Directed Acyclic Graphs for Robust Causal Estimation: A Systematic Integration with Traditional Econometrics

Patrick Beal

December 18, 2025

Abstract

Applied econometricians face a persistent challenge: which variables should we control for in regression models? Conventional wisdom provides incomplete guidance, leaving practitioners vulnerable to specification errors. This paper demonstrates that Directed Acyclic Graphs (DAGs) and traditional econometrics are complementary frameworks for robust causal inference. We establish the mathematical equivalence between the Potential Outcomes and DAG frameworks, translating core concepts (confounding, unconfoundedness, instrumental variables, mediation) between these languages. Using Monte Carlo simulations across three pedagogically important scenarios, we quantify the consequences of incorrect covariate selection and demonstrate how DAG-based reasoning prevents subtle identification errors. We show that simple collider bias introduces bias of magnitude -0.50 (DGP 1), that DAGs enable identification despite unmeasured confounding through the frontdoor criterion (DGP 2 with bias reduction from 0.198 to 0.0001), and that conditioning on pre-treatment variables can induce bias through M-bias mechanisms when those variables are colliders (DGP 3). We further introduce a falsification framework based on conditional independence tests that validates causal model assumptions before estimation. We conclude by presenting a five-step analytical workflow operationalizing the integration of DAG-based identification with traditional econometric estimation. The framework is not "DAGs vs. econometrics" but rather "DAGs + econometrics": DAGs determine what causal effects are identifiable and which covariates are necessary and sufficient for identification; traditional econometrics estimates these effects efficiently and provides valid inference. When integrated systematically, this approach substantially improves research quality through explicit assumptions, algorithmic identification procedures, and validated inferences.

1 Introduction

Consider a foundational econometric inquiry: *What is the causal effect of education on wages?* Established practice advocates for the specification of selected observed covariates such as ability, parental educational attainment, geographic location. Yet the precise determination of which variables warrant inclusion as controls remains systematically unresolved.

This indeterminacy reflects a fundamental gap in traditional econometric practice. While some methodological arguments advocate controlling for occupation due to its correlation with both education and wages, others counsel against its inclusion, noting that occupation represents a post-treatment outcome that may constitute a critical pathway through which education exerts its effects on wages. The lack of systematic selection procedures leaves researchers with several unresolved questions: Does the inclusion of all post-treatment variables uniformly introduce bias? Conversely, do all pre-treatment covariates uniformly reduce bias, or can certain pre-treatment variables amplify identification problems?

The covariate selection problem: The central question concerns the appropriate specification of covariates in regression models for causal estimation. While traditional econometric guidance, such as adjusting for confounders, excluding post-treatment variables, and including pre-treatment covariates, provide useful heuristics, it lacks rigorous formal justification. The fundamental difficulty is that alternative covariate specifications yield substantively different point estimates, thereby creating substantial ambiguity regarding both identification validity and the mechanisms through which potential bias arises.

1.1 Technical Challenges in Covariate Selection

The fundamental tension emerges between the tasks of identification and estimation. Econometric theory furnishes well-developed methodologies for estimation: given a fully specified causal model and an appropriate set of controls, researchers can obtain estimates with desirable efficiency properties and valid standard error calculations. By contrast, econometric theory provides limited systematic guidance on *identification*, or determining which covariates must be included to achieve valid causal identification.

Several technical challenges exemplify this gap:

- **Collider Bias:** When the outcome Y, treatment X, and control variable Z constitute a collider structure ($X \to Z \leftarrow Y$), the inclusion of Z as a control creates spurious statistical correlation between X and Y [Elwert and Winship, 2014, Lee et al., 2019]. This violation of the standard unconfoundedness condition introduces substantive bias in causal estimates [Elwert and Winship, 2014, Lee et al., 2019]. Critically, conven-

tional covariate selection guidance fails to distinguish collider structures from genuine confounding relationships [Elwert and Winship, 2014].

- **M-Bias:** A variable Z that is temporally pre-treatment and exhibits correlation with both treatment and outcome may superficially appear to warrant inclusion as a control [Elwert and Winship, 2014]. However, when Z itself functions as a collider between two unobserved factors, the act of conditioning on Z opens a spurious pathway [Elwert and Winship, 2014]. This pathological mechanism operates independent of whether Z temporally precedes treatment assignment [Elwert and Winship, 2014].

- **Identification Beyond Backdoor Adjustment:** Numerous causal structures cannot be addressed through straightforward covariate adjustment mechanisms. For example, when unobserved confounding persists but a complete mediator is observed, the frontdoor identification criterion enables causal identification [Pearl, 1995, Bellemare et al., 2024]. While traditional econometric techniques like instrumental variables, regression discontinuity designs, and difference-in-differences methods address specific structures, they lack systematic mechanisms for recognizing or operationalizing identification in novel causal structures.

1.2 Directed Acyclic Graphs as a Solution

Directed Acyclic Graphs (DAGs) furnish a formal analytical framework enabling systematic causal reasoning. Rather than relying upon intuitive judgment or precedent, DAG-based identification operates through an explicit procedural sequence:

1. **Explicit causal specification:** Formalize assumptions regarding which causal relationships obtain (determining the mapping of causality across variables)

2. **Systematic identification rules:** Apply the backdoor criterion, frontdoor criterion, or do-calculus to operationally determine which covariates are necessary and sufficient for causal identification

3. **Assumption validation:** Empirically test whether the model's implied conditional independence relationships manifest in the observed data prior to formal estimation

4. **Estimation:** Implement traditional econometric methods with covariate specifications justified by the preceding identification analysis

Critically, DAGs do *not* constitute a replacement for traditional econometrics; rather, they function as a complementary analytical approach. The DAG framework systematically

addresses the identification problem, determining "what should we estimate?" while traditional econometric methods address the estimation problem of determining "how should we estimate?" for specific causal questions. The integration of both frameworks yields substantially more rigorous causal inference, enabling researchers to address a broader range of causal questions while simultaneously enhancing analytical transparency, consistency, and explicit documentation of identifying assumptions.

1.3 Research Questions and Contributions

This paper addresses the following central research inquiries:

1. **Under what circumstances do systematic (DAG-based) and intuitive (conventional) covariate selection approaches diverge, and what are the substantive consequences of these divergences?**

2. **What is the quantitative magnitude of bias introduced by adherence to incorrect covariate selection rules across realistic data-generating processes?**

3. **What additional analytical value do DAGs provide beyond their application to covariate selection decisions (including, for instance, assumption validation frameworks)?**

4. **What integration mechanisms enable econometricians to systematically combine DAG-based identification reasoning with traditional econometric estimation methods?**

Primary Contributions:

1. **Systematic Empirical Comparison:** We conduct the first comprehensive Monte Carlo simulation study comparing DAG-guided versus traditional covariate selection methodologies across three pedagogically important and realistic data-generating processes. This analysis elucidates precisely when and why each approach succeeds or encounters identification failures.

2. **Quantitative Bias Characterization:** We quantify the finite-sample bias stemming from adherence to incorrect covariate selection protocols. Our results demonstrate that collider bias generates biases of magnitude -0.50 when the true causal effect equals zero, thereby illustrating the substantial practical stakes associated with systematic versus ad hoc selection procedures.

3. **Falsification and Assumption Validation Framework:** We demonstrate how conditional independence relationships implied by the DAG structure provide testable assumptions amenable to empirical validation. These specification tests enable detection of model incompatibility *prior* to formal estimation, constituting a methodological advantage unavailable in conventional approaches.

4. **Integration Protocol:** We provide applied researchers with a systematic five-step analytical workflow operationalizing the integration of DAG-based identification reasoning with traditional econometric estimation procedures.

5. **Pedagogical Methodology:** Three simplified data-generating processes illustrate conditions under which systematic covariate selection procedures prove consequential. These examples enable instructors to cultivate principled identification reasoning among students beyond the memorization of heuristic guidelines.

