Redesigning Policies for Lasting Impact with Causal Graphs: Evidence from Uganda's Youth Opportunities Program

Arielle Anderer
Samuel Curtis Johnson Graduate School of Management, Cornell University, aanderer@cornell.edu

Nur Kaynar*
Samuel Curtis Johnson Graduate School of Management, Cornell University, nur.kaynar@cornell.edu

Dmitry Mitrofanov
Carroll School of Management, Boston College, dmitry.mitrofanov@bc.edu.

Problem definition. Anti-poverty programs — cash transfers, vocational training, asset grants — aim to produce durable gains in income, health, and education for poor and underemployed populations. These gains take years, often a decade, to materialize, while the operational decisions that shape them are made on much shorter cycles. To iterate effectively, policymakers need methods that estimate long-term impact from short-term data and identify which short-term changes carry that impact, so the program can be redesigned around the mechanisms that drive durable gains. **Methodology/results.** We develop a framework that combines short-term experimental data with long-term observational data to estimate the long-term effect of a program and rank design modifications by their predicted long-term gain. We learn the causal mechanisms linking the treatment, intermediate (surrogate) outcomes, and long-term outcomes. These mechanisms are encoded in a *causal graph*, recovered using our COMB-PC algorithm. Using the learned graph, we identify a valid surrogate set, yielding a consistent estimator for the long-term effect. We then decompose the long-term effect into directed pathways using a graph-constrained structural model in an interpretable way, and rank design modifications by their predicted long-term gain. Applied to the Uganda Youth Opportunities Program (YOP), our method identifies the welfare effects that persist and those that fade, recovering the nine-year impact from only two-year surrogates which are based on two years of post-disbursement data. We find that most of the lasting effect operates through a single mechanism: enrollment in vocational training. Counterfactual analysis on the learned graph shows that strengthening this channel is expected to produce the largest long-term gains. **Managerial implications.** For policymakers designing the next pilot or scale-up, our framework converts short-horizon experimental evidence into actionable guidance on both whether a program's gains will persist and what design changes improve the next iteration. For YOP, the decomposition isolates a single high-leverage channel — early enrollment in vocational training — and points to a redesign built around raising enrollment: bundling the cash grant with enrollment vouchers, transportation support, or direct connections to training providers.

Key words: causal inference; surrogates; long-term policy effects; machine learning; development operations; natural experiment; empirical findings

* Corresponding author.

1. Introduction

Whether an anti-poverty or human-development intervention should be launched, refined, scaled, or abandoned hinges on its effect on welfare outcomes such as durable gains in income, health and education that typically materialize over years or decades, long after these decisions must be made. These interventions, including cash transfers, vocational programs, microfinance, and preventive health campaigns, are deployed at scale by governments, NGOs, and multilateral organizations across the developing world (Banerjee et al. 2015, Parker and Todd 2017, McKenzie 2017, Baird et al. 2016). However, long-term evaluation presents an operational challenge. Tracking subjects over long horizons is costly in environments with limited data infrastructure and fragmented administrative records, and attrition compounds the problem as study populations move, migrate, or become unreachable (Molina-Millán and Macours 2025). Moreover, even when long-term measurement is feasible in principle, the funding and political cycles that determine whether a program is piloted, scaled, modified, or wound down move on much shorter timescales than the outcomes themselves.

The Uganda Youth Opportunities Program (YOP), which we return to throughout this paper as a running example, is a case in point. Beginning in 2008, the Government of Uganda ran a program offering one-time cash grants to groups of young adults to fund vocational training and start-up costs for skilled trades, with the long-term goal of raising employment and incomes among an underemployed rural population (Blattman et al. 2014). Four years in, it seemed the program had a large impact: relative to a control group, treated individuals reported 38% higher earnings, 11% higher consumption, and substantially greater investment in business capital and skilled work. Nine years in, however, these effects had mostly dissipated; the two groups had converged on employment, earnings, and consumption (Blattman et al. 2020). Lasting effects appeared only on a narrower set of outcomes: durable assets and sustained occupation in a skilled trade. The conclusion was that the grants acted more as a "kick start" than a lift out of poverty.

This pattern of fading effects is not unique to YOP. Similar programs that grant cash, physical assets, or training to poor or underemployed populations have become widespread over the past two decades. They vary widely in what is granted, to whom, and what services are bundled, and there is little consensus on which design choices drive long-term outcomes. Across this body of evaluations, short-term results are often encouraging, but long-term data remain rare (Blattman et al. 2020) — and the decision to fund these programs typically rests on whether their short-term gains persist (Banerjee et al. 2015).

Decision-makers selecting among variants for the next pilot or scale-up therefore need two things from short-term data. First, they need faster and more reliable estimates of long-term impact, so

that resources are not committed or withdrawn on the basis of an inaccurate picture of long-term impact. Second, they need to understand the mechanism through which an intervention produces (or fails to produce) durable effects, i.e., which short-term changes mediate the long-term impact, and which of these the current treatment fails to engage. These two needs are complementary. An estimate of long-term impact tells decision-makers what the effect of an intervention will be, while understanding the mechanism identifies which features of the intervention to change or expand. This motivates the research question we address in this paper: How can we use short-term data to forecast an intervention's long-term effects, identify the mechanisms that drive them, and guide the design of follow-on interventions?

Existing methods do not address the question of understanding the mechanism. The most prominent of these methods is the surrogacy framework introduced by Athey et al. (2019). The framework relies on surrogate variables: intermediate outcomes (i.e., outcomes measured between the treatment and the long-term outcome) that mediate the treatment's effect on that outcome. In the YOP setting, candidate surrogates might include earnings, business capital, or occupational status measured two or four years after the grant. The surrogacy framework combines short-term experimental data, in which the treatment and the surrogates are observed but the long-term outcome is not, with longer-term observational data containing both the surrogates and the outcome, together yielding an estimate of the long-term average treatment effect before it is realized. A growing literature has extended the framework in several directions (Yang et al. 2024, Athey et al. 2020, Imbens et al. 2022, Huang et al. 2026, Battocchi et al. 2021, Anderer et al. 2022). Across these extensions, however, the analyst is assumed to arrive with prior knowledge about which short-term variables carry information about the long-term outcome, whether as strict surrogates that mediate the treatment effect, or as informative proxies whose relationship to the long-term outcome has been characterized in advance. In settings where these relationships are themselves well understood, this assumption is reasonable. In the development context, however, in-depth knowledge of such relationships is often a binding constraint. The channels through which interventions generate long-term outcomes are often poorly understood and contested, especially in developing-economy contexts where data infrastructure, administrative records, and information systems are less advanced. This makes it difficult not only to measure long-term impacts, but also to understand why they arise and how programs should be redesigned. As a result, identifying which mechanism is actually doing the work is often precisely what investigators, policymakers, and local authorities most need to learn.

We address this constraint by learning the surrogate outcomes and the underlying mechanism from data rather than assuming them. Specifically, we combine causal structure learning with the surrogacy framework to recover both a set of valid surrogates and the causal pathways connecting the treatment to the long-term outcome. Causal structure learning, unlike causal inference, focuses

on identifying the presence or absence of causal relationships among variables using the framework of graphical causal models (Pearl 2000, Spirtes et al. 2000). Building on this literature, we develop COMB-PC, an algorithm that combines short-term experimental data with longer-term observational data to recover a directed acyclic graph (DAG) representing the causal relationships among the treatment, candidate short-term variables, and the long-term outcome. The resulting causal graph helps determine (*i*) which short-term variables are influenced by the treatment, and (*ii*) which of those variables, in turn, affect the long-term outcome — that is, which short-term variables serve as valid surrogates. Given the learned graph, we derive a closed-form, non-parametric expression for the long-term average treatment effect and show that it can be consistently estimated from the experimental and observational data under standard assumptions in the surrogacy literature. The graphical representation also has a practical advantage in development-policy settings: at a glance, the policymaker sees which channel carries the program's effect and how to strengthen it.

Beyond enabling consistent estimation by identifying valid surrogates, the learned graph also enables counterfactual reasoning about the design of potential interventions. Combining the learned graph with structural equation models, we can estimate how the long-term outcome would respond to changes in the treatment's effect on each short-term variable. For example, this would allow us to quantify what would happen to the long-term outcome if we strengthen the treatment's effect on a surrogate by a specific amount. This in turn opens up two avenues by which to refine the design of further interventions to improve these long-term outcomes. First, the graph identifies short-term variables that carry the treatment's effect to the long-term outcome, enabling incremental improvements along established intervention pathways. Second, and perhaps more distinctively, the graph can also identify short-term variables that affect the long-term outcome but may not currently be engaged by the treatment, uncovering novel opportunities for intervention.

We apply the framework to YOP, using 2- and 4-year measurements of earnings, business capital, and occupational choice as candidate surrogates and 9-year outcomes as long-term targets. We construct two samples matching the surrogacy data structure: an experimental sample with treatment and short-term measurements, and an observational sample with all three timepoints but no treatment information. COMB-PC recovers a directed graph linking the cash grant to long-term outcomes through specific short-term pathways. Our long-term treatment effect estimates match the original 9-year follow-up: treatment and control groups converge on earnings and consumption, while persistent effects remain on durable assets and skilled-trade engagement (Blattman et al. 2020). We achieve this match using only the first two years of data, demonstrating that long-term effects can be estimated from short-term measurements. We also find that naively using all short-term variables as surrogates without graph-based selection produces false-positive estimates. The learned graph identifies year-2 vocational training enrollment as the dominant pathway, carrying

most of the long-term effect on its own, with training hours and skilled-trade participation acting as downstream amplifiers. Counterfactual analysis confirms that strengthening the treatment's effect on enrollment yields the largest long-term gains, suggesting that future interventions should boost early vocational training enrollment.

1.1. Contributions

This paper develops a framework for using short-term experimental evidence to both estimate and improve the long-term impact of development interventions. Our central contribution is to turn the surrogacy framework from a tool for predicting the long-term effect of a fixed treatment into a tool for learning the causal mechanisms through which a program creates durable welfare gains and using those mechanisms to guide redesign. We make three contributions.

- *Causal structure learning for two-sample surrogacy.* Building on Imbens (2020), who highlights surrogacy as a natural setting in which graphical causal models can clarify the assumptions underlying long-term treatment-effect estimation, we push this synthesis further by learning the causal structure from two-sample data and using it to select surrogates, identify long-term effects, and guide program redesign. We introduce COMB-PC, a causal-structure learning algorithm for the two-sample surrogacy setting in which treatment assignment and long-term outcomes are never jointly observed. This setting arises naturally in development-policy evaluation: a short-term experiment observes treatment assignment and early outcomes, while a separate observational sample contains early and long-term outcomes. Standard causal-discovery algorithms cannot be applied directly because no single sample contains all variables. COMB-PC resolves this problem by combining conditional-independence information from the experimental and observational samples to recover the causal structures up to an equivalence class linking treatment, intermediate outcomes, and long-term outcomes.
- *Graph-selected surrogates, long-term effect identification, and redesign.* Building on the learned graph, we develop a non-parametric identification strategy for long-term treatment effects. The graph delivers two objects that existing surrogate-index approaches typically take as given or do not distinguish: a mediating surrogate set that intercepts directed treatment effects, and a surrogate–outcome adjustment set that blocks noncausal paths between the surrogates and the long-term outcome. This distinction is important because short-term variables may be predictive of long-term outcomes without being valid mediators of the treatment effect. The resulting closed-form estimand combines the short-term experimental sample with the long-term observational sample to estimate long-term treatment effects before the long-term outcome is observed in the experimental sample. We then use the same learned structure to move from evaluation to design. By combining the graph with a structural model, we decompose long-term effects into directed pathways and evaluate mechanism-level counterfactuals, asking how

long-term impact would change if a redesigned program strengthened its effect on a particular short-term mechanism. When the graph is only partially identified, we evaluate these redesign counterfactuals across the graph equivalence class, yielding robust recommendations under uncertainty about the causal structure.

- *Evidence from Uganda's Youth Opportunities Program.* We apply the framework to Uganda's Youth Opportunities Program, a cash-grant intervention whose short-run gains largely faded by the nine-year follow-up except for durable assets and skilled-trade attachment. We construct a two-sample design by splitting districts into an experimental sample containing treatment and short-term outcomes and an observational sample containing short- and long-term outcomes but no treatment assignment. Using two- and four-year outcomes as candidate surrogates, the graph-selected estimator reproduces the nine-year patterns: effects fade for earnings, consumption, and employment, but persist for durable assets and skilled-trade outcomes. We then show that the same qualitative pattern can be recovered using only year-two experimental outcomes, substantially shortening the experimental learning cycle. In contrast, using all candidate short-term variables as surrogates without graph-based selection identifies significant effects on outcomes that are not significant in the nine-year randomized benchmark. Finally, the learned mechanism identifies year-two vocational training enrollment as the dominant early pathway behind persistent effects, with training hours and skilled-trade participation acting mainly as downstream amplifiers. Model-based counterfactuals suggest that follow-on programs should prioritize increasing early vocational-training enrollment rather than only increasing training intensity.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 formally defines the problem and the assumptions on which our results rest. Section 4 presents COMB-PC, our algorithm for learning the causal structure from short-term experimental and longer-term observational data. Section 5 derives the long-term treatment effect identification strategy using the learned graph. Section 6 evaluates counterfactual redesigns under a graph-constrained structural model. Section 7 applies the framework to the Uganda Youth Opportunities Program. Section 8 concludes with a discussion of implications and limitations.

2. Literature Review

This paper builds on three streams of literature: (i) long-term welfare effects of development programs, (ii) causal and prescriptive methods for decision-making, and (iii) surrogate-based estimation of long-term treatment effects.

Long-term welfare effects of development programs. This literature studies large-scale anti-poverty programs, where the central puzzle is that short-term gains often fail to predict long-term impacts.

The Youth Opportunities Program (YOP) in northern Uganda, which we revisit in Section 7, gave one-time grants of roughly USD 400 to underemployed young adults; Blattman et al. (2014) found earnings and hours gains at four years, but by the nine-year follow-up these had faded as the control group caught up, leaving only gains in durable assets and skilled-trade attachment (Blattman et al. 2020). Multifaceted "Graduation" programs show a related pattern: one-year gains in consumption, food security, and income (Banerjee et al. 2015) were preserved or amplified in some countries and faded in others by year ten (Banerjee et al. 2021). Compounding this, few development RCTs are followed over long horizons (Bouguen et al. 2019). What these evaluations cannot provide is a way to tell, from short-term data, which gains will persist and which early changes carry them.

Operations management has examined related questions in developing-country settings, asking how operational and process design choices shape welfare outcomes. Deo et al. (2013), Natarajan and Swaminathan (2014), and Aflaki and Pedraza-Martinez (2016) study operational decisions in chronic-care capacity and donor-funded humanitarian health programs, where welfare consequences arrive long after operational choices are locked in. Closer to our development setting, Uppari et al. (2024) combine field experiments in Rwanda with a structural model to show that operations-based strategies for an off-grid lighting business yield far larger improvements in lamp usage than price-based alternatives. Ramdas and Sungu (2024) run a field experiment in Mumbai showing that a simple operational lever (shorter data-replenishment cycles for mobile plans) increases access to healthcare information and attendance at health camps. Other work examines welfare-relevant interventions in resource-limited settings, including cash assistance for refugees (Kotsi et al. 2022), agricultural market reforms (Levi et al. 2020), food subsidies (Aouad et al. 2024), drone delivery of blood products in Rwanda (Jeon et al. 2026), and spatial resource allocation during epidemics (Long et al. 2018). Other work documents welfare inequities, including post-disaster price gaps borne by low-income communities (Barriola and Schmidt 2025) and gender disparities in public-service access (Kaaua and Virudachalam 2025). However, this work measures effects only after they materialize, leaving a structural mismatch between the cycle on which long-term welfare is measured and the cycle on which program-design decisions are made. We address this mismatch by combining short-term experimental data with long-term observational data to inform program-design decisions, compressing the cycle on which decisions are made.

Causal and prescriptive methods for decision-making. Within operations management, decision-making has been supported by heterogeneous treatment-effect estimation in healthcare (Wang et al. 2021), prescriptive analytics with observational data and covariates (Bertsimas and Kallus 2020), and decision rules that rely on structural assumptions when experimentation is limited (Bastani et al. 2021). Causal discovery addresses a different but complementary problem: learning the causal graph from conditional-independence patterns and structural restrictions. This literature is well developed

in statistics and computer science (Spirtes et al. 2000, Chickering 2002, Heinze-Deml et al. 2018, Cussens et al. 2017), with extensions to heterogeneous and overlapping data sources (Bareinboim and Pearl 2016, Mooij et al. 2020, Triantafillou and Tsamardinos 2015, Huang et al. 2020) and applications to instrumental-variable validation (Eberhardt et al. 2025). We develop COMB-PC, a causal discovery algorithm for the two-sample surrogacy setting in which treatment and long-term outcomes are never jointly observed. COMB-PC adapts the PC algorithm (Spirtes et al. 2000) by assigning each conditional-independence query to the sample in which the relevant variables are jointly observed. The recovered graph yields a graph-selected surrogate set, a surrogate–outcome adjustment set, and a closed-form identification formula for the long-term treatment effect, extending the backdoor logic of Pearl (1995) and equivalence-class adjustment results (Maathuis et al. 2009, Maathuis and Colombo 2015, Perković et al. 2018) to the two-sample surrogacy structure.

To our knowledge, this is the first framework for two-sample surrogate identification that learns both the surrogate set and the surrogate–outcome adjustment set from a causal graph, and then uses that graph to support mechanism-guided counterfactual program design. The framework builds on the perspective in Imbens (2020) that graphical causal models can clarify the identifying assumptions in surrogacy settings, and uses this structure to discipline the choice of surrogate variables. It also addresses a common critique of causal discovery — that the methodology has remained largely separate from applied decision-making — by using the recovered graph as the foundation for real-world policy design.

