel.47.1.5>
- Janzing D, Schölkopf B (2010). Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*, 56(10): 5168–5194. <https://doi.org/10.1109/TIT.2010.2060095>
- Jiao L, Wang Y, Liu X, Li L, Liu F, Ma W, et al. (2024). Causal inference meets deep learning: A comprehensive survey. *Research*, 7: 0467. <https://doi.org/10.34133/research.0467>
- Jun SJ, Lee S (2023). Identifying the effect of persuasion. *Journal of Political Economy*, 131(8): 2032–2058. <https://doi.org/10.1086/724114>
- Kuang K, Li L, Geng Z, Xu L, Zhang K, Liao B, et al. (2020). Causal inference. *Engineering*, 6: 253–263. <https://doi.org/10.1016/j.eng.2019.08.016>
- Kusner MJ, Loftus J, Russell C, Silva R (2017). Counterfactual fairness. In: *Advances in Neural Information Processing Systems 30* (I Guyon, U von Luxburg, S Bengio, H Wallach, R Fergus, SVN Vishwanathan, R Garnett, eds.), 4066–4076. Curran Associates, Inc. Fergus, Rob and Vishwanathan, S. V. N. and Garnett, Roman.
- Laplace P (1825). *Essai Philosophique sur les Probabilités*. Bachelier, Paris, 5 edition. Originally published in 1814.
- Lauritzen SL (1996). *Graphical Models*, volume 17. Clarendon Press.
- Lauritzen SL (2004). Discussion on causality. Discussion of A. P. Dawid’s “Probability, causality, and the empirical world: A Bayes–de Finetti–Popper–Borel synthesis”. *Scandinavian Journal of Statistics*, 31(2): 189–192. <https://doi.org/10.1111/j.1467-9469.2004.03-200A.x>
- Lauritzen SL, Dawid AP, Larsen BN, Leimer H-G (1990). Independence properties of directed Markov fields. *Networks*, 20(5): 491–505. <https://doi.org/10.1002/net.3230200503>
- Lehmann EL, Casella G (1998). *Theory of Point Estimation*. Springer, New York, 2 edition.
- Li F, Ding P, Mealli F (2023). Bayesian causal inference: A critical review. *Philosophical Transactions of the Royal Society A*, 381(2247): 20220153.
- Malinsky D, Shpitser I, Richardson T (2019). A potential outcomes calculus for identifying conditional path-specific effects. In: *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics* (K Chaudhuri, M Sugiyama, eds.), volume 89 of *Proceedings of Machine Learning Research*, pages 3080–3088. PMLR.
- Manski CF (1993). Identification problems in the social sciences. *Sociological Methodology*, 23: 1–56. <https://doi.org/10.2307/271005>
- Morgan SL, Winship C (2014). *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press, Cambridge, 2 edition.
- Neison F (1844). On a method recently proposed for conducting inquiries into the comparative sanatory condition of various districts, with illustrations, derived from numerous places in Great Britain at the period of the last census. *Journal of the Statistical Society of London*, 7(1): 40–68. <https://doi.org/10.2307/2337745>

- Neyman J (1923). On the application of probability theory to agricultural experiments. Essay on principles. Section 9. *Statistical Science*, 5(4): 465–472. Translated and edited by DM Dabrowska and TP Speed from the 1923 Polish original.
- Pearl J (1985). Bayesian networks: A model of self-activated memory for evidential reasoning. In: *Proceedings of the Seventh Conference of the Cognitive Science Society* (DC Berger, ed.), 329–334. Lawrence Erlbaum Associates.
- Pearl J (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.
- Pearl J (1995). Causal diagrams for empirical research. *Biometrika*, 82(4): 669–688. <https://doi.org/10.1093/biomet/82.4.702>
- Pearl J (2001). Direct and indirect effects. In: *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI-2001)* (JS Breese, D Koller, eds.), pages 411–420. Morgan Kaufmann, San Francisco, CA.