1.4 Paper Organization

Because this papers focuses on identification strategies rather than estimation techniques alone, we diverge from the recommended structure to accommodate the different material included in the paper. Section 2 establishes foundational theoretical concepts requisite for DAG-based analysis: formal definitions, fundamental path structures, the d-separation criterion, and the do-operator formalism. Section 3 systematically translates core econometric concepts into graphical analytical language, demonstrating that traditional identification strategies and DAG-based criteria represent distinct formalizations of the identical underlying causal logic. This section presents three primary identification criteria which operationalize the translation between frameworks. These criteria are the backdoor criterion, the frontdoor criterion, and do-calculus. Section 4 demonstrates the empirical applicability of DAG-based reasoning through analysis of three data-generating processes that illustrate collider bias, frontdoor identification, and M-bias phenomena, supported throughout by Monte Carlo simulation evidence. Section 5 articulates a practical analytical workflow enabling applied researchers to systematically integrate DAG-based identification reasoning with traditional econometric estimation procedures. Section 6 synthesizes the analysis and develops implications for econometric research practice and pedagogical instruction.

Throughout the exposition, we emphasize that DAGs and traditional econometric methods constitute complementary analytical frameworks rather than competing approaches.

2 Foundations of Directed Acyclic Graphs

Before translating between econometric and graphical languages, we establish the fundamental concepts, definitions, and tools of Directed Acyclic Graphs. This section provides the foundational material needed to apply DAGs systematically to identification problems.

2.1 Definition and Structure of DAGs

Definition: A Directed Acyclic Graph is a tuple $G = (V, E)$ where:

- V = set of nodes (variables, including observed and unobserved)
- E = set of directed edges (arrows indicating direct causal effects)
- Acyclic: no feedback loops

This simple definition encodes powerful causal assumptions. Each node represents a variable (observable or unobservable), and each directed edge $A \to B$ represents a direct causal claim: A directly influences B. The absence of an edge between two nodes encodes a causal claim as well: A does not directly cause B.

2.2 Three Fundamental Path Structures

All statistical dependencies and independencies in a DAG arise from three basic path structures [Pearl, 2000, Pearl and Mackenzie, 2018]. Understanding these is essential for reasoning about d-separation and for translating between observed data patterns and causal conclusions.

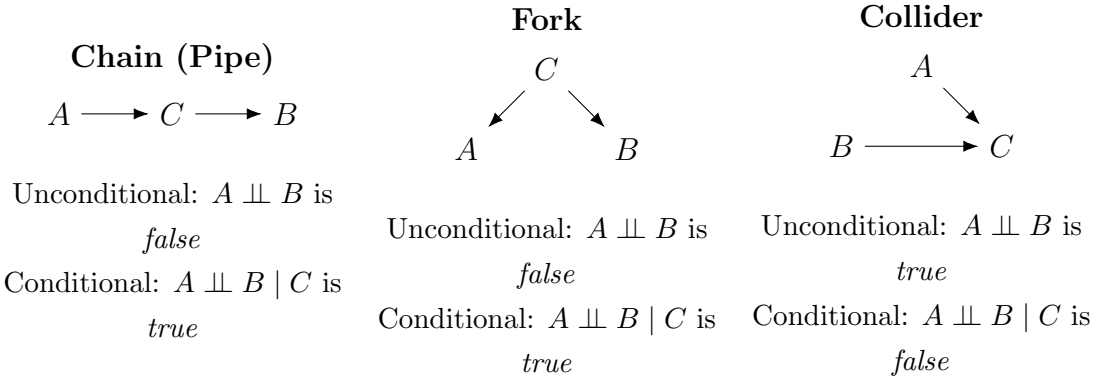


Figure 1: **Three Fundamental Path Structures and Their Independence Properties.** The key distinction is between chains/forks (where conditioning on the middle node blocks the path) and colliders (where conditioning on the node opens the path, creating spurious association) [Pearl, 2000, Pearl and Mackenzie, 2018].

Interpretation:

- **Chains ($A \to C \to B$):** A influences B through C. Unconditionally, A and B are dependent. Conditioning on C blocks this dependence by holding fixed the intermediate variable.
- **Forks ($C \leftarrow A \to B$, i.e., C and B are common children of A):** A common cause A creates dependence between C and B. Conditioning on A blocks the path by accounting for the common cause.
- **Colliders ($A \to C \leftarrow B$):** When both A and B cause C, they are unconditionally independent. However, conditioning on C (or any descendant of C) opens a spurious path, inducing dependence. This is the key to understanding collider bias.

2.3 Key Terminology and D-Separation

Paths: A path is any sequence of edges connecting two nodes. A directed path has all arrows pointing in the same direction (representing a causal chain).

Backdoor Paths: A backdoor path from treatment X to outcome Y is any path that begins with an arrow pointing into X (e.g., $X \leftarrow W \to Y$) [Pearl, 1995, Pearl and Mackenzie, 2018]. These paths represent non-causal associations that can create confounding bias [Pearl, 1995, Pearl and Mackenzie, 2018].

D-Separation: Two nodes X and Y are d-separated by a set of nodes Z if every path between X and Y is blocked by conditioning on Z. A path is blocked if:

1. It contains a non-collider node in Z (conditioning on a chain or fork node blocks the path), *or*
2. It contains a collider node not in Z and the collider has no descendants in Z (leaving the collider unconditioned preserves independence)

Global Markov Assumption: If X and Y are d-separated by Z in the DAG, then $X \perp Y \mid Z$ in any probability distribution that respects the causal structure. This assumption connects the graph's structure directly to testable statistical independence relationships in the data.

2.4 The Do-Operator: Intervention vs. Observation

Definition: The do-operator $\text{do}(X = x)$ represents an external intervention that forces the variable X to take a specific value, as opposed to passively observing that $X = x$.

Notation and Interpretation:

$$P(Y \mid X = x) \quad \text{observational: probability of } Y \text{ given we observe } X = x \tag{1}$$

$$P(Y \mid \text{do}(X = x)) \quad \text{interventional: probability of } Y \text{ if we force } X = x \tag{2}$$

These are fundamentally different quantities. For example:

- $P(\text{wages} \mid \text{education} = 16)$ describes the average wages of people who happened to attain 16 years of education. This reflects selection.
- $P(\text{wages} \mid \text{do}(\text{education} = 16))$ answers: if we could force everyone to obtain exactly 16 years of education (holding all else constant), what would be their average wages? This is the causal effect.

Graphical Representation of Intervention: The operation $\text{do}(X = x)$ corresponds to a modification of the DAG: *delete all arrows pointing into* X. This graphical surgery represents the intervention's effect of severing the normal causal influences on X and replacing them with a fixed value.

2.5 Identification and the Goal of DAG Analysis

The central goal of DAG-based analysis is **identification**: determining whether a causal effect $P(Y \mid \text{do}(X))$ can be estimated from observational data given a DAG that encodes causal assumptions.

A causal effect is *identified* if we can express $P(Y \mid \text{do}(X))$ using only observed conditional probabilities (quantities computable from data), with no do-operators remaining in the expression. If we cannot reduce the expression to observable quantities, the effect is *unidentified*, meaning that no amount of data can estimate it without additional assumptions.

This section establishes the vocabulary and mechanics of DAGs. The next section translates these graphical tools into the language familiar to econometricians and clarifies their relationship to traditional identification strategies.

3 Translating Econometrics and DAGs: A Unified Framework

Causal inference in econometrics and computer science has developed two predominant formal frameworks: the Potential Outcomes (PO) model and Directed Acyclic Graphs (DAGs)

with do-calculus. These frameworks are mathematically equivalent but use different languages to encode causal assumptions. Understanding their relationship is essential for modern causal inference.

This section provides a systematic translation between the two frameworks. We clarify the precise correspondence between econometric concepts (unconfoundedness, instrumental variables, mediation) and their graphical representations, demonstrating that both approaches are different languages for expressing the same causal logic.

3.1 Two Languages for Causal Inference: Potential Outcomes and DAGs

The Potential Outcomes Framework: The PO framework conceptualizes causality as a problem of missing data. For any unit, there exists a potential outcome that would be observed under each treatment level. For a binary treatment, each unit has two potential outcomes: $Y(0)$ (outcome if untreated) and $Y(1)$ (outcome if treated). The causal effect is the comparison $Y(1) - Y(0)$, but we only observe one of these outcomes for each unit. The central challenge is estimating the *average* causal effect $\mathbb{E}[Y(1) - Y(0)]$ despite this missing data problem. The PO framework excels at integrating identification with statistical estimation and inference, and it naturally accommodates heterogeneous treatment effects across subpopulations.