Surrogate-based estimation of long-term treatment effects. These methods use short-term measurements as surrogates for long-term outcomes, building on the framework introduced by Prentice (1989). Most relevant to our setting, Athey et al. (2025) treat short-term measures as a joint surrogate index, combining an experimental sample (treatment and surrogates observed) with an observational sample (surrogates and outcome observed). Subsequent work extends this framework: Imbens et al. (2025) address persistent confounding using sequential short-term outcomes; Huang et al. (2026) extend to long-term treatments whose effects unfold over future horizons; Kallus and Mao (2025) derive efficiency bounds without requiring strong surrogacy assumptions. Operations management has adopted these tools for decisions made before long-term outcomes arrive: Yang et al. (2024) learn targeting policies for long-term retention using surrogate-imputed outcomes, and Anderer et al. (2022) combine imperfect surrogate information with true outcomes in Bayesian adaptive trial designs. Across these approaches, the surrogate set is taken as given. This is defensible when mechanisms are understood, but less so in development settings, where the channels producing long-term welfare gains are themselves the object of inquiry. Our framework learns the surrogate set from data so that surrogacy follows from the recovered graph rather than being assumed, augments Athey et al. (2025)'s surrogate index with a graph-derived adjustment set that closes post-treatment

paths, and moves from forecasting long-term effects to designing for them by decomposing those effects along learned causal paths.

3. Problem Setup

Consider a setting with two samples: an experimental sample (denoted as E) and an observational sample (denoted as O). The experimental sample comprises N_E observations and includes a binary treatment variable W together with a set of covariates, but does not include the long-term outcomes $\mathbf{Y}$, which are observed only at a longer horizon.[1] The observational sample comprises N_O observations and includes both $\mathbf{Y}$ and the covariates, but does not include the treatment W. We use the indicator $D \in \{E, O\}$ to denote the sample under consideration.

We partition the variables by the time at which they are collected. Let $\mathbf{X}$ denote the set of *pre-treatment* (baseline) covariates, collected before treatment assignment, and let $\mathbf{S} = \{S_1, \ldots, S_m\}$ denote the set of *post-treatment short-term variables*, collected after treatment but before the long-term outcomes $\mathbf{Y} = \{Y_1, \ldots, Y_q\}$. Variables in $\mathbf{S}$ act as candidate surrogates that may or may not mediate the effect of W on $\mathbf{Y}$. Both $\mathbf{X}$ and $\mathbf{S}$ are observed in both samples; only W and $\mathbf{Y}$ are sample-specific. The variables in each sample are $\mathbf{V}^E = \mathbf{X} \cup \{W\} \cup \mathbf{S}$ and $\mathbf{V}^O = \mathbf{X} \cup \mathbf{S} \cup \mathbf{Y}$, and the full set across both samples is $\mathbf{V} = \mathbf{X} \cup \{W\} \cup \mathbf{S} \cup \mathbf{Y}$. We assume that both samples are generated by the same underlying causal mechanism, represented by a directed acyclic graph $\mathcal{G}^*$ over $\mathbf{V}$. We do not observe $\mathcal{G}^*$; we observe only the data it generates, recorded separately in the experimental and observational samples.

3.1. Identifying Assumptions for Combining Experimental and Observational Data

We now state the assumptions on the structure of $\mathcal{G}^*$ and on the relationship between the two samples, following standard conditions in the surrogate-index literature (Athey et al. 2025).

Assumption 1 **(Common Data-Generating Process)**. *The experimental and observational samples share the same underlying causal mechanism $\mathcal{G}^*$ and the structural equations governing every variable in $\mathbf{V} \setminus \{W\}$ are identical across the two samples.*

This is what allows us to combine information across samples. The experimental sample identifies the part of $\mathcal{G}^*$ involving W, while the observational sample identifies the part involving $\mathbf{Y}$; the common mechanism ties them together.

The next assumption formalizes the experimental randomization.

Assumption 2 **(Randomization on Observables)**. *The treatment W is randomized conditional on the pre-treatment covariates $\mathbf{X}$, implying that any incoming edge to W in the true underlying graph $\mathcal{G}^*$ originates from a variable in $\mathbf{X}$.*

[1] In this paper, we focus on a binary treatment variable W and discrete outcomes $\mathbf{Y}$, with $\mathbf{X}$ and $\mathbf{S}$ representing sets of discrete variables. Our results can be extended to the continuous setting.

In particular, the experimental design ensures that W has no unobserved causes, so the experimental sample identifies the causal effect of W on $\mathbf{S}$ once we condition on $\mathbf{X}$.

Because our experimental sample lacks the long-term outcomes, we further restrict the causal structure over $\mathbf{V}$ so that long-term effects can be inferred from observational data.

ASSUMPTION 3 **(Causal Structure of $\mathcal{G}^*$).** *The true causal graph $\mathcal{G}^*$ satisfies:*

(i) No directed edge points from a later stage to an earlier stage.

(ii) There is no direct edge from W to any long-term outcome $Y_k \in \mathbf{Y}$.

Part (i) encodes the structural ordering of the variable categories: short-term variables cannot affect $\mathbf{X}$ or W, and long-term outcomes cannot affect any earlier variable. Part (ii) imposes the surrogacy restriction—the treatment's effect on the long-term outcomes is mediated entirely through the short-term variables $\mathbf{S}$. This is plausible in development settings, where such interventions are one-time and do not recur, so their effect on long-run welfare must be transmitted through the intermediate states they alter—precisely the channels $\mathbf{S}$ is constructed to measure.

4. Learning Treatment-Effect Mechanisms and Surrogates

This section presents COMB-PC, a causal structure learning algorithm that recovers the directed acyclic graph linking the treatment, the short-term variables, and the long-term outcomes, using a combined short-term experimental sample and long-term observational sample. Section 4.1 provides a brief background on causal structure learning, including the standard identification assumptions and terminology we use throughout the rest of the paper.[2] Readers familiar with this material may skip directly to Section 4.2.

4.1. Background: Causal Structure Learning

Graphical causal models (Pearl 2000, Spirtes et al. 2000) use directed graphs to represent causal relationships among variables. Let $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ be a directed graph with node set $\mathbf{V}$ and edge set $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$, where each edge signifies a direct causal connection between the corresponding nodes. Two nodes V_i and V_j are *adjacent* if there is an edge $V_i \to V_j$ or $V_i \leftarrow V_j$ in $\mathcal{G}$. The *parents* and *children* of a node V_i are its direct causes and effects, respectively. A *path* between V_i and V_j is a sequence of distinct nodes starting at V_i and ending at V_j in which consecutive nodes are adjacent; a *directed path* is one in which every edge points in the direction of traversal. A *descendant* of V_i is any node reached from V_i by a directed path, and an *ancestor* of V_i is any node from which V_i can be reached by a directed path. A *directed acyclic graph (DAG)* is a directed graph with no

[2] Appendix A gives a more detailed background on graphical causal models, including graphical terminology, the causal Markov, faithfulness, and sufficiency assumptions, d-separation, Markov equivalence, and back-door adjustment, which form the basis for structure-learning and the identification results of the paper.

directed cycles. A node V_j is a *collider* on a path if its adjacent edges on that path point into V_j, i.e., $\rightarrow V_j \leftarrow$; otherwise it is a *noncollider*, either a *mediator* ($\rightarrow V_j \rightarrow$) or a *common cause* ($\leftarrow V_j \rightarrow$). A *v-structure* (or *unshielded collider*) consists of two non-adjacent nodes that both point into a common child, forming a "V." The *skeleton* of $\mathcal{G}$ is the undirected graph obtained by replacing each directed edge with an undirected edge.

Causal modeling involves associating each graph $\mathcal{G}$ with a probability distribution $P_{\mathcal{G}}(\mathbf{V})$ that factorizes according to the graph structure: $P_{\mathcal{G}}(\mathbf{V}) = \prod_{V_i \in \mathbf{V}} P_{\mathcal{G}}(V_i \mid Pa(V_i))$, where $Pa(V_i)$ denotes the parents of V_i in $\mathcal{G}$ (Eberhardt 2017). Here the terms "node" and "variable" can be used interchangeably, as they correspond to the graphical structure and the probability distribution, respectively. The following three assumptions bridge the observed data and the causal structure.

Assumption 4 **(Causal Markov)**. *Each variable $V_i \in \mathbf{V}$ in a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ is probabilistically independent of its non-descendants given its parents.*

Assumption 5 **(Faithfulness)**. *The only independences present in the probability distribution are those that are implied by the graph structure through the causal Markov conditions.*

Assumption 6 **(Causal Sufficiency)**. *For any pair of variables in $\mathbf{V}$, all common causes of those variables are also contained within $\mathbf{V}$.*

The causal Markov condition permits us to transition from the causal graph to the observed probabilistic independencies. Conversely, faithfulness enables us to deduce the structure of the causal graph from observed data independencies. Causal sufficiency ensures that all common causes of any pair of variables in $\mathbf{V}$ are contained within $\mathbf{V}$, thereby excluding hidden or unobserved confounders. Together, these assumptions allow causal structure learning algorithms to recover the graph from observational data by exploiting the conditional independence signatures that different graph structures leave in the joint distribution.

Two DAGs are *Markov equivalent* if they entail the same conditional independencies, and *Markov equivalence class* of a DAG is the set of DAGs Markov equivalent to it. Verma and Pearl (1990) characterize this relation graphically: two DAGs are Markov equivalent if and only if they share the same skeleton and the same v-structures. Causal structure learning algorithms therefore typically return the Markov equivalence class of the underlying graph rather than a single DAG. Identifying a unique DAG from observational data alone requires additional parametric restrictions on the structural equations, such as non-Gaussian noise (Shimizu et al. 2006) or additive noise models (Peters et al. 2014).

4.2. The COMB-PC Algorithm

To learn the causal mechanism by which the treatment W affects the long-term outcomes $\mathbf{Y}$, we develop COMB-PC, a causal structure learning algorithm tailored to the two-sample surrogacy setting. The main challenge is that W and $\mathbf{Y}$ are never jointly observed: the experimental sample contains $\mathbf{X}, W, \mathbf{S}$, while the observational sample contains $\mathbf{X}, \mathbf{S}, \mathbf{Y}$. Standard causal structure learning methods, which assume a single sample with all variables jointly observed, cannot be applied directly to this setting. We show that, under standard correctness conditions on the independence tests, COMB-PC asymptotically recovers the Markov equivalence class of the true data-generating graph.

In the first phase of the COMB-PC algorithm, we focus on identifying the *skeleton* of the causal graph; Algorithm 1 provides the pseudocode. This phase begins with the initialization of an undirected graph $\mathcal{G}_1 = (\mathbf{V}, \mathbf{U})$ where $\mathbf{U}$ consists of all possible undirected edges between variables in $\mathbf{V}$, except for the edges between the treatment W and each long-term outcome $Y_k \in \mathbf{Y}$, which are ruled out by Assumption 3 (*ii*). For each $V_i \in \mathbf{V}$, we let $\mathcal{N}(V_i)$ denote the neighborhood of V_i in the current graph, i.e., the set of variables adjacent to V_i via an undirected edge. The neighborhoods are initialized to mirror $\mathbf{U}$. Specifically, the initial neighborhoods are $\mathcal{N}(W) = \mathbf{X} \cup \mathbf{S}$, $\mathcal{N}(X_i) = \mathbf{V} \setminus \{X_i\}$ for every $X_i \in \mathbf{X}$, $\mathcal{N}(S_i) = \mathbf{V} \setminus \{S_i\}$ for every $S_i \in \mathbf{S}$, and $\mathcal{N}(Y_k) = \mathbf{V} \setminus \{W, Y_k\}$ for every $Y_k \in \mathbf{Y}$.

Throughout Phase 1 we write $V_i \perp\!\!\!\perp_E V_j \mid \mathbf{C}$ and $V_i \perp\!\!\!\perp_O V_j \mid \mathbf{C}$ to denote the conditional independence of V_i and V_j given the conditioning set $\mathbf{C}$ in the experimental (E) and observational (O) samples, respectively.[3] For every pair $V_i, V_j \in \mathbf{V}^E$ currently adjacent in $\mathcal{G}_1$, we check whether V_i and V_j are conditionally independent given a conditioning set $\mathbf{C}$ in $\mathcal{C}^E_{V_iV_j} = \left\{\mathbf{C} \mid \mathbf{C} \subseteq \mathbf{V}^E \setminus \{V_i, V_j\}\right\}$. [4] If some $\mathbf{C} \in \mathcal{C}^E_{V_iV_j}$ satisfies $V_i \perp\!\!\!\perp_E V_j \mid \mathbf{C}$, we remove the undirected edge between V_i and V_j from $\mathcal{G}_1$, update the corresponding neighborhoods, and record this conditioning set as the *separating set* $SepSet_{V_iV_j} = \mathbf{C}$. We then perform the analogous test for each long-term outcome $Y_k \in \mathbf{Y}$ and each currently adjacent $V_i \in \mathbf{V}^O \setminus \{Y_k\}$ in the observational sample using $\mathcal{C}^O_{V_iY_k} = \left\{\mathbf{C} \mid \mathbf{C} \subseteq \mathbf{V}^O \setminus \{V_i, Y_k\}\right\}$; if $V_i \perp\!\!\!\perp_O Y_k \mid \mathbf{C}$ for some $\mathbf{C}$, we remove the edge, update the neighborhoods, and set $SepSet_{V_iY_k} = \mathbf{C}$. The separating sets and final neighborhoods from both rounds of edge removal are retained as input to Phase 2.

The second phase of the COMB-PC algorithm focuses on orienting the undirected edges within the skeleton returned in the first phase; Algorithm 2 provides the pseudocode. We first apply the collider rule. For each triple $V_i, V_j, V_k \in \mathbf{V}$ with $V_j \neq W$, $V_i \in \mathcal{N}(V_j)$, $V_k \in \mathcal{N}(V_j)$, and $V_i \notin \mathcal{N}(V_k)$, if $V_j \notin SepSet_{V_iV_k}$ then V_j must be a collider between V_i and V_k Pearl (2000). Hence, we orient the

[3] Our procedure is agnostic to the choice of conditional independence test, provided it is consistent for the data-generating process at hand; standard options include Fisher's Z for Gaussian continuous data, the G^2 or χ^2 test for discrete data, and kernel-based tests such as KCI (Zhang et al. 2011) for nonparametric or mixed-type settings.

[4] Testing every such subset scales with $|\mathbf{V}|$; Appendix B gives a stage-aware implementation that uses the ordering $\mathbf{X}$ before W before $\mathbf{S}$ before $\mathbf{Y}$ to restrict the conditioning sets, and proves that later-stage variables can be excluded without affecting the recovered skeleton.

Algorithm 1: COMB-PC - Phase 1 (Skeleton Discovery)

Input: experimental sample E over $\mathbf{V}^E$ and observational sample O over $\mathbf{V}^O$.
Output: $\mathcal{G}_1 = (\mathbf{V}, \mathbf{U})$, $\mathcal{N}(V_i)$ for $V_i \in \mathbf{V}$, $SepSet_{V_iV_j}$ for $V_i, V_j \in \mathbf{V}, V_i \neq V_j$.
Initialization: $\mathbf{U} = \{(V_i — V_j) \mid \forall\, V_i, V_j \in \mathbf{V}^E, V_i \neq V_j\} \cup \{(V_i — Y_k) \mid \forall\, V_i \in \mathbf{V}^O \setminus \{Y_k\}, Y_k \in \mathbf{Y}\}$
$\mathcal{N}(W) = \mathbf{X} \cup \mathbf{S}$, $\mathcal{N}(X_i) = \mathbf{V} \setminus \{X_i\}\ \forall X_i \in \mathbf{X}$, $\mathcal{N}(S_i) = \mathbf{V} \setminus \{S_i\}\ \forall S_i \in \mathbf{S}$,
$\mathcal{N}(Y_k) = \mathbf{V} \setminus \{W, Y_k\}\ \forall Y_k \in \mathbf{Y}$,
$SepSet_{V_iV_j} = \emptyset\ \forall V_i, V_j \in \mathbf{V}, V_i \neq V_j$, $\ell = 0$.

1. **Experimental sample:**
   for a pair $V_i, V_j \in \mathbf{V^E}$ where $V_i \in \mathcal{N}(V_j)$:
      for $\mathbf{C} \in C^E_{V_iV_j}$ where $|\mathbf{C}| = \ell$:
         if $V_i \perp\!\!\!\perp_E V_j \mid \mathbf{C}$:
            Update $SepSet_{V_iV_j} = \mathbf{C}$ and $\mathbf{U} \leftarrow \mathbf{U} \setminus \{(V_i — V_j)\}$.
            Update $\mathcal{N}(V_i) \leftarrow \mathcal{N}(V_i) \setminus \{V_j\}$ and $\mathcal{N}(V_j) \leftarrow \mathcal{N}(V_j) \setminus \{V_i\}$.
            Break.
2. **Long-term outcome integration:**
   for $Y_k \in \mathbf{Y}$:
      for $V_i \in \mathbf{V}^O \setminus \{Y_k\}$ where $V_i \in \mathcal{N}(Y_k)$:
         for $\mathbf{C} \in C^O_{V_iY_k}$ with $|\mathbf{C}| = \ell$:
            if $V_i \perp\!\!\!\perp_O Y_k \mid \mathbf{C}$:
               Update $SepSet_{V_iY_k} = \mathbf{C}$ and $\mathbf{U} \leftarrow \mathbf{U} \setminus \{(V_i — Y_k)\}$.
               Update $\mathcal{N}(Y_k) \leftarrow \mathcal{N}(Y_k) \setminus \{V_i\}$ and $\mathcal{N}(V_i) \leftarrow \mathcal{N}(V_i) \setminus \{Y_k\}$.
               Break.
3. **if** $\ell < |\mathbf{V}| - 3$:
   Update $\ell = \ell + 1$. Go to Step 1.
4. **Return**.

edge $V_i — V_j$ as $V_i \to V_j$ and the edge $V_j — V_k$ as $V_k \to V_j$. We then apply the structural ordering implied by the variable categories: every variable in $\mathbf{X}$ precedes W, both precede every variable in $\mathbf{S}$, and all of these precede every variable in $\mathbf{Y}$. Accordingly, any surviving $X_i — W$ edge is oriented as $X_i \to W$, any surviving edge from $\mathbf{X} \cup \{W\}$ to $\mathbf{S}$ is oriented into the surrogate, and any surviving edge from $\mathbf{V} \setminus \mathbf{Y}$ to a long-term outcome Y_k is oriented into Y_k. The algorithm then applies Meek's rules repeatedly until no further orientations are made, propagating the orientations already determined without introducing new v-structures or directed cycles (Meek 1995). These rules are illustrated in Appendix C. The resulting graph $\mathcal{G}_2 = (\mathbf{V}, \mathbf{M})$ can include both directed and undirected edges. Directed edges represent causal directions that the algorithm successfully identified, whereas undirected edges represent relations where directionality is not conclusively identified, allowing for different DAGs consistent with $\mathcal{G}_2$ to orient those edges in different ways. In the last step of Phase 2, we identify all DAGs consistent with $\mathcal{G}_2$ and store them in $\mathbf{G}$.