- Pearl J (2009). *Causality*. Cambridge University Press, Cambridge, 2nd ed., first printing edition.
- Pearl J (2010a). An introduction to causal inference. *The International Journal of Biostatistics*, 6(2): 7. <https://doi.org/10.2202/1557-4679.1203>
- Pearl J (2010b). On the consistency rule in causal inference: Axiom, definition, assumption, or theorem? *Epidemiology*, 21(6): 872–875. <https://doi.org/10.1097/EDE.0b013e3181f5d3fd>
- Pearl J (2019). The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3): 54–60. <https://doi.org/10.1145/3241036>
- Pearl J, Glymour M, Jewell NP (2016). *Causal Inference in Statistics: A Primer*. John Wiley & Sons.
- Pearl J, Mackenzie D (2018). *The Book of Why: The New Science of Cause and Effect*. Basic Books, New York.
- Pearl J, Paz A (1986). Graphoids: A graph-based logic for reasoning about relevance relations. In: *Advances in Artificial Intelligence II* (BD Boulay, DHD Warren, HE Kyburg, eds.), pages 357–363. North-Holland.
- Perković E, Textor J, Kalisch M, Maathuis MH (2018). Complete graphical characterization and construction of adjustment sets in Markov equivalence classes of ancestral graphs. *Journal of Machine Learning Research*, 18(220): 1–62.
- Peters J, Bühlmann P (2014). Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101(1): 219–228. <https://doi.org/10.1093/biomet/ast043>
- Peters J, Bühlmann P, Meinshausen N (2016). Causal inference by using invariant prediction: Identification and confidence intervals. *Journal of the Royal Statistical Society, Series B, Statistical Methodology*, 78(5): 947–1012. <https://doi.org/10.1111/rssb.12167>
- Peters J, Janzing D, Schölkopf B (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press.
- Peters J, Mooij JM, Janzing D, Schölkopf B (2014). Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15(1): 2009–2053.
- Richardson TS (2003). Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1): 145–157. <https://doi.org/10.1111/1467-9469.00323>
- Richardson TS, Evans RJ, Robins JM, Shpitser I (2023). Nested Markov properties for acyclic directed mixed graphs. *The Annals of Statistics*, 51(1): 334–361. <https://doi.org/10.1214/22-AOS2253>
- Richardson TS, Robins JM (2013). Single world intervention graphs (SWIGs): A unification of

- the counterfactual and graphical approaches to causality, *Technical Report Working Paper 128*, Center for Statistics and the Social Sciences, University of Washington.
- Richardson TS, Robins JM (2014). ACE bounds; SEMs with equilibrium conditions. *Statistical Science*, 29(3): 363–366. <https://doi.org/10.1214/14-STS485>
- Richardson TS, Robins JM (2023). Potential outcome and decision theoretic foundations for statistical causality. *Journal of Causal Inference*, 11(1): 20220012. <https://doi.org/10.1515/jci-2022-0012>
- Richardson TS, Spirtes P (2002). Ancestral graph Markov models. *The Annals of Statistics*, 30(4): 962–1030. <https://doi.org/10.1214/aos/1031689015>
- Robins J, Greenland S (1989). The probability of causation under a stochastic model for individual risk. *Biometrics*, 45(4): 1125–1138. <https://doi.org/10.2307/2531765>
- Robins J, Greenland S (1991). Estimability and estimation of expected years of life lost due to a hazardous exposure. *Statistics in Medicine*, 10(1): 79–93. <https://doi.org/10.1002/sim.4780100113>
- Robins JM (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9–12): 1393–1512. [https://doi.org/10.1016/0270-0255\(86\)90088-6](https://doi.org/10.1016/0270-0255(86)90088-6)
- Robins JM (2003). Semantics of causal DAG models and the identification of direct and indirect effects. In: *Highly Structured Stochastic Systems* (PJ Green, NL Hjort, S Richardson, eds.), pages 70–81. Oxford University Press, Oxford.