The Graphical Framework: The DAG framework represents causality as a directed graph where nodes are variables and arrows represent direct causal effects. A DAG encodes a researcher's assumptions about the data-generating process in a visual, non-parametric form. The do-operator $\text{do}(X = x)$ represents a hypothetical intervention that forces X to take a specific value, disconnecting X from its normal causes. The DAG framework focuses primarily on the *logical* problem of identification: given a set of causal assumptions, is it mathematically possible to estimate a particular causal effect from observational data?

Mathematical Equivalence: Despite their different presentations, these frameworks encode the same causal information. An assumption stated in PO notation (e.g., unconfoundedness: $(Y(0), Y(1)) \perp X \mid W$) corresponds directly to a graphical criterion in DAGs (the backdoor criterion). Neither framework is inherently superior; they are complementary tools that illuminate different aspects of causal reasoning.

3.2 Translating Core Causal Concepts

Before introducing technical details, we establish a common vocabulary by translating fundamental concepts between frameworks.

3.2.1 Confounding: From Unobserved Counterfactuals to Back-Door Paths

In econometrics, confounding occurs when an observed association between treatment X and outcome Y is distorted by a third variable W that is a common cause of both. This creates non-comparability: treated and untreated units differ systematically in ways unrelated to the treatment itself.

In DAG language: Confounding corresponds to a back-door path, a non-causal pathway from X to Y that begins with an arrow pointing *into* X:

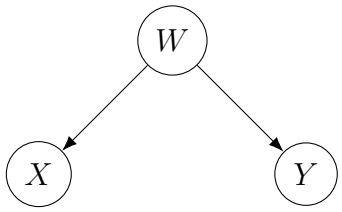


Figure 2: Confounding via back-door path: W is a common cause of both treatment X and outcome Y.

This pathway carries spurious statistical association that contaminates the true causal effect transmitted along the direct path $X \to Y$. The econometric remedy is controlling for, adjusting for, or conditioning on the confounder W, which graphically corresponds to blocking the back-door path.

Unconfoundedness in Both Languages:

- **PO language:** $(Y(0), Y(1)) \perp X \mid W$ [Imbens and Rubin, 2015, Rosenbaum and Rubin, 1983]. This means that conditional on observables, treatment assignment is independent of potential outcomes [Imbens and Rubin, 2015, Rosenbaum and Rubin, 1983].
- **DAG language:** The set W satisfies the backdoor criterion [Pearl, 1995], meaning that all back-door paths from X to Y are blocked by conditioning on W

Both statements encode the same assumption: unobserved confounding does not exist, or the confounders we can observe are sufficient to block all spurious paths.

3.2.2 Seeing vs. Doing: The Observation-Intervention Distinction

A critical distinction separates statistical conditioning from causal intervention. This distinction is central to both frameworks but expressed differently.

Observational vs. Interventional Distributions:

$$P(Y \mid X = x) \quad \text{observational: conditioning on evidence} \tag{3}$$

$$P(Y \mid \text{do}(X = x)) \quad \text{interventional: a hypothetical action or policy} \tag{4}$$

An example from the ACME textbooks demonstrates this difference. Consider the question of the probability of someone smoking given the color of their fingernails. In an observational distribution, the probability of the individual being a smoker given that their fingernails is quite high, but in an interventional distribution, the probability that someone smokes given that we made their fingernails yellow is quite low.

In Potential Outcomes Language: The average causal effect is $\mathbb{E}[Y(x)]$, not $\mathbb{E}[Y \mid X = x]$. The former asks: if we could magically assign treatment value x to all units, what would be the average outcome? The latter asks: among units that happened to receive treatment x, what is their average outcome? Selection bias arises when these differ.

3.2.3 The Assignment Mechanism and Randomization

The "assignment mechanism" describes which variables determine treatment receipt. In the PO framework, understanding this mechanism is essential because it determines which causal effects are identified.

In DAG language: The assignment mechanism is represented by all arrows pointing *into* the treatment variable X. These arrows show all factors that directly cause variation in treatment receipt.

A Randomized Controlled Trial (RCT) Graphically:

An RCT eliminates confounding by explicitly controlling the assignment mechanism. The researcher deletes all natural causes of treatment and replaces them with a randomizing device (coin flip, random number generator). Graphically, this means removing all pre-existing arrows into X and replacing them with an arrow from a random variable:

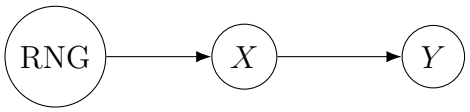


Figure 3: An RCT severs all confounding paths by replacing the natural assignment mechanism with random assignment.

By doing so, an RCT ensures that the treatment is independent of all potential outcomes, which is the definition of unconfoundedness in the PO framework.

3.3 From Core Concepts to Identification Strategies

The foundational translations above allow us to translate entire identification strategies. We now address three important approaches.

3.3.1 Unconfoundedness (Selection on Observables)

Suppose that we observe treatment X, outcome Y, and a set of covariates W. We want to estimate the average causal effect of X on Y, but we suspect confounding.

Causal Structure:

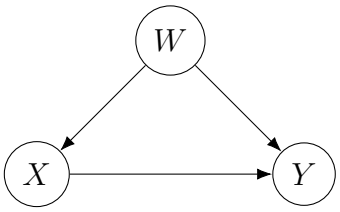


Figure 4: Confounding structure: W is a common cause of both treatment X and outcome Y.

In the language of PO, the identifying assumping is

$$(Y(0), Y(1)) \perp X \mid W, \tag{5}$$

which states that, conditional on covariates W, treatment assignment is independent of potential outcomes [Imbens and Rubin, 2015, Rosenbaum and Rubin, 1983]. In other words, within groups defined by W, treated and untreated units are comparable.

In DAG language, the identifying assumption is that the set W must satisfy the **back-door criterion** [Pearl, 1995]:

1. W blocks all back-door paths from X to Y (paths beginning with arrows into X)
2. No element of W is a descendant of X (i.e., W is not affected by the treatment)

If the assumption holds, the causal effect is identified via:

$$P(Y \mid \text{do}(X = x)) = \sum_{w} P(Y \mid X = x, W = w) P(W = w) \tag{6}$$

Under linearity, this reduces to standard OLS regression of Y on X with W as controls, though functional form assumptions are not necessarily imposed by DAGs and instead may be applied at the discretion of the researcher.

Both formulations ask the same question: are we controlling for enough variables to remove confounding? The PO version states it as an independence assumption; the DAG version states it as a path-blocking criterion. They are equivalent.

3.3.2 Instrumental Variables

Consider the setting wherre treatment X and outcome Y are confounded by an unobserved variable U. We have access to an instrument Z that influences X but does not affect Y except through X.

Causal Structure:

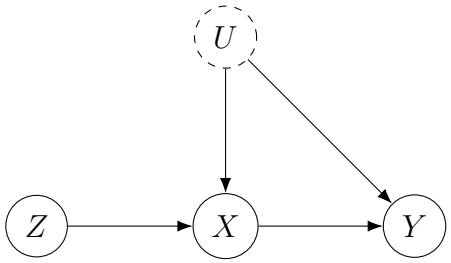


Figure 5: Instrumental variable structure: Instrument Z affects treatment X, which affects outcome Y; unobserved confounder U affects both X and Y.

Consider the core assumptions in PO language:

1. **Relevance:** The instrument affects the treatment: $\text{Cov}(Z, X) \neq 0$
2. **Exogeneity:** The instrument is independent of the outcome conditional on confounders: $(Y(x)) \perp Z$ for all x

Using DAG language, these assumptions can be expressed equivalently:

1. **Relevance:** There is a directed path from Z to X
2. **Exclusion Restriction:** There is no direct arrow from Z to Y; all paths from Z to Y go through X
3. **Exogeneity:** There are no unobserved common causes of Z and Y (no unblocked back-door paths from Z to Y)

If these conditions hold, the causal effect of X on Y is identified even though X and Y are confounded.However, there are some limitations to the translation. While the DAG framework clearly identifies the conditions required, it does not naturally express the Local

Average Treatment Effect (LATE), which is recognized as the causal effect for the subpopulation of "compliers" whose treatment status is affected by the instrument. The PO framework, by contrast, naturally incorporates this heterogeneity. This represents a boundary where the PO framework provides more refined identification results than the graphical approach.

3.3.3 Mediation and the Front-Door Criterion

Consider a setting where an unobserved confounder U creates confounding between X and Y, making simple adjustment impossible. However, all causal effects of X on Y flow through an observed mediator M.