Figure 1 illustrates COMB-PC on an example with a single pre-treatment covariate X, four candidate surrogates $\mathbf{S} = \{S_1, S_2, S_3, S_4\}$, and two long-term outcomes $\mathbf{Y} = \{Y_1, Y_2\}$. The leftmost panel shows the initial graph, near-complete except for the $W — Y_k$ adjacencies excluded by surrogacy. Phase 1 prunes adjacencies via conditional independence tests in the two samples, leaving the

Algorithm 2: COMB-PC - Phase 2 (Edge Orientation)

Input: $\mathcal{G}_1 = (\mathbf{V}, \mathbf{U})$, $\mathcal{N}(V_i)$ for $V_i \in \mathbf{V}$, $SepSet_{V_iV_j}$ for $V_i, V_j \in \mathbf{V}, V_i \neq V_j$.
Output: $\mathbf{G}$.
Initialization: $\mathbf{M} = \mathbf{U}$, $\mathcal{G}_2 = (\mathbf{V}, \mathbf{M})$.
1. **Collider orientation:**
   for $V_i, V_j, V_k \in \mathbf{V}$ where $V_i \in \mathcal{N}(V_j), V_j \in \mathcal{N}(V_k)$ and $V_i \notin \mathcal{N}(V_k)$:
      if $\{V_i, V_k\} \neq \{W, Y_\ell\}$ for every $Y_\ell \in \mathbf{Y}$, $V_j \notin SepSet_{V_iV_k}$, and $V_j \neq W$:
         Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_i \to V_j), (V_k \to V_j)\}\big) \setminus \{(V_i — V_j), (V_j — V_k)\}$.
2. **Structural orientation:**
   for $X_i \in \mathbf{X}$ where $(X_i — W) \in \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(X_i \to W)\}\big) \setminus \{(X_i — W)\}$.
   for $V_i \in \mathbf{X} \cup \{W\}, S_j \in \mathbf{S}$ where $(V_i — S_j) \in \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_i \to S_j)\}\big) \setminus \{(V_i — S_j)\}$.
   for $V_i \in \mathbf{V} \setminus \mathbf{Y}, Y_k \in \mathbf{Y}$ where $(V_i — Y_k) \in \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_i \to Y_k)\}\big) \setminus \{(V_i — Y_k)\}$.
3. **Meek rules: repeat** until $\mathbf{M}$ is unchanged in a full pass over (R1)–(R4):
   (R1) **for** $V_i, V_j, V_k \in \mathbf{V}$ where $V_i \notin \mathcal{N}(V_k)$ and $\{(V_i \to V_j), (V_j — V_k)\} \subseteq \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_j \to V_k)\}\big) \setminus \{(V_j — V_k)\}$.
   (R2) **for** $V_i, V_j, V_k \in \mathbf{V}$ where $\{(V_i \to V_j), (V_j \to V_k), (V_i — V_k)\} \subseteq \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_i \to V_k)\}\big) \setminus \{(V_i — V_k)\}$.
   (R3) **for** $V_i, V_j, V_k, V_l \in \mathbf{V}$ where $V_j \notin \mathcal{N}(V_l)$ & $\{(V_i — V_j), (V_j \to V_k), (V_i — V_l), (V_l \to V_k), (V_i — V_k)\} \subseteq \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_i \to V_k)\}\big) \setminus \{(V_i — V_k)\}$.
   (R4) **for** $V_i, V_j, V_k, V_l \in \mathbf{V}$ where $V_i \notin \mathcal{N}(V_k)$ & $\{(V_i \to V_j), (V_j \to V_k), (V_i — V_l), (V_l — V_k), (V_j — V_l)\} \subseteq \mathbf{M}$:
      Update $\mathbf{M} \leftarrow \big(\mathbf{M} \cup \{(V_l \to V_k)\}\big) \setminus \{(V_l — V_k)\}$.
4. Identify all DAGs consistent with $\mathcal{G}_2 = (\mathbf{V}, \mathbf{M})$, store them in $\mathbf{G}$.
5. **Return**.

skeleton in the middle panel. Phase 2 orients each surviving edge using structural ordering, the collider rule, and Meek's rules, producing the directed graph on the right. The only directed path from W to Y_1 runs through S_1 and S_3, while every short-term variable lies on a directed path from W to Y_2. The true graph is uniquely identified here, though this is not always guaranteed.

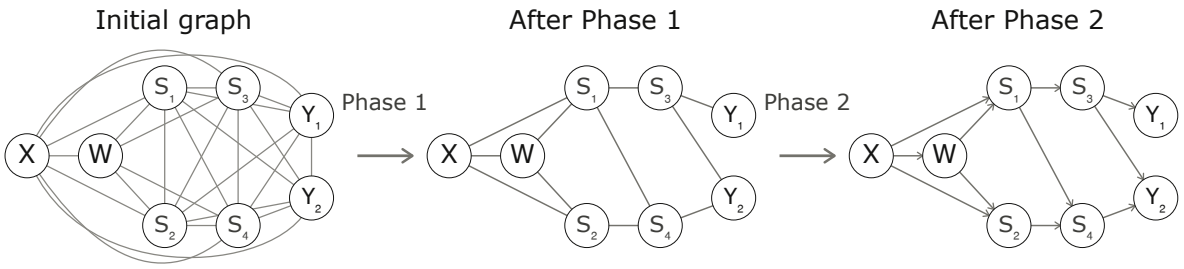

Figure 1 **Illustration of the two phases of COMB-PC. The leftmost panel shows the initial graph $\mathcal{G}_1$. The middle panel shows the skeleton recovered by Phase 1 via conditional independence tests in the experimental and observational samples. The rightmost panel shows $\mathcal{G}_2$, the output of Phase 2.**

To state our main correctness result, we assume that the conditional independence tests used by COMB-PC correctly reflect the conditional independence structure of the distribution induced by the true underlying DAG.

ASSUMPTION 7 **(Correctness of Independence Tests)**. *Let $\mathcal{G}^*$ be the true data-generating graph and let $P_{\mathcal{G}^*}$ denote the joint distribution it induces. Every conditional independence test conducted in Phase 1 correctly reflects the conditional independence structure of $P_{\mathcal{G}^*}$:*

(i) For each query in the experimental sample, $V_i \perp\!\!\!\perp_E V_j \mid \mathbf{C}$ if and only if V_i and V_j are conditionally independent given $\mathbf{C}$ under $P_{\mathcal{G}^*}$.

(ii) For each query in the observational sample, $V_i \perp\!\!\!\perp_O Y \mid \mathbf{C}$ if and only if V_i and Y are conditionally independent given $\mathbf{C}$ under $P_{\mathcal{G}^*}$.

Assumption 7 isolates the causal discovery problem from finite-sample conditional independence testing: it describes the relations that would obtain asymptotically as the sample size grows.

We can now state the main result of this section.

Theorem 1 *Let $\mathcal{G}^*$ be the true underlying DAG and let $\mathbf{G}$ be the set of DAGs returned by the COMB-PC algorithm. Under Assumptions 1-7, we have $\mathcal{G}^* \in \mathbf{G}$ and $\mathcal{G}^*$ is Markov equivalent to every $\mathcal{G} \in \mathbf{G}$.*

All proofs are provided in Appendix G. Under appropriate assumptions, Phase 1 recovers the true skeleton and Phase 2 orients every edge identifiable within the Markov equivalence class of $\mathcal{G}^*$. These guarantees are asymptotic; Appendix F evaluates COMB-PC's finite-sample recovery of the skeleton and edge orientations across a range of graph densities.

COMB-PC makes two contributions. First, to our knowledge it is the first causal discovery algorithm for the two-sample surrogacy setting, where W and $\mathbf{Y}$ are never jointly observed—a configuration that neither single-sample methods such as PC nor existing two-sample causal discovery methods accommodate. Second, by recovering the causal graph from data, it removes the standard surrogacy-literature requirement that a valid surrogate set be specified in advance: the surrogates and the mechanisms through which they mediate the treatment effect are read off the learned graph rather than assumed.

The output of COMB-PC is central to the surrogacy setting for two reasons. First, it identifies the surrogates and the mechanisms through which the treatment affects each long-term outcome, allowing us to estimate the long-term treatment effect without waiting for the long-term outcome to materialize. Second, it encodes the causal pathways linking the treatment, the surrogates, and the long-term outcome: the aggregate effect can be decomposed along its constituent mechanisms, and counterfactual program modifications can be evaluated using the learned graph, as will be discussed in the remainder of the paper.

5. Long-term Treatment Effect Identification with Surrogates

Building on the equivalence class $\mathbf{G}$ returned by COMB-PC, this section identifies the long-term average treatment effect of W on each outcome $Y_k \in \mathbf{Y}$. We first identify two objects off the learned graph: a set of mediating surrogates and a post-treatment adjustment set that closes the remaining noncausal paths to the outcome. We then combine them into a closed-form, non-parametric identification formula for the interventional distribution of Y_k (Proposition 1). Finally, we show that even though the true graph is unknown, $\mathbf{G}$ delivers bounds on the long-term effect (Theorem 2).

5.1. Surrogate and Adjustment Sets from the Graph

We begin by formalizing two graphical objects derived from $\mathcal{G}$: a *surrogate set* that intercepts every directed effect of W on Y_k, and a *surrogate–outcome adjustment set* that closes the remaining noncausal paths between those surrogates and Y_k.

Definition 1 (Valid Surrogate Set). A set $\mathbf{S}_{\mathcal{G},Y_k} \subseteq \mathbf{V} \setminus (\{W\} \cup \mathbf{Y})$ is a *valid surrogate set* for the effect of W on a long-term outcome $Y_k \in \mathbf{Y}$ in a DAG $\mathcal{G}$ if it satisfies:

(i) *Sufficiency.* $\mathbf{S}_{\mathcal{G},Y_k}$ intercepts every directed path from W to Y_k in $\mathcal{G}$.

(ii) *Relevance.* Every $S \in \mathbf{S}_{\mathcal{G},Y_k}$ lies on some directed path from W to Y_k in $\mathcal{G}$.

Sufficiency ensures that conditioning on the surrogates captures all of the treatment's directed effect on Y_k; relevance prevents the set from including variables that contribute nothing to that effect. A valid surrogate set need not be unique: when the graph contains parallel mediating paths, several distinct sets may satisfy both conditions, and Definition 1 admits all of them.

A valid surrogate set is necessary for identification but not sufficient. The relationship between the surrogates in $\mathbf{S}_{\mathcal{G},Y_k}$ and the outcome Y_k may be confounded by other short-term variables, opening *backdoor paths* from the surrogates to Y_k that pre-treatment covariates $\mathbf{X}$ alone cannot block, where a backdoor path is a path whose first edge has an arrowhead into S_j.[5] For each $S_j \in \mathbf{S}_{\mathcal{G},Y_k}$, let $\mathcal{B}_{\mathcal{G}}(S_j, Y_k)$ denote the set of backdoor paths from S_j to Y_k in $\mathcal{G}$. We close these paths by augmenting the surrogate set with a graph-derived adjustment set $\mathbf{Z}_{\mathcal{G},Y_k}$, defined next.

Definition 2 (Valid Surrogate–Outcome Adjustment Set). Let $\mathbf{S}_{\mathcal{G},Y_k} \subseteq \mathbf{S}$ be a valid surrogate set for $Y_k \in \mathbf{Y}$ in graph $\mathcal{G}$ in the sense of Definition 1. A set $\mathbf{Z}_{\mathcal{G},Y_k} \subseteq \mathbf{S} \setminus \mathbf{S}_{\mathcal{G},Y_k}$ is a *valid surrogate–outcome adjustment set* for $(\mathbf{S}_{\mathcal{G},Y_k}, Y_k)$ in $\mathcal{G}$ if the following conditions hold:

(i) For every $S_j \in \mathbf{S}_{\mathcal{G},Y_k}$, no element of $\mathbf{Z}_{\mathcal{G},Y_k}$ is a descendant of S_j in $\mathcal{G}$.

(ii) For every $S_j \in \mathbf{S}_{\mathcal{G},Y_k}$, every path in $\mathcal{B}_{\mathcal{G}}(S_j, Y_k)$ is blocked by $\mathbf{X} \cup \mathbf{Z}_{\mathcal{G},Y_k} \cup \left(\mathbf{S}_{\mathcal{G},Y_k} \setminus \{S_j\}\right)$.

(iii) Every path from W to Y_k in the graph obtained from $\mathcal{G}$ by deleting all incoming arrows into W is blocked by $\mathbf{X} \cup \mathbf{Z}_{\mathcal{G},Y_k} \cup \mathbf{S}_{\mathcal{G},Y_k}$.

[5] Due to space limitations, we formally define blocked paths, open paths, and backdoor paths in Appendix A.2.

Definition 2 is the surrogacy analogue of Pearl's backdoor criterion (Pearl 1995). Condition (i) prevents adjustment on descendants of the selected surrogates. Condition (ii) closes surrogate–outcome backdoor paths using $\mathbf{X}$, the remaining selected surrogates, and $\mathbf{Z}_{\mathcal{G},Y_k}$. Condition (iii) is the adjusted surrogacy condition: after conditioning on $\mathbf{X}$, $\mathbf{S}_{\mathcal{G},Y_k}$, and $\mathbf{Z}_{\mathcal{G},Y_k}$, treatment carries no remaining open path to Y_k in the post-intervention graph.

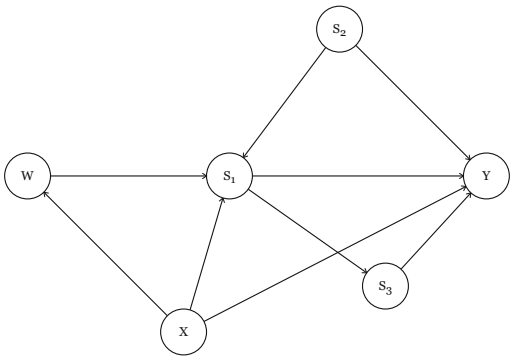

Figure 2 A DAG where $\{S_1\}$ is a valid surrogate set for Y in the sense of Definition 1, but there is a noncausal path $S_1 \leftarrow S_2 \rightarrow Y$ that the pre-treatment covariates X alone cannot block. Any valid surrogate–outcome adjustment set $\mathbf{Z}_{\mathcal{G},Y}$ for $(\{S_1\},Y)$ must therefore contain S_2.

5.2. Identification under a Known Graph

We can now state the main identification result of this section. Following the surrogate-index literature (Athey et al. 2025), Assumption 8 imposes the standard support and comparability conditions needed to combine the experimental and observational samples. Together with a valid surrogate set $\mathbf{S}_{\mathcal{G},Y_k}$ and a valid surrogate–outcome adjustment set $\mathbf{Z}_{\mathcal{G},Y_k}$, these identify the interventional distribution $P(Y_k \mid D = E, \mathrm{do}(W = w))$ from the experimental and observational samples, with no joint observation of W and Y_k required.[6] Here $\mathrm{do}(W = w)$ denotes the intervention that sets the treatment to w, the standard notation in graphical causal models (Pearl 2000); in the potential-outcomes framework it corresponds to the potential outcome $Y_k(w)$, so that $P(Y_k \mid D = E, \mathrm{do}(W = w))$ is the distribution of $Y_k(w)$ in the experimental population. We use the *do operator* throughout, as is common in graphical causal models, but it can be interchanged with the potential-outcomes notation $Y_k(w)$. See Appendix A.2 for details on the do operator.

Assumption 8 **(Comparability of Samples)**. *For a graph $\mathcal{G}$ and an outcome $Y_k \in \mathbf{Y}$, let $\mathbf{S}_{\mathcal{G},Y_k} \subseteq \mathbf{S}$ be a valid surrogate set for Y_k in $\mathcal{G}$ (Definition 1), and let $\mathbf{Z}_{\mathcal{G},Y_k} \subseteq \mathbf{S} \setminus \mathbf{S}_{\mathcal{G},Y_k}$ be a valid surrogate–outcome adjustment set for $(\mathbf{S}_{\mathcal{G},Y_k}, Y_k)$ in $\mathcal{G}$ (Definition 2). For every $w \in \{0,1\}$ and every $(\mathbf{x},\mathbf{z},\mathbf{s})$ such that $P(\mathbf{X} = \mathbf{x}, \mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s} \mid D = E, W = w) > 0$, the following two conditions hold:*

[6] We state the result for discrete variables for notational clarity. For continuous variables, the sums are replaced by integrals and the conditional probability $P(Y_k \mid \mathbf{X}, \mathbf{Z}_{\mathcal{G},Y_k}, \mathbf{S}_{\mathcal{G},Y_k})$ by a conditional expectation or density.

(i) The conditional distribution of Y_k given $\mathbf{X}$, $\mathbf{Z}_{\mathcal{G},Y_k}$, and $\mathbf{S}_{\mathcal{G},Y_k}$ is the same in the experimental and observational samples:

$$P(Y_k = y \mid D = E, \mathbf{X} = \mathbf{x}, \mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s}) = P(Y_k = y \mid D = O, \mathbf{X} = \mathbf{x}, \mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s}).$$

(ii) The observational sample has support at $(\mathbf{x}, \mathbf{z}, \mathbf{s})$: $P(\mathbf{X} = \mathbf{x}, \mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s} \mid D = O) > 0$.