- Robins JM, Greenland S (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(2): 143–155. <https://doi.org/10.1097/00001648-199203000-00013>
- Robins JM, Hernán MA, Siebert U (2004). Effects of multiple interventions. In: *Comparative Quantification of Health Risks: Global and Regional Burden of Disease Attributable to Selected Major Risk Factors* (M Ezzati, AD Lopez, A Rodgers, CJL Murray, eds.), pages 2191–2230. World Health Organization, Geneva.
- Robins JM, Richardson TS (2011). Alternative graphical causal models and the identification of direct effects. In: *Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures* (P Shrout, K Keyes, K Ornstein, eds.), chapter 6, pages 1–52. Oxford University Press.
- Rosenbaum PR (2010). *Design of Observational Studies*. *Springer Series in Statistics*. Springer, New York.
- Rothman KJ (1976). Causes. *American Journal of Epidemiology*, 104(6): 587–592. <https://doi.org/10.1093/oxfordjournals.aje.a112335>
- Rothman KJ, Greenland S, Lash TL, et al. (2008). *Modern Epidemiology*, volume 3. Wolters Kluwer Health/Lippincott Williams & Wilkins, Philadelphia.
- Roy AD (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2): 135–146. <https://doi.org/10.1093/oxfordjournals.oep.a041827>
- Rubin DB (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5): 688–701. <https://doi.org/10.1037/h0037350>
- Rubin DB (1980). Randomization analysis of experimental data: The Fisher randomization test comment. *Journal of the American Statistical Association*, 75(371): 591–593. <https://doi.org/10.2307/2287653>
- Rubin DB (2004). Direct and indirect causal effects via potential outcomes. *Scandinavian Journal of Statistics*, 31(2): 161–170. <https://doi.org/10.1111/j.1467-9469.2004.02-123.x>
- Rubin DB (2005). Causal inference using potential outcomes: Design, modeling, deci-

- sions. *Journal of the American Statistical Association*, 100(469): 322–331. <https://doi.org/10.1198/016214504000001880>
- Schölkopf B (2022). Causality for machine learning. In: *Probabilistic and Causal Inference: The Works of Judea Pearl* (H Geffner, R Dechter, JY Halpern, eds.), pages 765–804. ACM Books. Association for Computing Machinery, New York, NY.
- Shimizu S, Hoyer PO, Hyvärinen A, Kerminen A (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(72): 2003–2030.
- Shpitser I, Pearl J (2006). Identification of joint interventional distributions in recursive semi-Markovian causal models. In: *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)* (A Cohn, ed.), 1219–1226. AAAI Press.
- Shpitser I, Pearl J (2007). What counterfactuals can be tested. In: *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI-2007)* (R Parr, L van der Gaag, eds.), 352–359. AUAI Press.
- Shpitser I, Pearl J (2008). Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research*, 9: 1941–1979.
- Shpitser I, Richardson TS, Robins JM (2022). Multivariate counterfactual systems and causal graphical models. In: *Probabilistic and Causal Inference: The Works of Judea Pearl* (H Geffner, R Dechter, JY Halpern, eds.), 813–852. ACM Books, New York.
- Simon HA (1953). *Causal Ordering and Identifiability Studies in Econometric Method*. Wiley, New York.
- Spirtes P (2010). Introduction to causal inference. *Journal of Machine Learning Research*, 11: 1643–1662.
- Spirtes P, Glymour CN, Scheines R (2000). *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition.
- Strotz RH, Wold HOA (1960). Recursive and nonrecursive systems of equations. *Econometrica*, 28(2): 417–427. <https://doi.org/10.2307/1907731>
- Tchetgen Tchetgen EJ, Wang L, Sun B (2018). Discrete choice models for nonmonotone nonignorable missing data: Identification and inference. *Statistica Sinica*, 28(4): 2069. <https://doi.org/10.5705/ss.202016.0325>
- Tian J, Pearl J (2002). A general identification condition for causal effects. In: *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)* (A Darwiche, N Friedman, eds.), 567–573. AAAI Press.