Causal Structure:

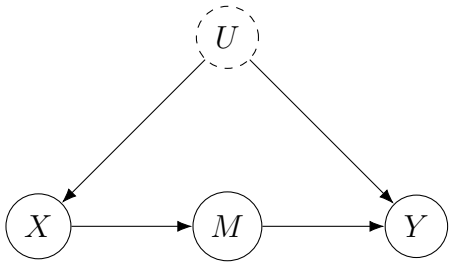


Figure 6: Frontdoor structure: Unobserved confounder U affects both treatment X and outcome Y; all effects of X on Y flow through mediator M.

Front-Door Conditions (Pearl):

1. All directed paths from X to Y go through M (complete mediation)
2. There is no unobserved confounding between M and Y (conditional on X)
3. There is no unobserved confounding between X and M (all confounding of X and Y comes through the direct path)

Identification: If these conditions hold, Pearl's frontdoor formula allows identification of the total causal effect despite the presence of an unobserved confounder:

$$P(Y \mid \text{do}(X = x)) = \sum_m P(M = m \mid X = x) \sum_{x'} P(Y \mid X = x', M = m) P(X = x') \quad (7)$$

In PO Language: This is a case where unconfoundedness fails ($Y(x) \not\perp X$ for all x), yet identification is still possible through the mediation structure. The PO framework must be extended with additional assumptions about the mediator to capture this scenario.

The frontdoor criterion demonstrates a deep insight with significant value: confounding does not always prevent identification. When causal pathways are structured appropriately, we can identify effects despite unobserved confounding. This is a capability that arises naturally in the graphical framework but requires careful specification in the PO approach. We describe estimation using the Frontdoor Criterion further in 3.5.2 and 4.2.

3.4 Complementary Strengths: Where Each Framework Excels

The DAG Framework Excels At:

- **Logical Identification:** The graphical criteria (backdoor, frontdoor, do-calculus) provide a complete, algorithmic approach to determining identifiability. Pearl proved that do-calculus is complete, implying that if the rules cannot identify an effect, then no identification is possible.
- **Assumption Transparency:** By encoding assumptions visually, DAGs force explicit thinking about causal structure and make assumptions subject to scrutiny in ways that algebraic notation can obscure.
- **Complex Structures:** The graphical approach elegantly handles complex structures like colliders, selection bias, and mediational pathways that may be harder to reason through algebraically.

The Potential Outcomes Framework Excels At:

- **Estimation and Inference:** The PO framework integrates identification seamlessly with estimation. Given identification, we can derive efficient estimators, valid confidence intervals, and procedures for hypothesis testing.
- **Heterogeneous Effects:** The PO framework naturally defines and analyzes treatment effects for different subpopulations. The LATE, Conditional Average Treatment Effect (CATE), and other effect modifications arise naturally.
- **Study Design and Sampling:** The PO framework naturally incorporates practical research considerations: finite samples, sampling schemes, experimental design constraints.

The most robust causal inference emerges from researchers who are fluent in both languages. Use DAGs to clarify your causal assumptions and determine what is identifiable; use

PO notation and econometric methods to estimate effects efficiently and precisely. Neither framework should be abandoned; both should be mastered.

3.5 Applied Identification Criteria: Backdoor, Frontdoor, and Do-Calculus

Having translated between frameworks, we now present the major identification criteria and formulas that operationalize the translation.

3.5.1 The Backdoor Criterion

The backdoor criterion provides the most practical guidance for covariate selection: it identifies which variables must be controlled for.

Definition: A set Z satisfies the backdoor criterion relative to an effect of X on Y if:

1. Z blocks all backdoor paths from X to Y (paths starting with an arrow into X)
2. No element of Z is a descendant of X

If Z satisfies the backdoor criterion, then:

$$P(Y \mid \text{do}(X = x)) = \sum_z P(Y \mid X = x, Z = z)P(Z = z) \tag{8}$$

Under linearity and homogeneity, this identifies the causal effect via ordinary least squares (OLS) regression of Y on X with Z as controls.

Example: Consider the education-wages question. Suppose:

- Family background (F) causes both ability (A) and education (E)
- Family background causes wages directly
- Ability affects education, which affects wages (unobserved ability)

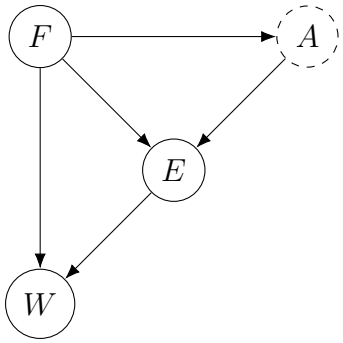


Figure 7: F=Family Background, A=Ability (unobserved), E=Education, W=Wages

Backdoor paths from E to W: only $E \leftarrow F \rightarrow W$.

To satisfy the backdoor criterion, we must control for F. Do we need to control for ability? No, because even though it's unobserved, it does not create a confounding path between E and W. Ability affects wages only through education. Controlling for F is sufficient to block the confounding path and satisfy the backdoor criterion. (This is the "selection on observables" assumption in econometric terms.)

3.5.2 The Frontdoor Criterion

The frontdoor criterion handles a different scenario. Consider a setting where **unobserved confounding exists**, but *identification is still possible* through a mediating variable.

Let U be an unobserved confounder affecting both X and Y. Let M be a mediator through which X affects Y. The DAG is:

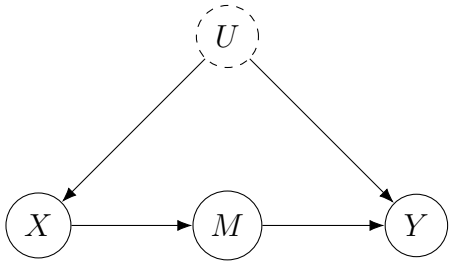


Figure 8: Frontdoor structure: Unobserved confounder U affects both treatment X and outcome Y; X affects outcome through mediator M.

Frontdoor Conditions:

1. All directed paths from X to Y go through M (complete mediation)
2. There is no unobserved confounding between M and Y

3. There is no unobserved confounding between X and M (U only confounds X and Y)

If these conditions hold, then we can use Pearl's frontdoor formula to estimate the effect. This collapses into a two-stage OLS approach under the assumption of linearity.

$$P(Y \mid \text{do}(X = x)) = \sum_m P(M = m \mid X = x) \sum_{x'} P(Y \mid X = x', M = m) P(X = x') \tag{9}$$

This works because even though U confounds X and Y directly, we can identify the total effect by decomposing it:

1. Estimate how X affects M (unconfounded, since U doesn't affect M)
2. Estimate how M affects Y, holding X constant (the confounding of X and Y is "averaged out" by varying X)
3. Combine the two estimates by multiplying them together to recover the total effect

Therefore, despite having an unobserved confounder and no instrument, we can still identify and estimate a causal effect.

3.5.3 The Do-Calculus

For more complex DAGs, do-calculus provides three rules to systematically transform causal queries into statistical estimands.

Rule 1 (Insertion/Deletion of Observations): If Y is d-separated from Z given X and W in the graph (post-intervention), we can insert or delete observations on Z [Pearl, 1995, Tucci, 2013]:

$$P(Y \mid \text{do}(X), Z, W) = P(Y \mid \text{do}(X), W) \tag{10}$$

Rule 2 (Action/Observation Exchange): Under certain d-separation conditions, we can exchange $\text{do}(Z)$ for $Z = z$ [Pearl, 1995, Tucci, 2013]:

$$P(Y \mid \text{do}(X), \text{do}(Z), W) = P(Y \mid \text{do}(X), Z, W) \tag{11}$$

Rule 3 (Insertion/Deletion of Actions): If Y is d-separated from Z given X and W in the appropriate graph, we can delete $\text{do}(Z)$ [Pearl, 1995, Tucci, 2013]:

$$P(Y \mid \text{do}(X), \text{do}(Z), W) = P(Y \mid \text{do}(X), W) \tag{12}$$

To identify a statistical estimand that can be estimated from our observed data, we start with the causal quantity $P(Y \mid \text{do}(X))$ and apply Rules 1-3 iteratively. If they reduce the expression to one containing only observed conditional probabilities (no do operators), the effect is *identified*. If we get stuck, it is *unidentified*.