PROPOSITION 1 **(Long-Term Treatment Effect Identification)**. *For a graph $\mathcal{G}$ and an outcome $Y_k \in \mathbf{Y}$, let $\mathbf{S}_{\mathcal{G},Y_k}$ and $\mathbf{Z}_{\mathcal{G},Y_k}$ be a valid surrogate set and a valid surrogate–outcome adjustment set for $(\mathbf{S}_{\mathcal{G},Y_k}, Y_k)$ in $\mathcal{G}$, respectively. Under Assumptions 1, 2, 3, 4, 6, and 8, the interventional distribution of Y_k in the experimental population is identified, for each $w \in \{0, 1\}$, by*

$$P(Y_k = y \mid D = E, \mathrm{do}(W = w)) = \sum_{\mathbf{x}} \sum_{\substack{\mathbf{z} \in \mathcal{Z}_{\mathcal{G},Y_k} \\ \mathbf{s} \in \mathcal{S}_{\mathcal{G},Y_k}}} P(Y_k = y \mid D = O, \mathbf{X} = \mathbf{x}, \mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s}) \times P(\mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s} \mid D = E, W = w, \mathbf{X} = \mathbf{x}) \times P(\mathbf{X} = \mathbf{x} \mid D = E), \tag{1}$$

where $\mathcal{S}_{\mathcal{G},Y_k}$ and $\mathcal{Z}_{\mathcal{G},Y_k}$ denote the supports of $\mathbf{S}_{\mathcal{G},Y_k}$ and $\mathbf{Z}_{\mathcal{G},Y_k}$, respectively.

Proposition 1 identifies the long-term effect of W on each Y_k by combining an outcome model from the observational sample, a joint surrogate–adjustment model from the experimental sample, and the marginal distribution of pre-treatment covariates, with no joint observation of W and Y_k required. This result extends the surrogate-index framework of Athey et al. (2019) in two respects. First, we augment the surrogate set with a graph-derived adjustment set $\mathbf{Z}_{\mathcal{G},Y_k}$ that blocks noncausal paths through post-treatment confounders, which no pre-treatment adjustment alone can block. Second, surrogacy is derived rather than assumed: Athey et al. (2019) take a surrogate set $\mathbf{S}$ together with the validity as primitives, whereas we split this set into a mediating surrogate set $\mathbf{S}_{\mathcal{G},Y_k}$ (Definition 1) and a surrogate–outcome adjustment set $\mathbf{Z}_{\mathcal{G},Y_k}$ (Definition 2), both read off the graph; the required conditional independence then follows from the graphical structure.

5.3. Identification under Graph Uncertainty

With Proposition 1 established, we next turn to recovering the average treatment effect using the graphs $\mathbf{G}$ returned by COMB-PC. Let $\tau_{\mathcal{G},Y_k}$ denote the average treatment effect of W on outcome $Y_k \in \mathbf{Y}$ in graph $\mathcal{G}$, defined as $\tau_{\mathcal{G},Y_k} = \sum_{y_k \in \mathcal{Y}_k} y_k \cdot \big(P(Y_k = y_k \mid D = E, \mathrm{do}(W = 1)) - P(Y_k = y_k \mid D = E, \mathrm{do}(W = 0))\big)$. If we had access to the true graph $\mathcal{G}^*$, $\tau_{\mathcal{G}^*,Y_k}$ would be directly identified by Proposition 1. Since $\mathcal{G}^*$ is unknown, we instead work with the equivalence class $\mathbf{G}$ returned by COMB-PC. Our next result bounds $\tau_{\mathcal{G}^*,Y_k}$ by the minimum and maximum of the candidate effects $\{\tau_{\mathcal{G},Y_k}\}_{\mathcal{G} \in \mathbf{G}}$.

Theorem 2 (Bounds on the True Average Treatment Effect) *Suppose Assumptions 1, 2, 3, 4, 5, 6, 7, and 8 hold. Let $\mathbf{G}$ be the set of graphs returned by the COMB-PC algorithm. For each*

$\mathcal{G} \in \mathbf{G}$ and each $Y_k \in \mathbf{Y}$, let $\tau_{\mathcal{G},Y_k}$ denote the candidate average treatment effect obtained by applying Proposition 1 to $\mathcal{G}$ with a valid surrogate set $\mathbf{S}_{\mathcal{G},Y_k}$ and a valid surrogate–outcome adjustment set $\mathbf{Z}_{\mathcal{G},Y_k}$. Then, for every $Y_k \in \mathbf{Y}$, we have $\min_{\mathcal{G}\in\mathbf{G}} \tau_{\mathcal{G},Y_k} \le \tau_{\mathcal{G}^*,Y_k} \le \max_{\mathcal{G}\in\mathbf{G}} \tau_{\mathcal{G},Y_k}$.

Theorem 2 bounds the true long-term effect $\tau_{\mathcal{G}^*,Y_k}$ using the equivalence class $\mathbf{G}$, enabling assessment of an intervention's long-term impact without waiting for Y_k to materialize. We next turn to improving that impact through program redesign, which is the focus of the rest of the paper.

6. Graph-Constrained Counterfactual Evaluation

A program's design choices are made years before its long-term outcomes can be measured, so in designing the next iteration the policymaker must decide which features to strengthen before knowing which ones carry the durable effects. Suppose the policymaker is considering a menu of feasible program modifications, each strengthening the treatment's effect on a particular short-term variable. Strengthening the channels behind the largest short-term gains is a natural starting point, but these are not always the channels that drive long-term effects: a short-term response can be large early on and still fade by the long-term horizon. The goal of this section is to identify the long-term impact of each modification and help select the one that yields the largest improvement in long-term outcomes.

The surrogate literature has so far been concerned with estimating the long-term effect of a fixed treatment. We instead compare alternative modifications of that treatment, which requires tracing how a change in a targeted short-term variable propagates to the long-term outcome. The graphical approach makes this explicit: the graph specifies the downstream relationships and identifies the paths through which a change in a targeted variable reaches the outcome. To exploit this, in Section 6.1 we layer a graph-constrained linear structural equation model on each candidate graph, assigning coefficients to the directed edges and yielding path-specific effects from short-term variables to long-term outcomes. We represent each modification as a modular change to the targeted variables' structural equations and, because COMB-PC returns an equivalence class of graphs, evaluate it across them, selecting the alternative with the highest worst-case improvement.

6.1. Graph-Constrained Decomposition of Treatment Mechanisms

The mechanism analysis proceeds in two steps. For a graph $\mathcal{G}$, we adopt a linear structural equations model (SEM) to quantify the strength of every directed edge and derive the model-implied effect of a program modification on the long-term outcomes. The SEM lets us trace how a change in the program's effect on a short-term variable S_j reaches a long-term outcome Y_k in $\mathcal{G}$. The next subsection handles graph uncertainty by comparing these graph-specific effects across all $\mathcal{G} \in \mathbf{G}$.

In a graphical causal model, the graph represents the data-generating process: each variable is generated by a structural equation whose arguments are its parents in the graph (Pearl 2000). For

a candidate graph $\mathcal{G} \in \mathbf{G}$ and each post-treatment variable $V_j \in \mathbf{S} \cup \mathbf{Y}$, the graph-constrained linear specification is

$$V_j = \alpha_j^{\mathcal{G}} + \sum_{V_i \in \mathrm{Pa}_{\mathcal{G}}(V_j)} \beta_{V_i,V_j}^{\mathcal{G}} V_i + \varepsilon_j^{\mathcal{G}}, \tag{2}$$

where $\alpha_j^{\mathcal{G}} \in \mathbb{R}$ is the graph-specific intercept and $\beta_{V_i,V_j}^{\mathcal{G}} \in \mathbb{R}$ is the graph-specific coefficient on the edge $V_i \to V_j$. The error term satisfies the conditional mean-zero condition $\mathbb{E}\left[\varepsilon_j^{\mathcal{G}} \mid \mathrm{Pa}_{\mathcal{G}}(V_j)\right] = 0$.

Under this specification, the graph $\mathcal{G}$ determines which variables enter each equation, and the linear SEM assigns a scalar coefficient to each directed edge. The intercepts and coefficients carry the superscript $\mathcal{G}$ because the regressors entering each equation are determined by the parent set $\mathrm{Pa}_{\mathcal{G}}(V_j)$, which depends on the graph. When $\mathcal{G} = \mathcal{G}^*$ and the true structural equations are linear in their parents, the coefficients recover the true direct causal effects (Pearl 2000). The equations are estimated from the sample that jointly observes each node and its parents. Equations for short-term variables are estimated from the experimental sample, where $\mathbf{X}$, W, and $\mathbf{S}$ are jointly observed. Equations for long-term outcomes are estimated from the observational sample, where $\mathbf{X}$, $\mathbf{S}$, and $\mathbf{Y}$ are jointly observed. Assumption 1 justifies this split by requiring the structural equations for variables in $\mathbf{V} \setminus \{W\}$ to be invariant across the two samples.

Evaluating the effect of a program modification on a long-term outcome requires two components: a downstream causal map from short-term variables to the outcome, and a characterization of how the modification alters the first-stage coefficients. We develop each in turn. For each pair (S_j, Y_k), let $\Gamma_{S_j,Y_k,\mathcal{G}}$ denote the downstream SEM-implied total causal effect of S_j on Y_k in $\mathcal{G}$. This quantity measures the change in Y_k induced by a one-unit increase in S_j. Let $\mathcal{P}_{\mathcal{G}}(S_j, Y_k)$ be the set of directed paths from S_j to Y_k in $\mathcal{G}$. For a path p, write $(U \to V) \in p$ to denote a directed edge that lies on that path. The contribution of path p is the product of the structural coefficients on its edges. Therefore, $\Gamma_{S_j,Y_k,\mathcal{G}} = \sum_{p \in \mathcal{P}_{\mathcal{G}}(S_j,Y_k)} \prod_{(U \to V) \in p} \beta_{U,V}^{\mathcal{G}}$. If there is no directed path from S_j to Y_k in $\mathcal{G}$, we set $\Gamma_{S_j,Y_k,\mathcal{G}} = 0$. Thus, $\Gamma_{S_j,Y_k,\mathcal{G}}$ aggregates both the direct effect of S_j on Y_k, when $S_j \to Y_k$ is present, and all indirect effects transmitted through directed paths from S_j to Y_k. This construction gives, for each candidate graph $\mathcal{G}$, the downstream causal effect from any short-term variable S_j to any long-term outcome Y_k.

We now use the SEM to evaluate program modifications. Let $\mathcal{R} = \{r^1, \ldots, r^H\}$ denote the menu of feasible program modifications, where each alternative $r^h = (\mathcal{A}^h, \boldsymbol{\delta}^h)$ specifies a set $\mathcal{A}^h \subseteq \{1, \ldots, m\}$ of targeted short-term variables and quantities $\delta_j^h \geq 0$ for $j \in \mathcal{A}^h$. Within the SEM, the first-stage effect of the program on S_j is the coefficient $\beta_{W,S_j}^{\mathcal{G}}$ on W in S_j's structural equation, and each modification r^h adjusts this coefficient by δ_j^h for the variables in $\mathcal{A}^h$. Under r^h, the structural equation for each targeted S_j becomes

$$S_j = \alpha_j^{\mathcal{G}} + \sum_{V_i \in \mathrm{Pa}_{\mathcal{G}}(S_j)} \beta_{V_i,S_j}^{\mathcal{G}} V_i + \delta_j^h W + \varepsilon_j^{\mathcal{G}}, \tag{3}$$

with all other structural equations unchanged. When $W \to S_j$ is in $\mathcal{G}$, the effective first-stage coefficient on W becomes $\beta^{\mathcal{G}}_{W,S_j} + \delta^h_j$, strengthening the existing first-stage effect. When $W \to S_j$ is not in $\mathcal{G}$ (so $\beta^{\mathcal{G}}_{W,S_j} = 0$), it becomes δ^h_j, introducing a new first-stage effect on S_j. With this, we can characterize the long-term consequence of a program modification under graph $\mathcal{G}$. Holding the downstream SEM fixed, the incremental SEM-implied effect of policy alternative r^h on outcome Y_k is therefore $\Delta_{Y_k,\mathcal{G}}(r^h) = \sum_{j \in \mathcal{A}^h} \delta^h_j \Gamma_{S_j,Y_k,\mathcal{G}}$.

When program modifications affect several long-term outcomes, the policymaker can either combine these outcomes into a single scalar via weights or evaluate each outcome separately. We adopt the former for the remainder of this section. The weights $w_k \in \mathbb{R}$ represent the policymaker's relative priorities across outcomes and are typically elicited directly from the policymaker. Under a fixed graph $\mathcal{G}$, the weighted gain from modification r^h is therefore $\Delta_{\mathcal{G}}(r^h) = \sum_{k=1}^{K} w_k \, \Delta_{Y_k,\mathcal{G}}(r^h)$.

6.2. Robust Counterfactual Program Selection

Now that we have characterized the gain from a program modification under a fixed graph $\mathcal{G}$, we must account for graph uncertainty. Because COMB-PC returns an equivalence class $\mathbf{G}$ rather than a single graph,the same policy alternative r^h must be evaluated over all candidate mechanisms. The alternative itself does not change across graphs: it targets the same short-term variables and specifies the same first-stage changes $\boldsymbol{\delta}^h$. What may change is how those targeted variables propagate to the long-term outcomes, because the downstream effects $\Gamma_{S_j,Y_k,\mathcal{G}}$ depend on the directed paths in $\mathcal{G}$.

To formalize this, for a fixed program modification r^h and outcome Y_k, evaluating the modification over the equivalence class gives the graph-indexed collection $\left\{\Delta_{Y_k,\mathcal{G}}(r^h) : \mathcal{G} \in \mathbf{G}\right\}$. These gains are identical only when the targeted short-term variables have the same downstream effects across all graphs in $\mathbf{G}$. Otherwise, different orientations of the undirected edges across $\mathbf{G}$ can produce different $\Gamma_{S_j,Y_k,\mathcal{G}}$ and therefore can yield different predicted gains from the same program modification. The following proposition formalizes this graph dependence for a single outcome.

Proposition 2 **(Graph dependence of the policy gain)**. *There exist data-generating processes satisfying Assumptions 1 - 8 for which COMB-PC returns a set of graphs* $\mathbf{G}$ *containing two graphs* $\mathcal{G}$ *and* $\mathcal{G}'$ *such that, under the graph-constrained linear SEM in* (2)*, there exist a program modification* $r^h = (\mathcal{A}^h, \boldsymbol{\delta}^h)$ *and an outcome* $Y_k \in \mathbf{Y}$ *for which* $\Delta_{Y_k,\mathcal{G}}(r^h) \neq \Delta_{Y_k,\mathcal{G}'}(r^h)$.

Proposition 2 reflects a limit of observational identification: the Markov equivalence class encodes the conditional-independence structure of the data but does not pin down the structural quantities on which counterfactual policy effects depend. Two graphs that are indistinguishable from the observed distribution can still disagree on $\Delta_{Y_k,\mathcal{G}}(r^h)$. Appendix E illustrates this with an example.

To account for this graph uncertainty, the policymaker can select the graph from $\mathbf{G}$ that best aligns with their understanding of the underlying mechanism, place a prior over $\mathbf{G}$ and rank modifications by their prior-weighted expected gain, or take the worst-case gain across $\mathbf{G}$ and select

the modification that maximizes it. We adopt the worst-case rule for the remainder of this section. For a program like YOP, where modifications are expensive to deploy and difficult to undo, this guards against committing to an alternative that looks promising under one candidate mechanism but underperforms under another that fits the data equally well. This is the same minimax logic emphasized in policy-learning work such as Kitagawa and Tetenov (2018), where treatment rules are evaluated by their worst-case welfare regret rather than by performance under a single data-generating process. Formally, the policymaker selects $h^\star \in \arg\max_{h\in\{0,1,\ldots,H\}} \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h)$. The status quo r^0 — defined by $\mathcal{A}^0 = \emptyset$, with $\Delta_{\mathcal{G}}(r^0) = 0$ for every $\mathcal{G} \in \mathbf{G}$ — is retained in the menu so that the selected modification has nonnegative worst-case gain by construction.

The following theorem formalizes this worst-case guarantee for the true graph $\mathcal{G}^*$.

Theorem 3 (Worst-case guarantee) *Suppose Assumptions 1 - 8 hold and the post-treatment structural equations are linear in their parents. Then for every program modification* r^h, $\min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h) \le \Delta_{\mathcal{G}^*}(r^h) \le \max_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h)$, *and for any* $h^\star \in \arg\max_{h\in\{0,1,\ldots,H\}} \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h)$, $\Delta_{\mathcal{G}^*}(r^{h^\star}) \ge \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^{h^\star}) \ge 0$.

Theorem 3 formalizes the implications of the decision rule: the worst-case gain is a lower bound on the impact a policymaker can expect under the true graph. The next section applies this framework — a graph-based mechanism decomposition and counterfactual evaluation– to the Uganda Youth Opportunities Program, where short-term data are used to recover the nine-year evaluation pattern.

7. From Surrogate Validation to Program Redesign: The Uganda Youth Opportunities Program

The Uganda Youth Opportunities Program (YOP) is a near-ideal setting in which to evaluate our proposed methodology. Implemented by the Government of Uganda in 2008 with World Bank support (Blattman et al. 2020), YOP gave one-time cash grants of approximately $400 per person (i.e., a sum roughly equal to recipients' annual income) to groups of underemployed young adults proposing to set themselves up as independent craftspeople in skilled trades such as carpentry, tailoring, and metalwork. Among eligible applicants, grants were randomly assigned: those selected received the cash transfer, while the rest served as controls. Surveys were conducted to measure livelihood outcomes two, four, and nine years after the grants were disbursed. Randomized panels of this length are rare in development field experiments, and the YOP panel is methodologically valuable for a reason that goes beyond data availability: because both short-term and long-term results are observed, we can test whether our framework would have raised an early warning that the program's initial gains might not persist.