- Tian J, Shpitser I (2010). On identifying causal effects. In: *Heuristics, Probability and Causality: A Tribute to Judea Pearl* (R Dechter, H Geffner, J Halpern, eds.), pages 415–444. College Publications, UK.
- Tjoa E, Guan C (2021). A survey on explainable artificial intelligence (XAI): Toward medical XAI. *IEEE Transactions on Neural Networks and Learning Systems*, 32(11): 4793–4813. <https://doi.org/10.1109/TNNLS.2020.3027314>
- Uhler C, Raskutti G, Bühlmann P, Yu B (2013). Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics*, 41(2): 436–463. <https://doi.org/10.1214/12-AOS1080>
- VanderWeele T (2015). *Explanation in Causal Inference: Methods for Mediation and Interaction*. Oxford University Press.
- VanderWeele TJ (2009). Concerning the consistency assumption in causal inference. *Epidemiology*, 20(6): 880–883. <https://doi.org/10.1097/EDE.0b013e3181bd5638>
- VanderWeele TJ, Richardson TS (2012). General theory for interactions in sufficient

- cause models with dichotomous exposures. *The Annals of Statistics*, 40(4): 2128. <https://doi.org/10.1214/12-AOS1019>
- VanderWeele TJ, Robins JM (2009). Minimal sufficient causation and directed acyclic graphs. *The Annals of Statistics*, 37(3): 1437–1465. <https://doi.org/10.1214/08-AOS613>
- Verma TS, Pearl J (1991). Equivalence and synthesis of causal models. In: *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence (UAI-91)* (P Bonissone, B D’Ambrosio, P Smets, eds.), 220–227. Morgan Kaufmann.
- Wachter S, Mittelstadt B, Russell C (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2): 841–887.
- Wang L (2022). On the homogeneity of measures for binary associations. arXiv preprint: <https://arxiv.org/abs/2210.05179>.
- Wang L, Richardson TS, Zhou X-H (2017a). Causal analysis of ordinal treatments and binary outcomes under truncation by death. *Journal of the Royal Statistical Society, Series B, Statistical Methodology*, 79(3): 719–735. <https://doi.org/10.1111/rssb.12188>
- Wang L, Tchetgen Tchetgen EJ (2018). Bounded, efficient and multiply robust estimation of average treatment effects using instrumental variables. *Journal of the Royal Statistical Society, Series B, Statistical Methodology*, 80(3): 531–550. <https://doi.org/10.1111/rssb.12262>
- Wang L, Zhou X-H, Richardson TS (2017b). Identification and estimation of causal effects with outcomes truncated by death. *Biometrika*, 104(3): 597–612. <https://doi.org/10.1093/biomet/asx034>
- Wright S (1921). Correlation and causation. *Journal of Agricultural Research*, 20: 557–585.
- Wu P, Wang L (2026). Position: A potential outcomes perspective on Pearl’s causal hierarchy. arXiv preprint: <https://arxiv.org/abs/2601.20405>
- Yang S, Wang L, Ding P (2019). Causal inference with confounders missing not at random. *Biometrika*, 106(4): 875–888. <https://doi.org/10.1093/biomet/asz048>
- Yao L, Chu Z, Li S, Li Y, Gao J, Zhang A (2021). A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data*, 15(5): 1–46. <https://doi.org/10.1145/3444944>
- Zhang K, Hyvärinen A (2009). On the identifiability of the post-nonlinear causal model. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI-2009)* (J Bilmes, A Ng, eds.), 647–655. AUAI Press.
- Zhou Y, Tang D, Kong D, Wang L (2024). Promises of parallel outcomes. *Biometrika*, 111(2): 537–550. <https://doi.org/10.1093/biomet/asae008>