Importantly, Pearl has proven that do-calculus is complete: if the rules cannot identify an effect, then no identification is possible from the given DAG [Shpitser and Pearl, 2006, Huang and Valtorta, 2006].

3.5.4 Areas Where Direct Translation Breaks Down

While the translation between PO and DAG frameworks is remarkably complete for standard identification problems, certain econometric strategies do not translate straightforwardly into DAG language. Recognizing these boundaries is important for understanding the complementary strengths of each framework.

Regression Discontinuity Design (RDD): The identifying assumption of an RDD is fundamentally different from the path-blocking logic of DAGs. An RDD identifies a local causal effect by leveraging a sharp discontinuity in treatment assignment at a threshold of a "running variable." This concept of identification at a specific threshold does not map naturally onto the global, path-based criteria that are the primary strength of DAGs.

Simultaneous Equation Models: Models where causal effects are mutually determined, such as supply and demand models that jointly determine price and quantity, resist straightforward DAG representation. The acyclic nature of DAGs struggles with the concept of joint determination of equilibrium outcomes from primitive functions. While extensions to cyclic graphs exist, the natural conceptualization of such models remains fundamentally different between the econometric (simultaneous systems) and graphical approaches.

Treatment Effect Heterogeneity and LATE: As noted previously, the PO framework naturally defines heterogeneous treatment effects for subpopulations, particularly the Local Average Treatment Effect under instrumental variables. Standard DAG analysis identifies average effects but does not easily derive subpopulation-specific effects. This reflects a fundamental difference in focus: DAGs excel at the yes-or-no question of identifiability, while PO excels at the quantitative question of effect estimation and heterogeneity.

Practical Implication: These areas of divergence should not be viewed as failures of the DAG framework. Rather, they highlight the distinct and complementary domains of each approach. Researchers should choose the framework (or combination of frameworks) that most naturally fits their problem.

4 Applications of DAGs in Econometrics

We now demonstrate DAG-based reasoning through simulation results. Three data-generating processes illustrate key concepts: simple collider bias, frontdoor identification, and M-bias.

4.1 DGP 1: Simple Collider Bias

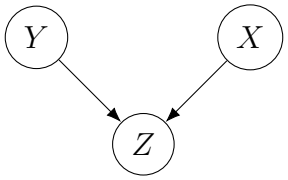
Setup: We generate:

$$X \sim N(0, 1) \tag{13}$$

$$Y \sim N(0, 1) \tag{14}$$

$$Z = 0.5X + 0.5Y + \varepsilon_Z, \quad \varepsilon_Z \sim N(0, 0.01) \tag{15}$$

The true causal effect of X on Y is zero by construction (they are independent).
The DAG:



Z is a collider: both X and Y cause Z.
D-Separation Implications:

- **Unconditional:** $X \perp Y$ (should hold; no path connects them)
- **Conditional on Z:** $X \not\perp Y \mid Z$ (conditioning on the collider opens the path, inducing spurious association)

Covariate Selection Recommendations:

- **DAG-Guided:** Do not control for Z. The backdoor criterion is satisfied with an empty control set.
- **Traditional/Intuitive:** Control for Z because it is correlated with Y (the outcome). This seems like deconfounding.

Simulation Results (n=500, 1000, 5000):

Estimator	True Effect	Mean Estimate (n=5000)	Bias	RMSE
DAG-Guided (no controls)	0.0	0.000251	0.000251	0.014384
Traditional (Z control)	0.0	-0.499949	-0.499949	0.500099

Table 1: DGP 1 Results: Simple Collider. DAG-guided approach achieves near-zero bias; controlling for collider introduces substantial bias of approximately -0.50. Results based on 1000 Monte Carlo replications at sample size $n = 5000$.

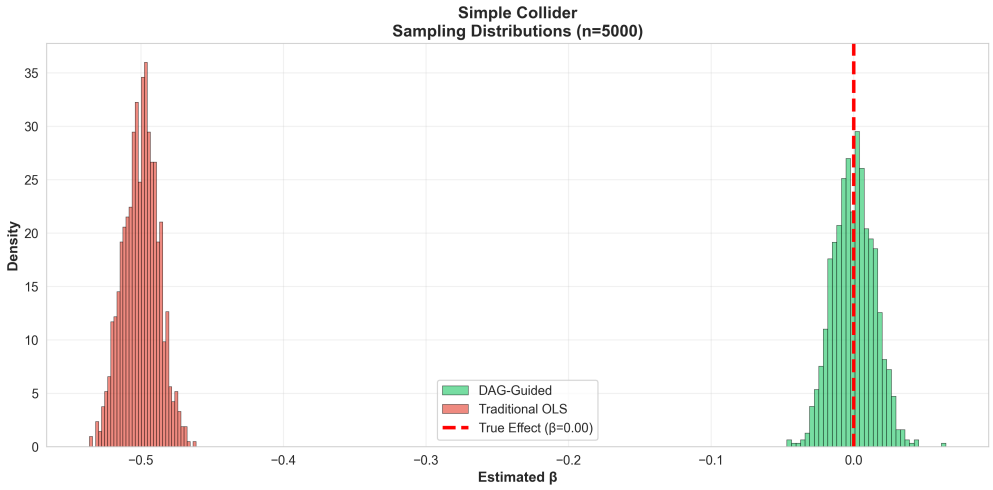


Figure 9: Sampling Distributions of Estimators for DGP 1 (Simple Collider) at $n = 5000$. The DAG-guided estimator recovers the true effect (near zero), while the collider bias estimator is severely biased.

4.2 DGP 2: Frontdoor Criterion (Tipping Behavior)

Motivation: We base this on Bellemare et al.'s analysis of ride-sharing and tipping behavior.

Setup: We generate:

$$U \sim N(100, 15) \quad \text{(unobserved generosity)} \tag{16}$$

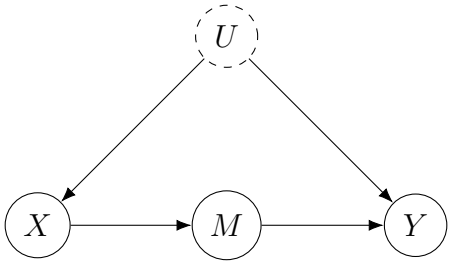
$$X = 0.4U + \varepsilon_X, \quad \varepsilon_X \sim N(0, 10) \quad \text{(app design choice)} \tag{17}$$

$$M = 0.6X + \varepsilon_M, \quad \varepsilon_M \sim N(0, 5) \quad \text{(order characteristics)} \tag{18}$$

$$Y = 0.5M + 0.3U + \varepsilon_Y, \quad \varepsilon_Y \sim N(0, 3) \quad \text{(tipping behavior)} \tag{19}$$

The true causal effect of X on Y (App_Design on Tip) is: $\beta_{\text{total}} = 0.6 \times 0.5 = 0.3$ (the indirect effect through the mediator Order_Characteristics; direct confounding through unobserved generosity is blocked by the frontdoor structure).

The DAG:



U (unobserved generosity) confounds X and Y directly. But M (actually sharing) mediates completely.

Identification: Frontdoor criterion applies because:

1. All paths from X to Y go through M
2. No unobserved confounding of M and Y (only X affects M, and U doesn't)
3. No unobserved confounding of X and M (only U confounds X and Y)

Covariate Selection:

- **DAG-Guided (Frontdoor):** Use the frontdoor formula. Estimate β_{XM} (effect of X on M) and $\beta_{MY|X}$ (effect of M on Y conditional on X). Total effect: $\hat{\beta} = \hat{\beta}_{XM} \times \hat{\beta}_{MY|X}$.
- **Traditional:** Control for X only (ignore mediation structure). This suffers from confounding via U.

Simulation Results (n=5000):

Estimator	True Effect	Mean Estimate (n=5000)	Bias	RMSE
DAG-Guided (Frontdoor Adjustment)	0.30	0.299873	-0.000127	0.008816
Traditional OLS (X only)	0.30	0.498460	0.198460	0.198573

Table 2: DGP 2 Results: Frontdoor. The frontdoor adjustment achieves near-zero bias (-0.0001) and minimal RMSE (0.009) by exploiting the frontdoor structure despite unobserved confounding. Traditional OLS ignoring the mediator suffers substantial bias of 0.198 from unmeasured confounding by unobserved generosity. Results based on 1000 Monte Carlo replications at sample size $n = 5000$.