At four years, YOP looked like a clear success. Blattman et al. (2020) report that, relative to the randomized control group, treated individuals reported 17% more work hours, 38% higher earnings,

and 11% higher consumption, alongside large increases in business capital and entry into skilled trades. By year 9, however, most of these gains had dissipated. Control-group participants eventually found wage work and accumulated capital on their own, and the two groups converged on earnings, consumption, and total employment. Lasting effects appeared only on a narrower set of outcomes such as durable assets and sustained occupation in a skilled trade. This raises the question our framework is designed to answer: could we have learned from the year-4 data alone what the year-9 follow-up eventually revealed? And can the short-term responses be decomposed into those that carry the program's durable effects and those that carry only its transient ones?

Specifically, we use the YOP data to address three questions. *First*, can the proposed framework recover the year-nine treatment effects from the year-two and year-four survey results alone, without waiting nine years? This is the basic identification test: using only the four-year data, the framework should reproduce the year-nine pattern of effects — both where lasting gains emerged and where they did not. *Second*, can we come to this conclusion based solely on two year data? If so, COMB-PC can allow policymakers to make informed decisions sooner. *Third*, how can we leverage the learned graph, which encodes the causal mechanisms, to redesign more effective future programs? Because the framework returns a graph in addition to a point estimate, it supports a counterfactual inference that ranks modifications by their predicted year-9 gain, identifying which short-term mechanisms are the highest-leverage targets when scaling similar programs.

7.1. Empirical Setup

The full sample contains 2,677 individuals in 535 groups across 14 northern Ugandan districts; 265 groups were randomly assigned to receive a grant.[7] Participants were surveyed at baseline and again at roughly two, four, and nine years post-disbursement. The pre-treatment covariate set $\mathbf{X}$ comprises 37 baseline variables. These include standard demographics (age, gender, urban residence), human-capital measures (education, literacy, numeracy, prior vocational training), economic status (wealth, savings, baseline earnings, credit access, work hours by sector), and group-level features of the grant application (group size and composition, prior group existence, leadership, district stratum). The candidate surrogate set $\mathbf{S}$ contains 17 intermediate outcomes from the year-two and year-four surveys. At year two: enrollment in vocational training, return to school, hours of training received, value of business assets, weekly hours worked, having a non-agricultural main occupation, working in any skilled trade, and self-reported monthly earnings. At year four: monthly net earnings, nondurable consumption, value of durable assets, weekly hours worked, having a non-agricultural main occupation, working in any skilled trade, working more than thirty hours per week in a skilled

[7] We use the replication data from Blattman et al. (2020), publicly archived at the Harvard Dataverse (Innovations for Poverty Action collection): `https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/VONOHA`.

trade, weekly hours spent on household chores, and hourly earnings. The long-term outcome set $\mathbf{Y}$ comprises the year-9 outcomes: monthly net earnings, nondurable consumption, durable assets, average weekly employment hours, employment outside of agriculture, any skilled trade work, full-time skilled-trade work (30+ hours per week), weekly hours spent on chores, and earnings per hour. [8] Throughout this section, W denotes randomized assignment to the grant.

The COMB-PC framework requires two samples: an *experimental sample* containing $(W, \mathbf{X}, \mathbf{S})$, which includes both the 2- and 4-year surrogates but withholds the long-term outcomes $\mathbf{Y}$, and an *observational sample* containing $(\mathbf{X}, \mathbf{S}, \mathbf{Y})$ but withholding W. We construct these by partitioning the 14 districts in the YOP sample into two disjoint groups of seven, with the experimental sample drawn from Arua, Kumi, Moroto, Moyo, Nakapiripirit, Nebbi, and Yumbe ($n = 1{,}017$) and the observational sample from Adjumani, Apac, Kaberamaido, Kotido, Lira, Pallisa, and Soroti ($n = 1{,}660$). Note that the geographic split is deliberately more demanding than an individual-level random split: the long-term outcome model must transport across geographic units rather than reuse units that appear in both samples. In fact, a similar experimental–observational separation underlies the surrogate-index analyses in seminal papers by Athey et al. (2019) and Imbens et al. (2022).

7.2. Can Short-Term Surrogates Recover the Year-Nine Benchmark?

First, we ask whether our framework, using the year-two and year-four outcomes, can recover the year-nine intent-to-treat (ITT) estimates of Blattman et al. (2020). We first apply the COMB-PC algorithm to the experimental and observational samples, testing conditional independence at significance level 0.05 using partial correlations, to learn the causal graph over the pre-treatment covariates, candidate surrogates, and long-term outcomes. COMB-PC returns a unique graph for the YOP data, so the bounds of Theorem 2 collapse to a point. We identify valid surrogate–control pairs from this graph via Definition 2 and estimate the year-nine treatment effect via Proposition 1. Standard errors for treatment effect estimates are computed from 500 bootstrap resamples.

Table 1 compares the nine-year treatment effects obtained from our estimation procedure to the benchmark. We also compare to a full-surrogate setting, where we use Proposition 1 with all candidate surrogates and all pre-treatment covariates without graphical selection. The first column reports the benchmark treatment effect estimates with their significance levels, the second column reports the same values estimated using our graph-selected surrogate and covariate set, and the third column reports these values estimated with the full set of surrogates and covariates.

In the field experiment, the short- and medium-term gains from YOP largely fade by year nine for earnings, consumption, and employment, while effects persist for durable assets and skilled-trade attachment (Blattman et al. 2020). Our COMB-PC-based estimator recovers this pattern. For the

[8] See Blattman et al. (2020) for full set of covariates and variable definitions.

Table 1 Recovering the Nine-Year ITT Estimates from a Geographic District Split

Year-9 outcome	Field Experiment ITT	Graph-selected estimator	Full-surrogate estimator
Economic outcomes			
Monthly net earnings	4.170	0.000	6.380
Nondurable consumption	2.730	0.000	15.030*
Durable assets	0.116**	0.053**	0.115***
Employment and occupation			
Avg. employment hrs/wk	0.513	0.000	2.140
Main occupation non-agri.	0.029	0.000	0.050**
Avg. chores hrs/wk	−0.117	0.000	0.114
Avg. earnings/hr	0.009	0.000	0.011
Skilled-trade attachment			
Engaged in any skilled trade	0.197***	0.117***	0.137***
Works 30+ hrs/wk skilled	0.029***	0.025**	0.029**

Notes. Experimental sample: (n = 1,017), drawn from seven districts containing $(W, \mathbf{X}, \mathbf{S})$. Observational sample:($n$ = 1,660), drawn from the remaining seven districts containing $(\mathbf{X}, \mathbf{S}, \mathbf{Y})$. No individual appears in both samples. A value of 0.000 in the graph-selected column indicates that COMB-PC identifies no directed causal path from treatment to the corresponding long-term outcome; the estimate is therefore a structural zero implied by the learned graph rather than a rounded numerical estimate. Significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

outcomes with statistically significant field experiment ITT estimates — durable assets, engagement in skilled trade, and full-time skilled-trade work — the graph-selected estimator returns statistically significant effects, though smaller in magnitude, a consequence of pruning weak channels during graph learning. For the remaining outcomes, where ITT estimates are statistically insignificant, COMB-PC does not identify a directed path from the treatment to the long-term outcome, and the estimator returns a structural zero, reflecting the absence of a treatment-to-outcome channel in the learned causal graph. The full-surrogate estimator produces a different, less reliable pattern. It implies statistically significant effects on nondurable consumption and employment outside of agriculture, though neither is significant in the randomized benchmark. This discrepancy illustrates the risk of treating all post-treatment variables as valid surrogates without imposing causal structure. Post-treatment variables may predict long-term outcomes, but need not mediate the treatment effect. An estimator that conditions on all candidate variables can therefore introduce bias by opening noncausal paths or adjusting for variables that do not belong in the surrogate set. The comparison shows that the learned causal graph plays a substantive role in separating valid surrogate channels from merely predictive post-treatment variables.

7.3. Earlier Decisions: Two-Year Surrogates Are Sufficient

We want to understand whether our procedure would enable policymakers to make consequential decisions on whether to continue, scale, redesign, or terminate interventions even earlier. To examine this directly, we restrict the candidate surrogate set **S** to the eight variables measured at year 2 and remove all year-4 variables from consideration. The aim is to understand whether we would be able to recover the same results on long-term efficacy even more quickly.

The graph learning and estimation procedures remain the same as in Section 7.2, except the available surrogate outcomes are restricted to the year-2 surrogates alone. Table 2 reports the results of this analysis. Once again, the first column reports the benchmark treatment effect which remains the same, while the second column reports the results from using this more restricted graph-selected surrogate and covariate set, and the third column reports the results from using all 2-year surrogates and all covariates.

Table 2 Recovering the Nine-Year ITT Estimates from a Geographic District Split Using Only Two-Year Surrogates

Year-9 outcome	Field Experiment ITT	Graph-selected estimator	Full-surrogate estimator
Economic outcomes			
Monthly net earnings	4.172	0.000	4.727
Nondurable consumption	2.726	0.000	8.624
Durable assets	0.116**	0.052**	0.070**
Employment and occupation			
Avg. employment hrs/wk	0.513	0.000	1.261
Main occupation non-agri.	0.029	0.000	0.049**
Avg. chores hrs/wk	−0.117	0.000	0.089
Avg. earnings/hr	0.009	0.000	−0.032
Skilled-trade attachment			
Engaged in any skilled trade	0.197***	0.095***	0.102***
Works 30+ hrs/wk skilled	0.029***	0.015*	0.015*

Notes. Experimental sample: ($n = 1{,}017$), drawn from seven districts containing $(W, \mathbf{X}, \mathbf{S}_{t=2})$, which only includes 2-year surrogate outcomes. Observational sample:($n = 1{,}660$), drawn from seven districts containing $(\mathbf{X}, \mathbf{S}_{t=2}, \mathbf{Y})$, again only including 2-year surrogate outcomes. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Our COMB-PC-based estimator once again recovers all three persistent effects identified in the randomized benchmark of Blattman et al. (2020) with the correct sign and statistical significance. For every outcome where the field experiment ITT is statistically insignificant, COMB-PC yet again identifies no directed path from the treatment to the long-term outcome, and the estimator returns a structural zero. Our framework therefore reproduces the conclusions of the nine-year evaluation using only data observed within two years of grant disbursement. In contrast, using all-candidate surrogates and covariates identifies one significant effect that does not appear in the benchmark.

Table 3 Most Important Directed Pathways from Treatment to Year-9 Outcomes Using Year-2 Surrogates

Directed pathway	Pathway contribution
Panel A. Durable assets at year 9	
$W \to$ vocational training$_{t=2} \to$ outcome	0.020
Panel B. Engagement in any skilled trade at year 9	
$W \to$ vocational training$_{t=2} \to$ outcome	0.045
$W \to$ vocational training$_{t=2} \to$ skilled trade$_{t=2} \to$ outcome	0.037
$W \to$ vocational training$_{t=2} \to$ training hours$_{t=2} \to$ skilled trade$_{t=2} \to$ outcome	0.012
$W \to$ training hours$_{t=2} \to$ skilled trade$_{t=2} \to$ outcome	0.005
Panel C. Full time employment in a skilled trade at year 9	
$W \to$ vocational training$_{t=2} \to$ skilled-trade work$_{t=2} \to$ outcome	0.012
$W \to$ vocational training$_{t=2} \to$ training hours$_{t=2} \to$ skilled-trade work$_{t=2} \to$ outcome	0.004
$W \to$ training hours$_{t=2} \to$ skilled-trade work$_{t=2} \to$ outcome	0.002

Notes. Each row reports the product of structural coefficients along a single directed path from treatment assignment, W, to the year-9 outcome named in the corresponding panel header.

Notably, point estimates are modestly attenuated when year-4 information is withheld, a pattern consistent with the additional predictive content carried by intermediate outcomes closer to year 9, but the pattern of significance is unchanged. The two-year surrogate structure is therefore sufficient for the identification task: restricting attention to the earliest measurements reduces precision somewhat but does not change the substantive conclusions of the nine-year-surrogate analysis.

7.4. What Drives the Long-Term Effects?

After validating the learned graph against the nine-year benchmark, we use it to examine the channels through which the persistent effects arise. Specifically, we estimate graph-constrained linear structural equation models as outlined in Section 6. We then look at the directed paths from W to each persistent outcome. To focus on the earliest available signals, we restrict the candidate surrogate set to year-2 variables.

Table 3 reports the mechanisms by which the treatment affects the outcomes via surrogates observable at the two-year mark as well as the size of their impact. For durable assets at year 9, the only mechanism is through vocational training at year 2. For engagement in a skilled trade at year 9, among the four mechanisms, three of them are mediated by the first-stage effect of treatment on vocational training at year 2. The remaining pathway runs through training hours at year 2, with a much smaller share of the total effect. We see a similar pattern for the long-term outcome of full-time skilled trade work. There are three mechanisms, two of which are mediated by the first-stage effect of treatment on vocational training at year 2, while the remaining mechanism is mediated by the first-stage effect of treatment on training hours at year 2. The decomposition reveals a consistent pattern

across the persistent outcomes. These mechanisms show that lasting impacts are mediated through vocational training enrollment at year 2. Other year-2 variables, such as training hours and early skilled-trade participation, operate mainly as downstream amplifiers that transmit and reinforce the effect from the initial training channel. These results for all outcomes show that vocational training and, to a lesser extent, training hours are essential to achieving lasting results. The graph also shows which short-term variables do *not* carry the program's lasting effect. The responses that made YOP look successful early on—more business assets, higher earnings, more hours worked, and a shift out of agriculture—are not on any causal path from the grant to the year-9 outcomes that persist.

7.5. Where Should a Redesigned Program Push Harder?

The decomposition in Section 7.4 identifies the mechanisms through which the treatment affects long-term outcomes. This naturally motivates the program-design counterfactual analysis that we illustrate in this subsection. Specifically, we ask how much the year-9 treatment effect would increase if the program had a 10% larger first-stage effect on a specific year-2 mechanism. This counterfactual has a direct operational interpretation. For example, strengthening the first-stage effect from treatment assignment, W, to vocational training enrollment corresponds to a redesigned YOP variant that more effectively encourages grant recipients to engage in vocational training. Such a redesign could involve enrollment vouchers, transportation support, or facilitating connections to local training providers. Strengthening the first-stage effect of W on training hours, by contrast, corresponds to a variant that increases training intensity without necessarily increasing training uptake.

Table 4 compares the impact of improving first-stage effects on vocational training and training hours for each persistent year-9 outcome. The table reports the percentage increase in the SEM-implied year-9 treatment effect when the relevant first-stage relationship is strengthened by 10%. Because vocational training enrollment is the only mechanism mediating the effect of treatment on durable assets, a 10% stronger first-stage effect on it translates into a 10% increase in the year-9 effect on durable assets. For the other two long-term outcomes, engagement in skilled trade and full time employment, vocational training mediates the majority of the effect, so with a 10% increase in first-stage effect on vocational training we see improvements of 9.8% and 9.6%, respectively. In contrast, training hours mediates less of the effect and therefore strengthening the program's effect on this variable produces only small additional gains, increasing the year-9 effects on skilled-trade engagement and full-time skilled-trade work by 0.2% and 0.4%, respectively.

This asymmetry is informative for policy design. The results suggest that the most impactful mechanism through which the program generates persistent effects is enrollment in vocational training. When we increase the impact of treatment on vocational training, that in itself will already induce a downstream effect on training hours. Once that enrollment decision has been made,

Table 4 Counterfactual Gains from Redesigning the YOP Program

Year-9 outcome	Increase in year-9 effect	
	10% stronger effect on vocational training$_{t=2}$	**10% stronger effect on training hours$_{t=2}$**
Durable assets	**10.0**%	*n/a*
Engaged in any skilled trade	**9.8**%	0.2%
Full time employment in skilled trade	**9.6**%	0.4%

Notes. Each cell reports the percentage increase in the SEM-implied year-9 treatment effect when the program's first-stage effect on the indicated year-2 mechanism is increased by 10%. The entry *n/a* indicates that the strengthened mechanism does not lie on a directed path to the corresponding outcome in the learned graph. Bold entries denote the larger counterfactual gain for each outcome.

marginal increases in training hours appear to add comparatively little to long-term outcomes. With these findings, as long as the cost of increasing vocational training uptake is not dramatically larger than the cost of increasing training hours, increasing vocational training uptake is also the more cost-effective of the two intervention pathways. Thus, committing resources to better enable grant recipients to enroll in training would yield the largest return on investment.

8. Discussion and Conclusions

Welfare outcomes materialize over many years, whereas decisions about the programs designed to improve them are made on much shorter time horizons. This timing gap is a central operational challenge of long-term program evaluation, particularly in development contexts, where fragmented administrative data, attrition, and short political and funding cycles are binding frictions. Our framework addresses these frictions by combining short-term experimental data with observational data to recover two objects: an estimate of the long-term treatment effect and a map of the causal mechanism through which the treatment produces this effect. Existing surrogacy methods recover only the first; recovering the second enables a shift from evaluation to improvements in design, allowing program designers to concentrate the next variant on the short-term responses that mediate the long-term effect. In the YOP application, the learned graph identifies year-2 vocational training enrollment as the dominant mediator of the program's persistent effects on durable assets and skilled-trade attachment, pointing to design components such as enrollment vouchers, transportation support, and direct matching with training providers. Two years of post-treatment measurement, combined with observational data on the long-term outcome, suffice to recover the persistent effects established at the nine-year follow-up — letting policymakers decide whether to continue, scale, redesign, or terminate an intervention seven years earlier than the randomized benchmark would permit. The framework is an instance of bottom-up theory building in development operations (Eftekhar et al. 2026): grounded in a concrete operational problem but applicable beyond it.

Our work on recovering mechanisms from short-term data is a starting point, and several extensions would make it more broadly applicable. First, the framework currently ranks a fixed menu of program modifications by worst-case welfare gain across the equivalence class; reformulating this as an optimization problem over a richer space of modifications, subject to cost and feasibility constraints, would let it propose optimal refinements rather than rank prespecified ones. Second, the YOP application shows that two-year measurement suffices to recover the nine-year effect pattern, but how this generalizes is not yet known; characterizing the minimum measurement window required for credible long-term forecasts would help policymakers gauge how quickly they can act. Finally, the mechanism-targeted recommendations are graph-constrained predictions; the natural validation is a randomized comparison of the original program against a variant designed to strengthen the identified mechanism.

References

Aflaki, Arian, Alfonso J. Pedraza-Martinez. 2016. Humanitarian funding in a multi-donor market with donation uncertainty. *Production and Operations Management* **25**(7) 1274–1291. doi:10.1111/poms.12563.

Anderer, Arielle, Hamsa Bastani, John Silberholz. 2022. Adaptive clinical trial designs with surrogates: When should we bother? *Management Science* **68**(3) 1982–2002. doi:10.1287/mnsc.2021.4096.

Aouad, Ali, Kamalini Ramdas, Alp Sungu. 2024. Food subsidies and substitution: A field experiment using digitized micro-grocery transactions in underserved communities. *The Wharton School Research Paper* .

Athey, Susan, Raj Chetty, Guido Imbens. 2020. Combining experimental and observational data to estimate treatment effects on long term outcomes. *arXiv preprint arXiv:2006.09676* .

Athey, Susan, Raj Chetty, Guido W Imbens, Hyunseung Kang. 2019. The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. Tech. rep., National Bureau of Economic Research.

Athey, Susan, Raj Chetty, Guido W. Imbens, Hyunseung Kang. 2025. The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. *The Review of Economic Studies* doi:10.1093/restud/rdaf087. Advance online publication, article rdaf087.

Baird, Sarah, Joan Hamory Hicks, Michael Kremer, Edward Miguel. 2016. Worms at work: Long-run impacts of a child health investment. *The Quarterly Journal of Economics* **131**(4) 1637–1680. doi:10.1093/qje/qjw022.

Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, Christopher Udry. 2015. A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science* **348**(6236) 1260799. doi:10.1126/science.1260799.

Banerjee, Abhijit, Esther Duflo, Garima Sharma. 2021. Long-term effects of the targeting the ultra poor program. *American Economic Review: Insights* **3**(4) 471–486. doi:10.1257/aeri.20200667.

Bareinboim, Elias, Judea Pearl. 2012. Causal inference by surrogate experiments: z-identifiability. *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*. 113–120.

Bareinboim, Elias, Judea Pearl. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* **113**(27) 7345–7352. doi:10.1073/pnas.1510507113.

Barriola, Xabier, William Schmidt. 2025. The high impact of hurricanes on prices in low-income communities. *Manufacturing & Service Operations Management* **27**(4) 993–1007.

Bastani, Hamsa, Mohsen Bayati, Khashayar Khosravi. 2021. Mostly exploration-free algorithms for contextual bandits. *Management Science* **67**(3) 1329–1349. doi:10.1287/mnsc.2020.3605.

Battocchi, Keith, Eleanor Dillon, Maggie Hei, Greg Lewis, Miruna Oprescu, Vasilis Syrgkanis. 2021. Estimating the long-term effects of novel treatments. *Advances in Neural Information Processing Systems* **34** 2925–2935.

Bertsimas, Dimitris, Nathan Kallus. 2020. From predictive to prescriptive analytics. *Management Science* **66**(3) 1025–1044. doi:10.1287/mnsc.2018.3253.

Blattman, Christopher, Nathan Fiala, Sebastian Martinez. 2014. Generating skilled self-employment in developing countries: Experimental evidence from Uganda. *The Quarterly Journal of Economics* **129**(2) 697–752. doi:10.1093/qje/qjt057.

Blattman, Christopher, Nathan Fiala, Sebastian Martinez. 2020. The long-term impacts of grants on poverty: Nine-year evidence from uganda's youth opportunities program. *American Economic Review: Insights* **2**(3) 287–304.

Bouguen, Adrien, Yue Huang, Michael Kremer, Edward Miguel. 2019. Using randomized controlled trials to estimate long-run impacts in development economics. *Annual Review of Economics* **11** 523–561. doi: 10.1146/annurev-economics-080218-030333.

Chickering, David Maxwell. 2002. Optimal structure identification with greedy search. *Journal of Machine Learning Research* **3** 507–554.

Cussens, James, Matti Järvisalo, Janne H. Korhonen, Mark Bartlett. 2017. Bayesian network structure learning with integer programming: Polytopes, facets and complexity. *Journal of Artificial Intelligence Research* **58** 185–229.

Deo, Sarang, Seyed M. R. Iravani, Tingting Jiang, Karen Smilowitz, Stephen Samuelson. 2013. Improving health outcomes through better capacity allocation in a community-based chronic care model. *Operations Research* **61**(6) 1277–1294. doi:10.1287/opre.2013.1214.

Eberhardt, Frederick. 2017. Introduction to the foundations of causal discovery. *International Journal of Data Science and Analytics* **3**(2) 81–91.

Eberhardt, Frederick, Nur Kaynar, Auyon Siddiq. 2025. Discovering causal models with optimization: Confounders, cycles, and instrument validity. *Management Science* **71**(4) 3283–3302.

Eftekhar, Mahyar, Serguei Netessine, Yanchong Zheng. 2026. Development operations management: From frictions to theory. *Available at SSRN 6636898* .

Geiger, Dan, Thomas Verma, Judea Pearl. 1990. Identifying independence in bayesian networks. *Networks* **20**(5) 507–534.

Heinze-Deml, Christina, Marloes H. Maathuis, Nicolai Meinshausen. 2018. Causal structure learning. *Annual Review of Statistics and Its Application* **5** 371–391. doi:10.1146/annurev-statistics-031017-100630.

Huang, Biwei, Kun Zhang, Mingming Gong, Clark Glymour. 2020. Causal discovery from multiple data sets with non-identical variable sets. *Proceedings of the AAAI conference on artificial intelligence*, vol. 34. 10153–10161.

Huang, Shan, Chen Wang, Yuan Yuan, Jinglong Zhao, Brocco (Jingjing) Zhang. 2026. Estimating effects of long-term treatments. *Management Science* doi:10.1287/mnsc.2023.02575. Articles in Advance.

Imbens, Guido, Nathan Kallus, Xiaojie Mao, Yuhao Wang. 2022. Long-term causal inference under persistent confounding via data combination. *arXiv preprint arXiv:2202.07234* .

Imbens, Guido, Nathan Kallus, Xiaojie Mao, Yuhao Wang. 2025. Long-term causal inference under persistent confounding via data combination. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **87**(2) 362–388. doi:10.1093/jrsssb/qkae095.

Imbens, Guido W. 2020. Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature* **58**(4) 1129–1179. doi:10.1257/jel.20191597.

Imbens, Guido W, Donald B Rubin. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Jeon, H Harriet, Claudio Lucarelli, Jean Baptiste Mazarati, Donatien Ngabo, Hummy Song. 2026. Frontiers in operations: Last-mile delivery in healthcare: Drone delivery for blood products in rwanda. *Manufacturing & Service Operations Management* .

Jung, Yonghan, Iván Díaz, Jin Tian, Elias Bareinboim. 2024. Estimating causal effects identifiable from a combination of observations and experiments. *Advances in Neural Information Processing Systems* **36**.

Kaaua, Dawson, Vanitha Virudachalam. 2025. Going the distance: The impact of commute on gender diversity in public service. *Manufacturing & Service Operations Management* **27**(2) 607–623.

Kallus, Nathan, Xiaojie Mao. 2025. On the role of surrogates in the efficient estimation of treatment effects with limited outcome data. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **87**(2) 480–509. doi:10.1093/jrsssb/qkae099.

Kitagawa, Toru, Aleksey Tetenov. 2018. Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* **86**(2) 591–616.

Kotsi, Telesilla O, Owen Q Wu, Alfonso J Pedraza-Martinez. 2022. Donations for refugee crises: In-kind vs. cash assistance. *Manufacturing & Service Operations Management* **24**(6) 3001–3018.

Lee, Sanghack, Juan D Correa, Elias Bareinboim. 2020. General identifiability with arbitrary surrogate experiments. *Uncertainty in artificial intelligence*. PMLR, 389–398.

Levi, Retsef, Manoj Rajan, Somya Singhvi, Yanchong Zheng. 2020. The impact of unifying agricultural wholesale markets on prices and farmers' profitability. *Proceedings of the National Academy of Sciences* **117**(5) 2366–2371.

Long, Elisa F, Eike Nohdurft, Stefan Spinler. 2018. Spatial resource allocation for emerging epidemics: A comparison of greedy, myopic, and dynamic policies. *Manufacturing & Service Operations Management* **20**(2) 181–198.

Maathuis, Marloes H., Diego Colombo. 2015. A generalized back-door criterion. *The Annals of Statistics* **43**(3) 1060–1088. doi:10.1214/14-AOS1295.

Maathuis, Marloes H., Markus Kalisch, Peter Bühlmann. 2009. Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics* **37**(6A) 3133–3164. doi:10.1214/09-AOS685.

Markus Kalisch, Martin Mächler, Diego Colombo, Marloes H. Maathuis, Peter Bühlmann. 2012. Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software* **47**(11) 1–26. doi:10.18637/jss.v047.i11.

McKenzie, David. 2017. How effective are active labor market policies in developing countries? a critical review of recent evidence. *The World Bank Research Observer* **32**(2) 127–154.

Meek, Christopher. 1995. Causal inference and causal explanation with background knowledge. *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*. 403–410.

Molina-Millán, Teresa, Karen Macours. 2025. Attrition in randomized controlled trials: Using tracking information to correct bias. *Economic Development and Cultural Change* **73**(2) 811–834.

Mooij, Joris M., Sara Magliacane, Tom Claassen. 2020. Joint causal inference from multiple contexts. *Journal of Machine Learning Research* **21**(99) 1–108.

Natarajan, Karthik V., Jayashankar M. Swaminathan. 2014. Inventory management in humanitarian operations: Impact of amount, schedule, and uncertainty in funding. *Manufacturing & Service Operations Management* **16**(4) 595–603. doi:10.1287/msom.2014.0497.

Parker, Susan W., Petra E. Todd. 2017. Conditional cash transfers: The case of Progresa/Oportunidades. *Journal of Economic Literature* **55**(3) 866–915. doi:10.1257/jel.20151233.

Pearl, Judea. 1995. Causal diagrams for empirical research. *Biometrika* **82**(4) 669–688. doi:10.1093/biomet/82.4.669.

Pearl, Judea. 2000. Causality: Models, reasoning and inference. *Cambridge, UK: Cambridge University Press* .

Perković, Emilija, Johannes Textor, Markus Kalisch, Marloes H. Maathuis. 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, Jonas, Joris M. Mooij, Dominik Janzing, Bernhard Schölkopf. 2014. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research* **15** 2009–2053.

Prentice, Ross L. 1989. Surrogate endpoints in clinical trials: Definition and operational criteria. *Statistics in Medicine* **8**(4) 431–440. doi:10.1002/sim.4780080407.

Ramdas, Kamalini, Alp Sungu. 2024. The digital lives of the poor: Entertainment traps and information isolation. *Management science* **70**(12) 9089–9100.

Rubin, Donald B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* **66**(5) 688.

Shimizu, Shohei, Patrik O. Hoyer, Aapo Hyvärinen, Antti Kerminen. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research* **7** 2003–2030.

Shpitser, Ilya, Tyler VanderWeele, James M Robins. 2010. On the validity of covariate adjustment for estimating causal effects. *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*. 527–536.

Spirtes, Peter, Clark N Glymour, Richard Scheines, David Heckerman, Christopher Meek, Gregory Cooper, Thomas Richardson. 2000. *Causation, prediction, and search*. MIT press.

Spirtes, Peter, Kun Zhang. 2016. Causal discovery and inference: concepts and recent methodological advances. *Applied informatics*, vol. 3. SpringerOpen, 1–28.

Triantafillou, Sofia, Ioannis Tsamardinos. 2015. Constraint-based causal discovery from multiple interventions over overlapping variable sets. *Journal of Machine Learning Research* **16**(1) 2147–2205.

Uppari, Bhavani Shanker, Serguei Netessine, Ioana Popescu, Rowan P Clarke. 2024. Design of off-grid lighting business models to serve the poor: Field experiments and structural analysis. *Management Science* **70**(5) 3038–3058.

van der Zander, Benito, Maciej Liśkiewicz, Johannes Textor. 2014. Constructing separators and adjustment sets in ancestral graphs. *Proceedings of the UAI 2014 Conference on Causal Inference: Learning and Prediction-Volume 1274*. 11–24.

Verma, Thomas, Judea Pearl. 1990. Equivalence and synthesis of causal models. *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*. 255–270.

Wang, Guihua, Jun Li, Wallace J. Hopp. 2021. An instrumental variable forest approach for detecting heterogeneous treatment effects in observational studies. *Management Science* **68**(5) 3399–3418.

Yang, Jeremy, Dean Eckles, Paramveer Dhillon, Sinan Aral. 2024. Targeting for long-term outcomes. *Management Science* **70**(6) 3841–3855. doi:10.1287/mnsc.2023.4881.

Zhang, Kun, Jonas Peters, Dominik Janzing, Bernhard Schölkopf. 2011. Kernel-based conditional independence test and application in causal discovery. *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI Press, 804–813.

Appendix
A. Graphical Causal Models

Graphical causal models (Pearl, 2000; Spirtes et al., 2000) use directed graphs to represent causal relationships among multiple variables. Let $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ be a directed graph, where $\mathbf{V}$ denotes a set of nodes and $\mathbf{E}$ represents the set of edges, such that $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$. In this graphical representation, each edge in $\mathcal{G}$ signifies a direct causal connection between the corresponding nodes.

We next define some graphical preliminaries that are used throughout the paper. Any two nodes X and Y are called adjacent if there is an edge $X \to Y$ or $X \leftarrow Y$ in the corresponding graph $\mathcal{G}$. The *parents* and *children* of a node X represent its direct causes and effects, respectively, in the graph $\mathcal{G}$. We say that a node Y is a *collider* on a path if its adjacent edges point into Y, i.e., $\to Y \leftarrow$. A *noncollider* on a path is a node Y that is either a *mediator* ($\to Y \to$) or a *common cause* ($\leftarrow Y \to$). In Figure A1(a), node Y is a collider on path $X \to Y \leftarrow Z$, in Figure A1(b) node Y is a mediator on path $X \to Y \to Z$, and in Figure A1(c) node Y is a common cause on path $X \leftarrow Y \to Z$. A *v-structure*, also known as an *unshielded collider*, is a specific configuration of nodes in a graph. A v-structure consists of two parent nodes directing edges towards a common child node, without an edge between the parents, forming a "V" shape. Any node that is connected to node X by a directed path is called a *descendant* of X, while any node connected to X by a directed path is an *ancestor* of X. We refer to the skeleton of a graph $\mathcal{G}$ as the undirected graph obtained by replacing directed edges in $\mathcal{G}$ with undirected edges. For instance, while the edge orientations vary among the graphs in Figure A1, their underlying skeletons are identical.

We define a *path* between two nodes X and Y as a sequence of nodes that starts with X, ends with Y, and where each consecutive pair of nodes in the sequence is connected by an edge in the graph. In this context, the direction of the edges does not matter; as long as there is a connection between any two consecutive nodes in the sequence, it forms a valid path. In other words, although the edge directions are considered, they do not impose any constraints on constructing a path between nodes. A *directed path* from node X to node Y is a path in which all edges point towards node Y. Then, a *directed acyclic graph (DAG)* is a directed graph without cycles.

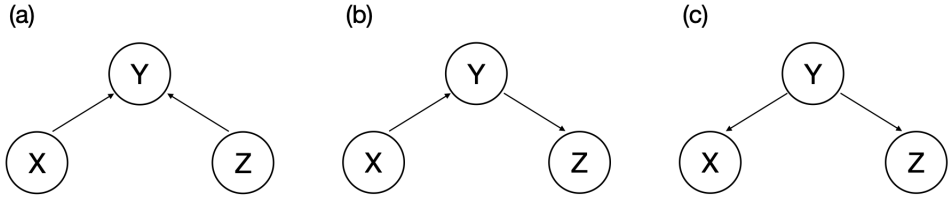

Figure A1 Illustrations of Collider, Mediator, and Common Cause Structures.

A.1. Causal Structure Learning

Causal modeling involves associating a probability distribution, denoted as $P_{\mathcal{G}}(\mathbf{V})$, with a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, which represents the causal relationships among the variables or nodes in the set $\mathbf{V}$. The underlying assumption is that the distribution $P_{\mathcal{G}}(\mathbf{V})$ is generated by the graph structure in a way that allows factorization: $P_{\mathcal{G}}(\mathbf{V}) = \prod_{X \in V} P_{\mathcal{G}}(X|Pa(X))$, where $Pa(X)$ represents the parents of node X in $\mathcal{G}$ (Spirtes and Zhang 2016, Eberhardt 2017). In this context, the terms "node" and "variable" can be used interchangeably, as they correspond to the graphical structure and the probability distribution, respectively.

The two key assumptions that bridge the observed data and the causal structure are stated below.

Assumption 9 **(Causal Markov)**. *Each variable $X \in \mathbf{V}$ in a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ is probabilistically independent of its non-descendants given its parents.*

Assumption 10 **(Faithfulness)**. *The only independences present in the probability distribution are those that are implied by the graph structure through the causal Markov conditions.*

Assumption 11 **(Causal Sufficiency)**. *For any pair of variables in $\mathbf{V}$, all common causes of those variables are also contained within $\mathbf{V}$.*

The first assumption, the causal Markov condition, permits us to transition from the causal graph to the observed probabilistic independencies. Conversely, the faithfulness condition enables us to deduce the structure of the causal graph from observed data independencies. Lastly, the causal sufficiency assumption ensures that all common causes of any pair of variables in the set $\mathbf{V}$ are also contained within $\mathbf{V}$, thereby excluding the existence of any hidden or unobserved confounders.