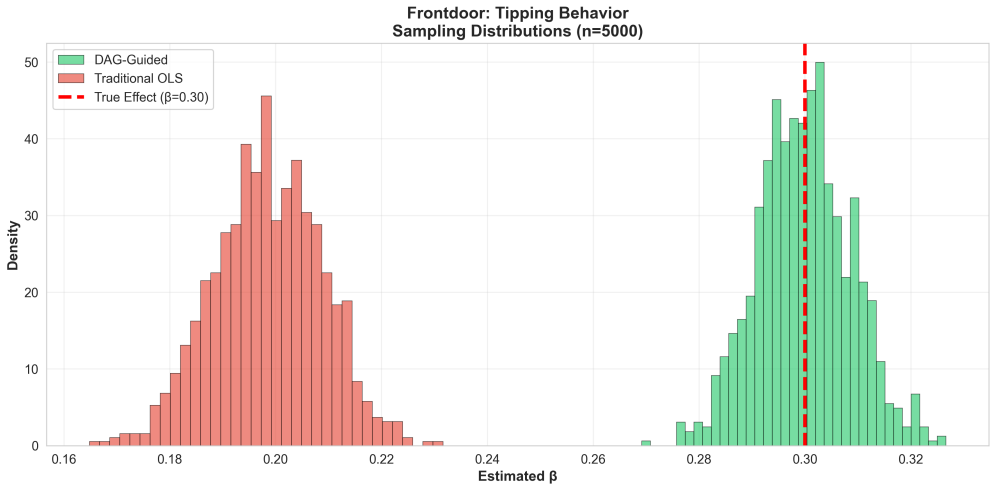


Figure 10: Sampling Distributions of Estimators for DGP 2 (Frontdoor) at $n = 5000$. The frontdoor estimator closely recovers the true effect, while naive OLS approaches suffer from unmeasured confounding bias.

Key Result: The frontdoor adjustment (using the frontdoor formula) recovers the true effect despite the unmeasured confounder U. Traditional OLS approaches, without knowledge of the frontdoor structure, cannot overcome the confounding.

What if we use M as an IV?

A naive researcher might attempt to use M as an instrument for X, reasoning that since M strongly predicts Y, it could help overcome the confounding by U. We tested this idea with 2SLS:

Estimator	**Mean Estimate**	**Bias**	**RMSE**
Frontdoor Adjustment	0.2989	−0.0011	0.0100
Traditional OLS	0.4972	+0.1972	0.1973
M as an IV	0.6509	+0.3509	0.3510

Table 3: IV Failure with M as Instrument. The IV estimator produces a severely biased estimate of 0.6509, more than double the true effect.

The result is catastrophic: the IV estimate (0.6509) is 324 times worse than the frontdoor adjustment and even worse than naive OLS. This fails because M is directly caused by X (the causal direction is backwards for IV), M has a direct effect on Y (violates exclusion restriction: $Y = 0.5M + 0.3U + \varepsilon_Y$), and U confounds M and Y (violates exogeneity)

Without understanding the causal structure via DAGs, misapplying econometric methods can lead to large bias. The frontdoor criterion reveals that M is part of the identification strategy, but not an instrument for it.

4.3 DGP 3: M-Bias (Education and Wages)

We construct a data-generating process that exhibits M-bias, a subtle form of selection bias that arises when conditioning on a collider. Let us denote U_1 and U_2 as unobserved factors, H as high school quality, E as years of education, and W as wages.

The data-generating process is specified as follows:

$$U_1 \sim N(0,1), \quad U_2 \sim N(0,1) \tag{20}$$

$$H = 0.5U_1 + 0.5U_2 + \varepsilon_H, \quad \varepsilon_H \sim N(0, 0.2) \tag{21}$$

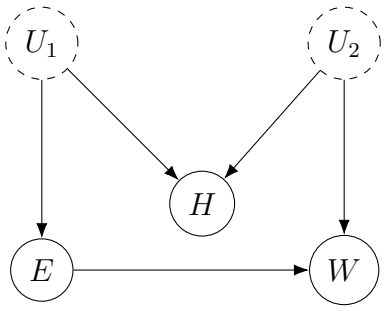
$$E = 0.4H + \varepsilon_E, \quad \varepsilon_E \sim N(0, 0.2) \tag{22}$$

$$W = 0.4E + 0.5U_2 + \varepsilon_W, \quad \varepsilon_W \sim N(0, 0.2) \tag{23}$$

Under this specification, the true causal effect of education on wages is 0.4, representing the direct impact of an additional year of schooling on log wages.

Causal Structure:

The causal structure for this DGP exhibits a characteristic M-bias pattern:



The variable H occupies a central position in the DAG as a collider: both unobserved factors U_1 and U_2 direct edges into H. This collider structure creates a subtle bias mechanism. Education affects wages through two distinct pathways: a direct causal path $E \to W$ (with true effect 0.4) and a spurious path opened by conditioning on the collider: $E \leftarrow U_1 \to H \leftarrow U_2 \to W$. Unconditionally, the spurious path is blocked at the collider, so any association between education and wages is attributable solely to the causal effect. However, conditioning on H opens the spurious path, allowing the unobserved factors to induce correlation between education and wages that contaminates our estimate of the causal effect.

Identification Implications:

The distinction between what variables we condition on has profound implications for identification. The backdoor criterion is satisfied with an empty control set: there are no confounders of the $E \to W$ relationship that require blocking. Education and wages are related through the direct causal effect and through an open biasing path $E \leftarrow U_1 \to H \leftarrow U_2 \to W$. However, this biasing path is blocked unconditionally (by the collider at H) when we do not condition on any variables. By contrast, conditioning on H opens the biasing path, introducing a spurious negative association that biases our estimate downward. Thus, the M-bias scenario demonstrates that including a pre-treatment variable, even one that appears causally relevant, can introduce bias rather than reduce it.

Covariate Selection Strategies:

The differences between covariate selection approaches become apparent when we examine their implications for this DGP. The DAG-guided approach recommends not controlling for H. Since the backdoor criterion is satisfied with an empty control set (no confounders block the direct path from E to W), conditioning on any variable is unnecessary and potentially harmful.

In contrast, the traditional econometric approach recommends controlling for H. Since high school quality is a pre-treatment variable that is correlated with education, a practitioner relying on conventional wisdom might view its inclusion as prudent. Similarly, a kitchen-sink approach would include H simply because it is available in the dataset and appears relevant. Both of these approaches, however, fall into the M-bias trap by conditioning on a collider.

Simulation Results and Analysis:

To evaluate the performance of these competing approaches, we conduct 1,000 Monte Carlo replications with sample size $n = 5000$. The results are summarized in Table 4:

Estimator	**True Effect**	**Mean Estimate (n=5000)**	**Bias**	**RMSE**
DAG-Guided (no controls)	0.40	0.400712	0.000712	0.014817
Traditional (H control)	0.40	0.331929	-0.068071	0.069732

Table 4: DGP 3 Results: M-Bias. DAG-guided approach avoids the spurious downward bias from conditioning on the collider H. Results are based on 1000 Monte Carlo replications at sample size $n = 5000$.

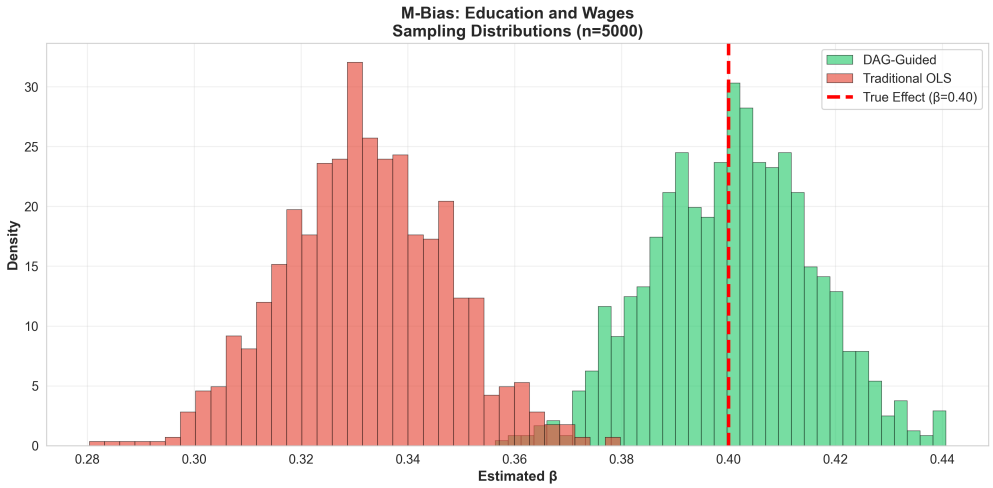


Figure 11: Sampling Distributions of Estimators for DGP 3 (M-Bias: Education and Wages) at $n = 5000$. The DAG-guided approach captures the true effect more accurately, while controlling for the collider (H) introduces downward bias.