Due to the close relationship between the causal structure and the resulting data distribution, many algorithms for causal structure learning leverage the identifiable independence structure in the data to make inferences about the underlying causal relationships. A key concept essential for this inference is *d-separation* (Geiger et al. 1990), often considered as the graphical equivalent of probabilistic independence. It is based on the notion of a *blocked path*:

Definition 3 (Blocked paths). A path between nodes X and Y is considered *unblocked* with respect to a set of nodes $\mathbf{C}$ if every collider Z on the path is in $\mathbf{C}$ or has a descendant in $\mathbf{C}$, and no other nodes on the path are in $\mathbf{C}$. If these conditions do not hold, the path is considered *blocked* with respect to $\mathbf{C}$ (Pearl 2000).

We can now introduce the concept of d-separation.

Definition 4 (d-separation). Two nodes X and Y are said to be *d-separated* with respect to a conditioning set $\mathbf{C}$ (denoted as $i \perp j|\mathbf{C}$) if all paths between them are blocked. Conversely, if there exists at least one unblocked path between X and Y given $\mathbf{C}$, they are considered *d-connected* (denoted as $i \not\perp j|\mathbf{C}$) (Pearl 2000).

REMARK 1. Under the causal Markov and faithfulness conditions, a (conditional) independence in $P_{\mathcal{G}}(\mathbf{V})$ is present if and only if there is a corresponding (conditional) d-separation in DAG $\mathcal{G}$ (Pearl 2000).

Remark 1 highlights the correspondence between (conditional) independence in $P_{\mathcal{G}}(\mathbf{V})$ and (conditional) d-separation in DAG $\mathcal{G}$ under the causal Markov and faithfulness conditions. This correspondence serves as the fundamental framework for a wide range of causal structure learning methods, providing a means to leverage observed independence patterns in data to infer underlying causal relations.

To illustrate the notion of blocked paths and their connection to the principles outlined in Remark 1, we refer to Figure A1. In Figure A1(a), we see the path $X \to Y \leftarrow Z$ where Y serves as a collider. Given the conditioning set $\mathbf{C} = \emptyset$, this path is considered 'blocked' since the collider, Y, is not included in $\mathbf{C}$. The path $X \to Y \leftarrow Z$ is the only path between X and Z within this figure. Since it is blocked, we establish that X and Z are d-separated, signifying a lack of information flow between the two nodes. By invoking Remark 1, we anticipate that X and Z to be marginally probabilistically independent, in line with their status of being d-separated within the graph. On the other hand, when $\mathbf{C} = Y$, the path $X \to Y \leftarrow Z$ transitions to being 'unblocked'. This is because Y, being the only collider on this path, is now included within the conditioning set $\mathbf{C}$. Hence by Remark 1, we expect that X and Z to be probabilistically dependent with respect to conditioning set $\mathbf{C} = Y$. Conversely, in Figure A1 (b) and (c), the paths $X \to Y \to Z$ and $X \leftarrow Y \to Z$ respectively are 'unblocked' when the conditioning set is $\mathbf{C} = \emptyset$. This is due to the absence of colliders on these paths and the fact that Y, acting as a noncollider, is not included in the conditioning set $\mathbf{C}$. Hence, we expect that X and Z to be marginally dependent. However, conditioning on $\mathbf{C} = Y$ blocks these paths and X and Z become marginally independent with respect to conditioning set $\mathbf{C} = Y$.

Despite different causal relationships in Figure A1(b) and A1(c), they imply the same independence relations. In contrast, the independence relations in Figure A1(a), i.e., $X \perp\!\!\!\perp Z$ and $X \not\perp\!\!\!\perp Z \mid Y$ uniquely identifies Y as a collider. These differences underscore the fact that colliders leave distinct signatures on conditional independence patterns. However, mediator and common causes can result in identical patterns in conditional independence relations. This demonstrates that the observable data cannot uniquely identify the underlying causal graph, as both mediators and common causes can generate identical patterns. This idea is captured by the concept of *Markov equivalence class*. Two graphs that have the same independence structure are said to be *Markov equivalent*. This means that the independence relations represented by each graph in the equivalence class are identical, even though the causal relationships they represent are not. Two DAGs are in the same Markov equivalence class if and only if they have the same skeleton and the same v-structures (Verma and Pearl 1990).

A.2. Identification of Average Treatment Effect using Causal Graphs

The previous section discusses how to discover causal graphs, which are essential for learning presence or absence of causal relationships among variables. However, understanding the structure of underlying causal relations is only the first step. Many policy decisions are also interested in quantifying the causal effects of specific interventions. From the joint distribution of two variables, W and Y, we can derive the conditional probability $P(Y \mid W)$ using observational data, which tells us the probability of Y given that W takes on a specific value w, i.e., $P(Y \mid W = w)$. However, what we often need for policy decisions is the causal effect of setting W to a specific value w, represented as $P(Y \mid \text{do}(W = w))$. This notation distinguishes the intervention $\text{do}(W = w)$ from mere observation. While we can calculate $P(Y \mid W = w)$ directly from the joint distribution of W and Y in an observational data, the challenge lies in determining whether this observational data, combined with the underlying causal structure, allows us to infer the causal effect $P(Y \mid \text{do}(W = w))$.

With this notation, given a graph $\mathcal{G}$, the average effect of a binary treatment can be defined as:

$$\tau_{\mathcal{G}} = E[Y|do(W = 1)] - E[Y|do(W = 0)].^{9} \tag{A1}$$

Identification of the average treatment effect $\tau_{\mathcal{G}}$ from observational data is challenging primarily due to the presence of confounders that can induce spurious correlations between the treatment W and the outcome Y. For example, let's consider the causal graph depicted in Figure A2. In this graph, there are three paths between W and Y: $W \to X_1 \to X_2 \leftarrow X_3 \to Y$, $W \to X_4 \to Y$, and $W \leftarrow X_5 \to X_6 \to Y$. By Definition 3, the path $W \to X_1 \to X_2 \leftarrow X_3 \to Y$ is *blocked* as the variable X_2 is a collider and both $W \to X_4 \to Y$ and $W \leftarrow X_5 \to X_6 \to Y$ are *unblocked* as there are no colliders on these paths. However, only the path $W \to X_4 \to Y$ represents a causal effect of treatment W on outcome Y. The path $W \leftarrow X_5 \to X_6 \to Y$, in contrast, reflects confounding. The association observed between W and Y along this path is not due to a causal mechanism from W to Y, but is a byproduct of their mutual associations through X_5 and X_6. Thus, without controlling for the confounders X_5 or X_6, the observed correlation in the data between W and Y will reflect both the actual causal effect from W to Y and the spurious correlation introduced by the path via X_5 and X_6. To accurately estimate the average treatment effect τ, it is vital to adjust for these confounders and thereby isolate the causal effect of W on Y.

As demonstrated in the example above, adjusting for relevant covariates is crucial to accurately estimate causal effects from observational data. This method, known as covariate adjustment, helps

[9] In the potential outcomes framework, $do(W = 1)$ and $do(W = 0)$ correspond to the notation $Y(1)$ and $Y(0)$, respectively. Therefore, $E[Y|do(W = 1)] - E[Y|do(W = 0)] = E[Y(1)] - E[Y(0)]$. For more details, see Rubin (1974) and Imbens and Rubin (2015).

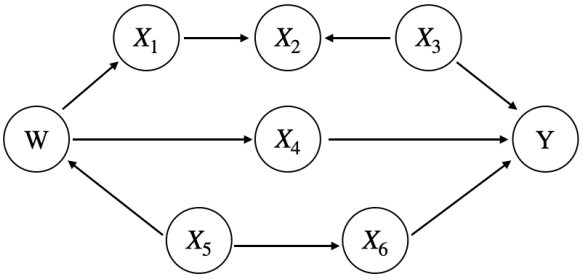

Figure A2 A figure to illustrate the back-door criterion.

isolate the true causal relationship between the treatment and the outcome by controlling for confounders that could otherwise bias the results (Pearl 1995, Shpitser et al. 2010). One of the most well-known methods for this is Pearl's back-door adjustment. When estimating the effect of W on Y, a *back-door path* is defined as any unblocked path connecting W to Y that begins with an arrow pointing towards W. Such paths introduce potential confounding by offering a non-causal route for information flow between W and Y. Pearl (2000) introduced the *back-door criterion* as a method for identifying sets of variables that, when conditioned on, can block the backdoor paths and allow for the estimation of causal effects from observational data. The back-door criterion provides a graphical test to determine a valid adjustment set to estimate the causal effect from observational data.

Definition 5. **Back-Door Criterion (Pearl 2000).**[10] Let $\mathbf{W}, \mathbf{Y}$ and $\mathbf{Z}$ be pairwise disjoint sets of vertices in a DAG $\mathcal{G}$. $\mathbf{Z}$ satisfies the *back-door criterion* relative to $\mathbf{W}, \mathbf{Y}$ in $\mathcal{G}$ if

(i) no node in $\mathbf{Z}$ is a descendant of any node in $\mathbf{W}$, and

(ii) for every $W \in \mathbf{W}$, the set $\mathbf{Z} \cup \mathbf{W} \setminus \{W\}$ blocks every back-door path from W to any member of $\mathbf{Y} \in \mathcal{G}$.

Theorem 4 ***Back-Door Criterion (Pearl 2000).*** *If a set of variables* $\mathbf{Z}$ *satisfies the back-door criterion relative to* $(\mathbf{W}, \mathbf{Y})$ *in a DAG* $\mathcal{G}$*, then the causal effect of* $\mathbf{W}$ *on* $\mathbf{Y}$ *is identifiable and is given by the formula*

$$P(\mathbf{Y} = \mathbf{y} | do(\mathbf{W} = \mathbf{w})) = \sum_{\mathbf{z} \in \mathcal{Z}} P(\mathbf{Y} = \mathbf{y} | \mathbf{W} = \mathbf{w}, \mathbf{Z} = \mathbf{z}) P(\mathbf{Z} = \mathbf{z}) \tag{A2}$$

where $\mathcal{Z}$ *is the support of* $\mathbf{Z}$*.*

[10] The definition is adapted from Pearl's to explicitly consider cases where W and Y are sets of vertices, in accordance with the set-based analysis framework in Maathuis and Colombo (2015).

Corollary 4.1 ***ATE with Back-Door Criterion****. Let* $\mathbf{W}$ *be a binary treatment variable. If a set of variables* $\mathbf{Z}$ *satisfies the back-door criterion relative to* $(\mathbf{W}, \mathbf{Y})$ *in a DAG* $\mathcal{G}$*, then the ATE* $\tau_{\mathcal{G}}$ *of* W *on* Y *is identifiable and is given by the formula*

$$\tau_{\mathcal{G}} = \sum_{y \in \mathscr{Y}} \sum_{z \in \mathscr{Z}} \mathbf{y} P(\mathbf{Y} = \mathbf{y} | \mathbf{W} = \mathbf{1}, \mathbf{Z} = \mathbf{z}) P(\mathbf{Z} = \mathbf{z}) - \sum_{\mathbf{y} \in \mathscr{Y}} \sum_{\mathbf{z} \in \mathscr{Z}} y P(\mathbf{Y} = \mathbf{y} | \mathbf{W} = \mathbf{0}, \mathbf{Z} = \mathbf{z}) P(\mathbf{Z} = \mathbf{z}) \tag{A3}$$

where $\mathscr{Y}$ *and* $\mathscr{Z}$ *are the supports of* $\mathbf{Y}$ *and* $\mathbf{Z}$*, respectively.*

This corollary demonstrates that if we have access to control variables satisfying the back-door criterion, we can accurately estimate the average treatment effect. This is achieved by calculating the expected outcomes conditioned on both the treatment and covariate values, and then weighting these outcomes by the probability of observing the covariate values. However, the back-door adjustment criterion is not complete (Pearl 2000). Thus, there are causal graphs where the back-door criterion does not identify a valid adjustment set, but adjusting for different covariate sets can still accurately estimate the causal effect.

There is a vast amount of research focused on finding necessary and sufficient graphical criteria for the selection of adjustment sets to accurately estimate causal effects. Shpitser et al. (2010) extend the back-door criterion and provide a necessary and sufficient graphical criterion for adjustment in DAGs. Other research has focused on constructing adjustment sets that are valid for more general graph classes, including those with unobserved confounders or structures that represent Markov equivalence classes (van der Zander et al. 2014, Maathuis and Colombo 2015, **?**). However, these methods are designed for single-sample settings and are not directly applicable when combining multiple datasets. Our work aligns with a recent stream of research in causal graphical models that combines multiple datasets collected under heterogeneous conditions, such as different populations and various sampling methods, with the possibility of sampling biases (Bareinboim and Pearl 2012, Lee et al. 2020, Jung et al. 2024). However, these studies *do not* consider the scenario of short-term experimental data with missing long-term outcomes combined with long-term historical data. We contribute to this literature by algorithmically selecting surrogates and covariates to derive a *novel* closed-form expression for the long-term treatment effect that is *guaranteed* to be computed from the available samples.

B. Stage-Aware Pruning of Conditioning Sets

Algorithm 1 as stated tests every subset of $\mathbf{V}^E \setminus \{V_i, V_j\}$ (respectively $\mathbf{V}^O \setminus \{V_i, Y_k\}$) as a candidate conditioning set, which scales with $|\mathbf{V}|$. The stage ordering $\mathbf{X} \prec W \prec \mathbf{S} \prec \mathbf{Y}$ admits a strict reduction: conditioning variables in stages strictly later than both endpoints of the tested pair are never required to establish conditional independence and may be excluded a priori. For instance, testing

independence between two pre-treatment covariates requires conditioning sets only from $\mathbf{X}$ itself; testing between a covariate and the treatment, only from $\mathbf{X} \setminus \{X_i\}$; testing between two surrogates, only from $\mathbf{X} \cup \{W\} \cup \mathbf{S}$. The following proposition makes this precise.

LEMMA 1 **(Later-Stage Variables in Conditioning Sets)**. *Let P be a distribution that is Markov and faithful to a DAG $\mathcal{G}$ respecting the stage ordering $\mathbf{X} \prec W \prec \mathbf{S} \prec \mathbf{Y}$. For any pair $V_i, V_j \in \mathbf{V}$, let $\mathbf{F}_{ij}$ denote the set of variables in stages strictly later than both V_i and V_j. If $V_i \perp\!\!\!\perp V_j \mid \mathbf{C}$ under P for some $\mathbf{C} \subseteq \mathbf{V} \setminus \{V_i, V_j\}$, then $V_i \perp\!\!\!\perp V_j \mid \mathbf{C} \setminus \mathbf{F}_{ij}$ under P as well. Consequently, whenever V_i and V_j are conditionally independent given some conditioning set, there exists one that contains no later-stage variables.*

The intuition is that any path between V_i and V_j that enters a stage strictly later than both endpoints must contain a collider in a maximal block of latest-stage nodes on that path. Since the stage ordering prevents directed edges from moving to earlier stages, every descendant of that collider also lies in a stage strictly later than both endpoints. Removing $\mathbf{F}_{ij}$ from the conditioning set therefore removes the collider and all of its possible conditioned descendants, leaving the path blocked. Hence no later-stage variable is ever needed to witness conditional independence between V_i and V_j.

Proof of Lemma 1. Let $\sigma(V)$ denote the stage of a variable V, with $\sigma(X) = 1$ for $X \in \mathbf{X}$, $\sigma(W) = 2$, $\sigma(S) = 3$ for $S \in \mathbf{S}$, and $\sigma(Y) = 4$ for $Y \in \mathbf{Y}$. Since $\mathcal{G}$ respects the stage ordering, every directed edge $A \to B$ in $\mathcal{G}$ satisfies $\sigma(A) \leq \sigma(B)$.

Fix $V_i, V_j \in \mathbf{V}$ and let $r_{ij} = \max\{\sigma(V_i), \sigma(V_j)\}$. Then $\mathbf{F}_{ij} = \{V \in \mathbf{V} : \sigma(V) > r_{ij}\}$. Define $\mathbf{C}^- = \mathbf{C} \setminus \mathbf{F}_{ij}$. If $\mathbf{F}_{ij} = \emptyset$, then $\mathbf{C}^- = \mathbf{C}$ and the result is immediate, so suppose $\mathbf{F}_{ij} \neq \emptyset$.

Because P is faithful to $\mathcal{G}$, the conditional independence $V_i \perp\!\!\!\perp V_j \mid \mathbf{C}$ implies that V_i and V_j are d-separated by $\mathbf{C}$ in $\mathcal{G}$. We show that V_i and V_j are also d-separated by $\mathbf{C}^-$. Consider an arbitrary path $p = (U_0, U_1, \ldots, U_m)$ from V_i to V_j, where $U_0 = V_i$ and $U_m = V_j$. It is enough to show that p is blocked by $\mathbf{C}^-$.

First suppose that p contains no node in $\mathbf{F}_{ij}$. If p were open given $\mathbf{C}^-$, then no noncollider on p would belong to $\mathbf{C}^-$, and every collider on p would have itself or a descendant in $\mathbf{C}^-$. Since p contains no node in $\mathbf{F}_{ij}$, the sets $\mathbf{C}$ and $\mathbf{C}^-$ agree on all noncolliders on p. Hence no noncollider on p would belong to $\mathbf{C}$. Also, since $\mathbf{C}^- \subseteq \mathbf{C}$, every collider activated by $\mathbf{C}^-$ would also be activated by $\mathbf{C}$. Therefore p would be open given $\mathbf{C}$, contradicting the fact that $\mathbf{C}$ d-separates V_i and V_j. Thus p is blocked by $\mathbf{C}^-$.

Now suppose that p contains at least one node in $\mathbf{F}_{ij}$. Let $L = \max_{0 \leq t \leq m} \sigma(U_t)$. Since p contains a node in $\mathbf{F}_{ij}$, we have $L > r_{ij}$. Choose a maximal contiguous block $U_a, U_{a+1}, \ldots, U_b$ of nodes on p whose stage is L. Because the endpoints $U_0 = V_i$ and $U_m = V_j$ have stages at most $r_{ij} < L$, this block

cannot contain either endpoint, so $0 < a \le b < m$. By maximality of the block, the neighboring nodes U_{a-1} and U_{b+1} have stages strictly below L. The stage ordering therefore forces the boundary edges to be oriented as $U_{a-1} \to U_a$ and $U_{b+1} \to U_b$.