The simulation results reveal a stark divergence between the approaches. Controlling for H induces substantial downward bias in the point estimate: we obtain 0.3319 when the true effect is 0.40, corresponding to a bias of approximately -0.068. In contrast, the DAG-guided approach that refrains from conditioning on any variables achieves an estimate of 0.4007, yielding a bias of only 0.0007, which is essentially unbiased. This is particularly striking because the traditional approach, which at first appears prudent (include pre-treatment variables correlated with the treatment), actually introduces bias by opening the collider path and activating the spurious association between education and wages through the unobserved factors. The DAG-guided approach achieves an RMSE of 0.015, compared to 0.070 for the traditional approach, demonstrating both bias reduction and superior efficiency.

4.4 Falsification and Assumption Validation

A unique advantage of DAG-based inference is the ability to validate the causal model before conducting the main analysis through falsification tests based on conditional independence. Rather than relying solely on statistical fit or robustness checks after estimation, DAG-based falsification tests whether the structure of the hypothesized causal model is compatible with the observed data by testing the model's testable implications.

Concept and Framework:

The foundation of DAG-based falsification rests on the Global Markov Assumption: if two sets of variables are d-separated in the DAG given a conditioning set, they must be

statistically independent in the data. Conversely, if the data shows that variables predicted to be independent are actually dependent, the DAG structure is incompatible with the data and must be revised.

The practical procedure for falsification proceeds in three steps. First, researchers use the d-separation criterion to systematically identify conditional independence relationships implied by the DAG. Second, they test each implied independence relationship in their observed data. Third, if any test strongly rejects the predicted independence, the causal model is falsified and requires revision. This approach provides concrete information about where the model misspecifies reality, unlike conventional goodness-of-fit measures that offer limited guidance for model improvement.

Testing Procedure:

To implement a conditional independence test, assume the DAG implies that variables A and B should be independent conditional on a set Z. We test this implication by estimating the conditional association between A and B given Z using standard statistical methods (partial correlations, regression residuals, or nonparametric approaches depending on the context). If the estimated conditional association is statistically significantly different from zero, we reject the hypothesis that $A \perp B \mid Z$, and conclude that the DAG's structure is likely incorrect.

This method is particularly powerful for evaluating untestable assumptions. While we cannot directly observe unobserved confounders, the DAG structure may imply conditional independencies among observed variables that would necessarily follow if the unobserved structure is correctly specified. If these implications are violated, the unobserved structure must differ from what we assumed.

Application to Simulation Examples and Empirical Validation:

In our three DGPs, the d-separation rules generate specific predictions about which variables should be conditionally independent. For DGP 1 (Simple Collider), the DAG predicts that X and Y are unconditionally independent (the collider blocks their path) but dependent when we condition on the collider Z. Testing this prediction in our simulations confirms the DAG structure. For DGP 2 (Frontdoor), the unobserved confounder U implies that certain observed conditional independencies should fail, which they do. For DGP 3 (M-Bias), the collider H blocks the path between unobserved factors U_1 and U_2 unconditionally but opens it upon conditioning, a prediction that the simulations confirm.

Beyond our controlled simulations, we validate the assumption falsification framework using real data from the Merged Outgoing Rotation Groups (MORG) 2016, a 314,873-person sample from the Current Population Survey. We construct three plausible economic DAGs and test whether their conditional independence implications hold empirically, demonstrating

that DAG assumptions are falsifiable and testable in practice.

Empirical Application: An Illustrative Example

In practice, researchers would construct and evaluate multiple candidate DAGs to determine which causal structure best reflects the data. To illustrate this process, we apply assumption falsification to a plausible economic DAG using MORG 2016 data. While we present one DAG here, the methodology generalizes to any number of candidate structures.

Illustrative DAG: Teen Employment with Confounding. Consider a model where teenage employment (Em) is influenced by gender (G), education level (E), and metropolitan status (M), with family background serving as an unobserved common cause. The graphical structure is:

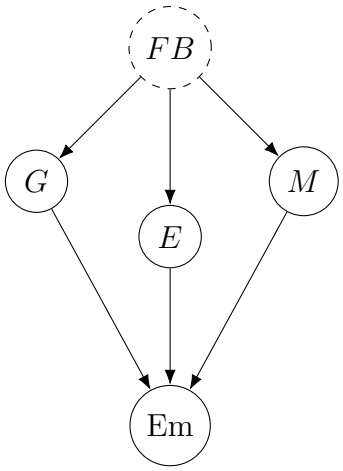


Figure 12: Illustrative DAG: Teen Employment. FB=Family Background (unobserved), G=Gender, E=Education, M=Metro Status, Em=Employment

Applying d-separation to this DAG yields four conditional independence implications. We test these implications against MORG 2016 data (teen subsample, $N = 19,907$) using partial correlation tests. Table 5 presents the results.

The fourth test falsifies its implied conditional independence: conditional on metro status, gender and education remain statistically dependent ($r = -0.0262$, $p = 0.0002$). This violation suggests that family background creates additional confounding between these variables beyond what metro status alone can explain. In a practical setting, this result would prompt DAG revision to better capture the confounding structure. The researcher would then re-test the revised model's implications against the same data, repeating this cycle until the DAG more closely aligns with the empirical evidence.

Test	Conditional Independence	Correlation	P-Value	Result
1	$G \perp M$ (unconditional)	0.0000	1.0000	Consistent
2	$G \perp \text{Em} \mid E, M$	0.0000	1.0000	Consistent
3	$M \perp \text{Em} \mid E$	0.0000	1.0000	Consistent
4	$G \perp E \mid M$	-0.0262	0.0002	**Falsified**
Summary: 3 of 4 tests consistent (75%)				

Table 5: Conditional Independence Tests for Illustrative DAG (Teen Employment). Test 1 shows zero variance in metro status variable within teen subsample. Test 4 falsifies the predicted independence, indicating that family background confounds gender and education beyond what metro status captures.

Interpretation and Extensions

This illustrative application demonstrates how falsification tests guide DAG development. The pattern of results indicates where the initial model succeeds (capturing the relationships between gender, employment, metro status, and education) and where it falls short (underestimating confounding through family background). Rather than accepting or rejecting the model wholesale, researchers systematically identify mismatch between assumed structure and observed patterns. In a full analysis, researchers would construct multiple candidate DAGs, evaluate each against the same data, and select the structure most compatible with the empirical evidence. Falsifications guide this iterative refinement, providing concrete direction for model improvement—a capability unique to the DAG framework compared to traditional goodness-of-fit measures or specification tests that offer limited guidance for systematic revision.

5 Integrating DAGs with Econometric Practice

DAGs are not a replacement for econometrics; they are a complement. This section presents a systematic workflow for integrating DAG-based identification with traditional econometric estimation.

5.1 Step 1: Specify Causal Structure

Task: Draw a DAG representing your causal assumptions.

Guidance:

1. List all relevant variables: treatment (X), outcome (Y), measured confounders, potential mediators, and unobserved factors

2. Include unobserved factors explicitly as dashed nodes or annotate them clearly

3. Draw arrows for direct causal effects only (if A affects B only through C, draw $A \to C \to B$, not $A \to B$)

4. Make causal order explicit: which variables are pre-treatment, simultaneous, or post-treatment?

Example: For education and wages:

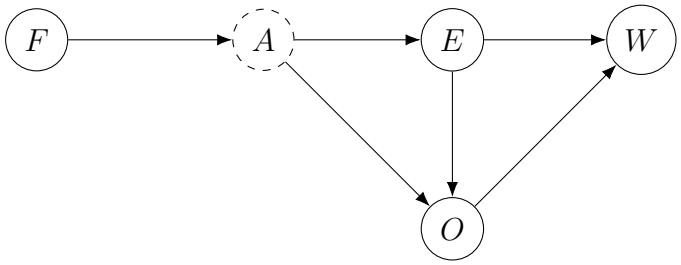


Figure 13: Note: Ability (A) is unobserved. F=Family Background, E=Education, W=Wages, O=Occupation

Explicit documentation of causal assumptions enables critical evaluation. Assumptions can be debated and refined before estimation rather than after results are reported. The careful reader may also note that the DAG used here differs slightly from the DAG used earlier that also seeks to quantify the effect of education on wages. This is standard practice, since typically many viable DAGs will be created (or even hypothesized through causal discovery techniques, which are beyond the scope of this project) and evaluated to identify which models are possible given the data.

5.2 Step 2: Test Assumptions via D-Separation

Task: Derive and test conditional independence implications from the DAG.

Process:

1. Use d-separation rules to identify which variables should be independent (conditionally or unconditionally) given the DAG

2. Test each implication in your data

3. If any test strongly rejects, revise the DAG and iterate

Example Tests:

- **Partial correlation test:** If X and Y should be independent conditional on Z, estimate $\text{corr}(X, Y \mid Z)$ via regression residuals and test if significantly different from zero
- **Conditional independence test:** Regress X on Z and save residuals $\tilde{X}$. Regress Y on Z and save residuals $\tilde{Y}$. Test if $\text{corr}(\tilde{X}, \tilde{Y}) = 0$.

This step validates assumptions before main estimation, preventing post-hoc surprises.

5.3 Step 3: Identify Estimation Strategy

Apply identification criteria to determine what and how to estimate.

For Standard Confounding (Backdoor Criterion):

1. Enumerate all backdoor paths from X to Y
2. Find a minimal set of variables blocking all paths (a valid adjustment set)
3. Verify no post-treatment variables are included

Example: For education-wages with Family Background confounding:

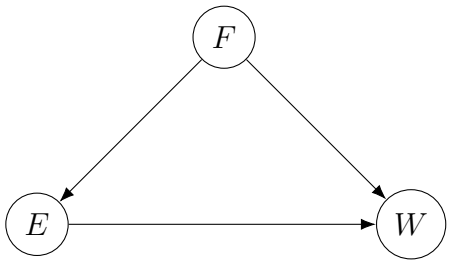


Figure 14: Backdoor path: $E \leftarrow F \rightarrow W$. Valid adjustment set: $\{F\}$. (F=Family Background, E=Education, W=Wages)

Proceed with OLS: `Wages` $\sim$ `Education` + `Family_Background`

For Unobserved Confounding (Frontdoor Criterion):

1. Check if a mediator M intercepts all directed paths from X to Y
2. Verify no confounding of X and M
3. Verify no confounding of M and Y (conditional on X)
4. Use frontdoor formula: decompose into β_{XM} and $\beta_{MY|X}$

For Complex Structures (Do-Calculus):

Apply Rules 1-3 iteratively. If the procedure reduces $P(Y \mid \text{do}(X))$ to observable quantities, identification is possible.

5.4 Step 4: Estimate Using Standard Methods

Use OLS, IV, or other estimators with the covariate set identified in Step 3. DAGs tell you *what* to estimate (identification); econometrics tells you *how* (estimation).

Example:

1. If backdoor satisfied with controls Z: Estimate via OLS $Y \sim X + Z$
2. If frontdoor applies: Estimate $M \sim X$ (first stage) and $Y \sim M + X$ (second stage), then form $\hat{\beta} = \hat{\gamma} \times \hat{\delta}$
3. Standard errors and confidence intervals follow usual econometric procedures

5.5 Step 5: Document the Analysis

Finally, present the causal model, assumptions, and identification strategy alongside results, including

- **Causal Model Diagram:** Include the DAG with labeled variables and causal direction clearly marked
- **Stated Assumptions:** Explicitly list which causal relationships are assumed to exist, which are blocked, and what unobserved factors are included
- **Identification Justification:** Reference the specific criterion (backdoor, frontdoor, do-calculus) and explain how the adjustment set satisfies it
- **Assumption Validation:** Present results from conditional independence tests showing whether the model's implications are consistent with the data
- **Estimation Results:** Report coefficients, standard errors, and confidence intervals using standard econometric procedures

This approach separates causal identification (the DAG and assumptions) from statistical estimation and inference. It makes the reasoning transparent and enables evaluation of both the causal model and the econometric methods independently.

6 Conclusion

6.1 Summary of Findings

This paper demonstrates that Directed Acyclic Graphs provide systematic, algorithmic guidance for covariate selection in causal inference. Our Monte Carlo simulations across three pedagogically important scenarios reveal:

1. **Where Approaches Diverge:** DAG-guided and traditional approaches diverge whenever colliders or complex causal pathways are present. In such cases, DAG reasoning prevents errors; intuition fails.

2. **Magnitude of Bias:** Collider bias introduces substantial distortion: DGP 1 shows bias of -0.50 (compared to true effect of 0.0), while M-bias introduces bias of -0.068 when the true effect is 0.4 (DGP 3). These misspecifications are not negligible.

3. **Frontdoor Power:** The frontdoor criterion enables identification despite unmeasured confounding, recovering true effects when traditional methods fail (DGP 2).

4. **Falsification Value:** Conditional independence tests derived from DAG structure validate assumptions before main estimation, a unique advantage.

6.2 Implications for Econometric Practice and Teaching

On Identification and Covariate Selection:

The separation of identification from estimation enables systematic reasoning about which covariates matter for causal inference. DAGs provide a formal language for this reasoning. The simulations show that ad hoc covariate choices (controlling for all available variables or following general rules) can introduce substantial bias in realistic causal structures.

On Integration of Methods:

DAGs provide tools for identification, which is determining which estimands are identifiable and what information is needed to identify them. Traditional econometrics provides tools for estimation and implementing these procedures efficiently with appropriate inference. These are complementary activities rather than competing approaches.

For Teaching Causal Inference:

The three DGPs illustrate concrete cases where intuitive reasoning fails and systematic reasoning succeeds. Cases like collider bias demonstrate that controlling for "confounders" without understanding the underlying causal structure can backfire. Conversely, the frontdoor example shows that identification is sometimes possible despite confounding. Teaching

these cases moves instruction beyond memorized rules toward principled reasoning about identification.

6.3 Future Research

Several natural extensions merit investigation:

- **Sensitivity Analysis:** How do results change if the DAG is slightly misspecified?
- **Automated DAG Specification:** Tools to help researchers formulate causal structures interactively
- **Dynamic Settings:** Extension to time-series and panel data with feedback loops
- **Complexity and Replication:** Can we understand why certain DAG structures or model features lead researchers to different specification choices?
- **Informing Estimation Procedures:** Beyond Pearl's formula or the two-stage front-door criterion estimation, what other estimation procedures follow from the DAG framework?

6.4 Final Thought

Causal inference requires clear thinking about causal structure. Directed Acyclic Graphs formalize this thinking, replacing ad-hoc intuition with systematic logic. When integrated with traditional econometric estimation, they improve research quality through:

1. **Systematic identification** of what to estimate
2. **Explicit assumptions** that are transparent and falsifiable
3. **Efficient estimation** using well-developed econometric methods
4. **Rigorous validation** via testable implications

The framework is not "DAGs vs. econometrics" but "DAGs + econometrics for robust causal estimation."

References

Marc F. Bellemare, Jeffrey R. Bloem, and Noah Wexler. The paper of how: Estimating treatment effects using the front-door criterion. *Oxford Bulletin of Economics and Statistics*, 2024. Forthcoming.

Felix Elwert and Christopher Winship. Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology*, 40:31–53, 2014.

Yimin Huang and Marco Valtorta. Pearl's calculus of intervention is complete. *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, pages 217–224, 2006.

Guido W Imbens and Donald B Rubin. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015.

H. Lee, J. K. Aronson, and D. Nunan. Collider bias. In *Catalogue Of Bias*. Catalogue of bias collaboration, 2019. URL https://catalogofbias.org/biases/collider-bias/.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

Judea Pearl and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018.

Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

Ilya Shpitser and Judea Pearl. Identification of joint interventional distributions in recursive semi-markovian causal models. *Proceedings of the 21st National Conference on Artificial Intelligence*, 2:1219–1226, 2006.

Robert R Tucci. Introduction to judea pearl's do-calculus. *arXiv preprint arXiv:1305.5506*, 2013.