We claim that at least one node in the block $U_a, \dots, U_b$ is a collider on p. If $a = b$, then both adjacent edges point into U_a, so U_a is a collider. If $a < b$, suppose for contradiction that no node in the block is a collider. Since $U_{a-1} \to U_a$, avoiding a collider at U_a requires the next edge to be $U_a \to U_{a+1}$. Then avoiding a collider at U_{a+1} requires $U_{a+1} \to U_{a+2}$. Continuing in this way, we obtain $U_a \to U_{a+1} \to \cdots \to U_b$. In particular, the edge adjacent to U_b from the left is $U_{b-1} \to U_b$. But the right boundary edge is also $U_{b+1} \to U_b$, so both adjacent edges on p point into U_b, making U_b a collider. This contradiction proves that the block contains a collider; call one such collider Q.

Since $\sigma(Q) = L > r_{ij}$, we have $Q \in \mathbf{F}_{ij}$. Moreover, every descendant of Q also belongs to $\mathbf{F}_{ij}$. Indeed, along any directed path starting from Q, stages cannot decrease, so every descendant D of Q satisfies $\sigma(D) \ge \sigma(Q) = L > r_{ij}$. Thus $\{Q\} \cup \mathrm{Desc}_{\mathcal{G}}(Q) \subseteq \mathbf{F}_{ij}$, and hence neither Q nor any descendant of Q belongs to $\mathbf{C}^- = \mathbf{C} \setminus \mathbf{F}_{ij}$. Therefore Q is a collider on p with no conditioned descendant under $\mathbf{C}^-$, so p is blocked by $\mathbf{C}^-$.

We have shown that every path from V_i to V_j is blocked by $\mathbf{C}^-$. Thus V_i and V_j are d-separated by $\mathbf{C}^-$ in $\mathcal{G}$. Since P is Markov with respect to $\mathcal{G}$, d-separation implies conditional independence under P. Therefore $V_i \perp\!\!\!\perp V_j \mid \mathbf{C} \setminus \mathbf{F}_{ij}$, as claimed.□

C. Meek Rules Illustration

The Meek rules provide a systematic method for orienting undirected edges within a partially directed acyclic graph to achieve a fully directed acyclic graph without introducing new unshielded collider or cycles (Meek 1995). Figure A3 demonstrates these rules visually, showcasing how they apply in different configurations of a graph.

D. Rules of Do-Calculus

Let $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, and $\mathbf{W}$ be arbitrary disjoint sets of nodes in a causal DAG $\mathcal{G}$ and let P be the probability distribution induced by graph $\mathcal{G}$. According to Theorem 3.4.1 in Pearl (2000), the following rules apply to any disjoint subsets of variables $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, and $\mathbf{W}$:

1. **Rule 1 (Insertion/deletion of observations):**

$$P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{z}, \mathbf{w}) = P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{w}) \text{ if } (\mathbf{Y} \perp \mathbf{Z} \mid \mathbf{X}, \mathbf{W})_{\mathcal{G}_{\overline{\mathbf{X}}}}.$$

2. **Rule 2 (Action/observation exchange):**

$$P(\mathbf{y} \mid do(\mathbf{x}), do(\mathbf{z}), \mathbf{w}) = P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{z}, \mathbf{w}) \text{ if } (\mathbf{Y} \perp \mathbf{Z} \mid \mathbf{X}, \mathbf{W})_{\mathcal{G}_{\overline{\mathbf{X}}\underline{\mathbf{Z}}}}.$$

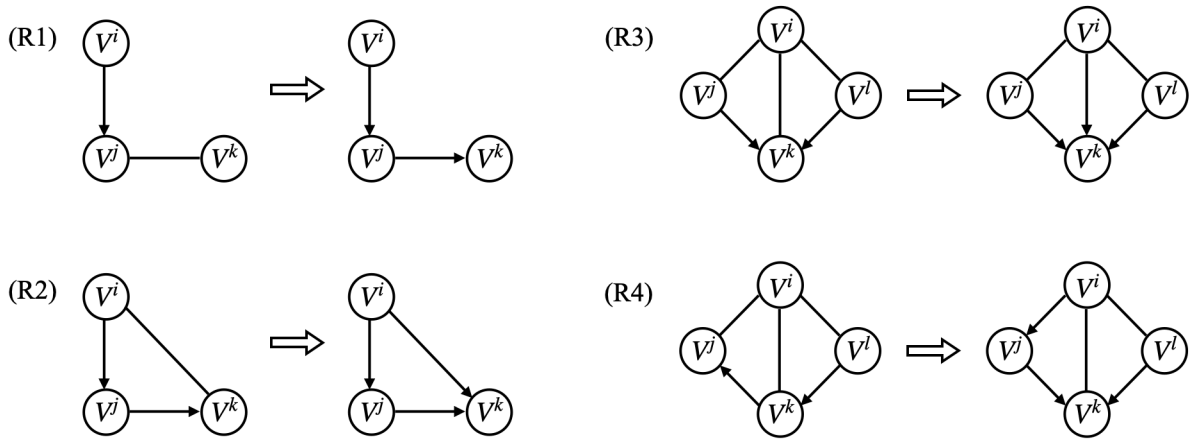

Figure A3 **Meek rules illustrated.**

3. **Rule 3 (Insertion/deletion of actions):**

$$P(\mathbf{y} \mid do(\mathbf{x}), do(\mathbf{z}), \mathbf{w}) = P(\mathbf{y} \mid do(\mathbf{x}), \mathbf{w}) \text{ if } (\mathbf{Y} \perp \mathbf{Z} \mid \mathbf{X}, \mathbf{W})_{\mathcal{G}_{\overline{\mathbf{X},\mathbf{Z}(\mathbf{W})}}}.$$

where $\mathcal{G}_{\overline{X}}$ denotes the graph obtained by deleting all incoming arrows to nodes in $\mathbf{X}$ from $\mathcal{G}$, $\mathcal{G}_{\underline{\mathbf{X}}}$ denotes the graph obtained by deleting all outgoing arrows from nodes in $\mathbf{X}$ from $\mathcal{G}$, and $\mathbf{Z}(\mathbf{W})$ is the set of $\mathbf{Z}$-nodes that are not ancestors of any $\mathbf{W}$-nodes in $\mathcal{G}_{\overline{\mathbf{X}}}$.

E. Equivalent Graphs Can Imply Different Policy Gains

This appendix illustrates Proposition 2: two graphs in the same Markov equivalence class can imply different gains from the same program modification. The equivalence class fixes the skeleton and the v-structures but not every edge orientation, so an undirected edge among the short-term variables admits more than one completion—and different completions route the treatment's effect to the outcome through different paths.

Focus on the downstream subgraph connecting the short-term variables $\{S_1, S_2, S_3\}$ to the long-term outcome Y. The edges $S_1 \to Y$, $S_2 \to Y$, and $S_3 \to Y$ are oriented, while the edge between S_1 and S_2 remains unoriented. The equivalence class contains two DAG completions: $\mathcal{G}$ (with $S_1 \to S_2$), and $\mathcal{G}'$ with $S_2 \to S_1$) that share the same skeleton and the same unshielded colliders. However, the unresolved orientation of the edge between S_1 and S_2 determines which surrogate has an indirect path to Y through the other. Under $\mathcal{G}$ S_1 reaches Y both directly and via S_2,:

$$\Gamma_{S_1,Y,\mathcal{G}} = \beta^{\mathcal{G}}_{S_1,Y} + \beta^{\mathcal{G}}_{S_1,S_2}\beta^{\mathcal{G}}_{S_2,Y} \qquad \text{and} \qquad \Gamma_{S_2,Y,\mathcal{G}} = \beta^{\mathcal{G}}_{S_2,Y},$$

and under $\mathcal{G}'$ their roles reverse

$$\Gamma_{S_1,Y,\mathcal{G}'} = \beta^{\mathcal{G}'}_{S_1,Y} \qquad \text{and} \qquad \Gamma_{S_2,Y,\mathcal{G}'} = \beta^{\mathcal{G}'}_{S_2,Y} + \beta^{\mathcal{G}'}_{S_2,S_1}\beta^{\mathcal{G}'}_{S_1,Y}$$

Therefore, a modification that increases the first-stage effect on S_1 by δ_1 has gain $\delta_1 \Gamma_{S_1,Y,\mathcal{G}}$ under $\mathcal{G}$ and $\delta_1 \Gamma_{S_1,Y,\mathcal{G}'}$ under $\mathcal{G}'$, which need not be equal.

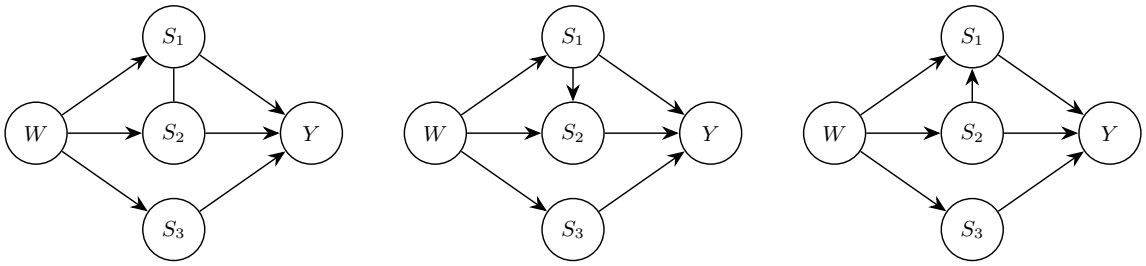

(a) CPDAG: $S_1 - S_2$ undirected (b) DAG $\mathcal{G}$: $S_1 \to S_2$ (c) DAG $\mathcal{G}'$: $S_2 \to S_1$

Figure A4 COMB-PC orients channels $W \to S_j$ for $j = 1, 2, 3$ but leaves the edge S_1–S_2 undirected (a). The two DAG completions $\mathcal{G}$ (b) and $\mathcal{G}'$ (c) constitute the equivalence class G. The channel set $C = \{1, 2, 3\}$ is invariant across G; the leverages of S_1 and S_2 on Y, however, differ between $\mathcal{G}$ and $\mathcal{G}'$.

F. Numerical Experiments with Synthetic Data

In this section, we evaluate the performance of COMB-PC algorithm and the proposed treatment effect identification strategy using synthetic data. We consider four distinct scenarios varying the edge inclusion probability p. The edge inclusion probability p is the fixed probability with which each potential edge between any pair of nodes in a graph is included, thereby determining the overall density of the graph. We use the following probabilities in our study: 0.2, 0.3, 0.4, and 0.5. In each of these scenarios, we generate 200 DAGs, with the number of nodes varying from 10 to 15. We ensure that the generated DAGs satisfy the Assumptions 2 and 3. These DAGs are parameterized as linear Gaussian models and we simulate experimental and observational samples with 50,000 observations. Then, we conduct conditional independence tests using partial correlations, applying a significance threshold of $\alpha = 0.05$, adjusted with a Bonferroni correction to account for multiple hypothesis testing. We use the *pcalg package* (Markus Kalisch et al. 2012) for both generation of DAGs and samples as well as for conducting conditional independence tests.

Table 1 assesses the performance of the COMB-PC algorithm at various edge inclusion probabilities. The MEC column shows the average number of graphs within the Markov equivalence class for each density scenario. As the edge inclusion probability increases from 0.2 to 0.5, the number of graphs within the same Markov equivalence class decreases. As the edge inclusion probability increases, paralleling the decrease in the number of graphs within the Markov equivalence class, we also observe a corresponding reduction in the number of undirected edges. Next, we calculate the average true positive rate (TPR), false positive rate (FPR), true negative rate (TNR), and

false negative rate (FNR) for the graphs returned by the COMB-PC algorithm. The true positive rate (TPR), which indicates the algorithm's accuracy in correctly identifying true causal directions, shows a decreasing pattern, dropping from 0.858 at a density of 0.2 to 0.601 at a density of 0.5. This decline suggests a reduced accuracy in identifying causal relationships in denser graphs. Conversely, the false positive rate (FPR) increases with graph density, rising from 0.005 to 0.033, which points to a higher rate of incorrectly identified causal directions in denser networks. Complementing these trends, the true negative rate (TNR) slightly decreases, and the false negative rate (FNR) significantly increases with higher densities.

Table A1 Evaluating COMB-PC Algorithm's Performance Across Different Edge Densities.

Edge inclusion probability	MEC	Unoriented Edges	Oriented Edges			
			TPR	**FPR**	**TNR**	**FNR**
$p = 0.2$	3.445	1.565	0.858	0.005	0.995	0.142
$p = 0.3$	3.260	1.255	0.813	0.01	0.99	0.187
$p = 0.4$	2.015	0.82	0.786	0.019	0.981	0.214
$p = 0.5$	1.405	0.37	0.601	0.033	0.967	0.399

Having established the accuracy of the COMB-PC algorithm for identifying causal relationships, we next evaluate our proposed treatment effect identification strategy. We begin by calculating the true average treatment effect $\tau_{\mathcal{G}^*}$ for each true underlying graph $\mathcal{G}^*$ by considering all directed paths from the treatment to the outcome and multiplying the true edge coefficients along these paths. We then empirically estimate the treatment effect $\tau_{\mathcal{G}}$ using short-term experimental and long-term observational data. To do this, we identify relevant surrogate and backdoor adjustment sets over the graphs returned by the COMB-PC algorithm. Table A2 summarizes the estimation errors, expressed as $(\tau - \tau^*)/\tau^*$, assessing the deviation of the estimated effects from the true effects across 200 simulations for varying graph densities. The results reveal that estimation errors are minimal in sparser graphs, with a mean error of 0.048 at $p = 0.2$. As graph density increases, mean errors rise, reaching 0.205 at $p = 0.5$. This trend underscores the crucial role of precise causal graph identification, as the diminishing accuracy in causal structure learning within denser graphs significantly compromises the reliability of treatment effect estimations.

Table A2 Average Treatment Effect Estimation Errors Across Varying Graph Densities.

	min	25%	mean	50%	75%	max
$p = 0.2$	-0.186	-0.009	0.048	0.001	0.011	1.113
$p = 0.3$	-0.524	-0.004	0.125	0.006	0.113	1.0
$p = 0.4$	-0.166	-0.005	0.126	0.002	0.115	1.065
$p = 0.5$	-0.233	-0.009	0.205	0.020	0.321	1.021

G. Proofs

LEMMA 2. *Two DAGs are equivalent if and only if they share the same skeleton and same v-structures.*

Proof of Lemma 2. The proof of this lemma can be found in Verma and Pearl (1990).

Proof of Theorem 1. Let skel($\mathcal{G}$) denote the skeleton of a graph $\mathcal{G}$. By Assumptions 1 and 7, every conditional-independence query made in either sample is a query about the same distribution $P_{\mathcal{G}^*}$ induced by the true DAG $\mathcal{G}^*$. By Assumptions 4, 5, and 6, conditional independence in $P_{\mathcal{G}^*}$ is equivalent to d-separation in $\mathcal{G}^*$ over the observed variable set $\mathbf{V}$.

We first show that Phase 1 recovers the true skeleton. The edge W—Y_k is not included in the initial graph for any $Y_k \in \mathbf{Y}$. This agrees with $\mathcal{G}^*$: Assumption 3(i) rules out $Y_k \to W$, and Assumption 3(ii) rules out $W \to Y_k$. Now consider any other pair A, B tested in Phase 1. If Algorithm 1 removes A—B, then for some tested set $\mathbf{C}$ it found $A \perp\!\!\!\perp B \mid \mathbf{C}$ in the relevant sample. By the oracle assumption and Markov-faithfulness, A and B are d-separated by $\mathbf{C}$ in $\mathcal{G}^*$. Hence A and B cannot be adjacent in $\mathcal{G}^*$, because the length-one path between adjacent vertices is open given any conditioning set that excludes the endpoints.

Conversely, suppose A and B are not adjacent in $\mathcal{G}^*$. If $A, B \in \mathbf{V}^E$, choose an endpoint, say B, such that A is not a descendant of B; such an endpoint exists because $\mathcal{G}^*$ is acyclic. Since A is not a parent of B, the local Markov property implies that A and B are d-separated by $\mathrm{Pa}_{\mathcal{G}^*}(B)$. By Assumption 3(i), no long-term outcome can be a parent of any node in $\mathbf{V}^E$, so $\mathrm{Pa}_{\mathcal{G}^*}(B) \subseteq \mathbf{V}^E \setminus \{A, B\}$. Moreover, $|\mathrm{Pa}_{\mathcal{G}^*}(B)| \le |\mathbf{V}^E| - 2 \le |\mathbf{V}| - 3$, so Algorithm 1 tests this set and removes A—B if the edge has not already been removed. If instead the tested pair is A, Y_k with $A \in \mathbf{V}^O \setminus \{Y_k\}$, use the same argument with a long-term-outcome endpoint. If A is not a descendant of Y_k, then $\mathrm{Pa}_{\mathcal{G}^*}(Y_k)$ d-separates A and Y_k; if $A = Y_\ell$ is a descendant of Y_k, then $\mathrm{Pa}_{\mathcal{G}^*}(Y_\ell)$ d-separates them. In either case, Assumption 3(ii) rules out W as a parent of a long-term outcome, so the separating set lies in $\mathbf{V}^O \setminus \{A, Y_k\}$ and has size at most $|\mathbf{V}^O| - 2 = |\mathbf{V}| - 3$. Therefore Algorithm 1 tests such a set and removes the edge. Thus $\mathrm{skel}(\mathcal{G}_1) = \mathrm{skel}(\mathcal{G}^*)$.

We next show that Phase 2 produces a partially directed graph $\mathcal{G}_2$ that is consistent with $\mathcal{G}^*$. Consider first the collider-orientation step. Let $A - B - C$ be an unshielded triple in $\mathcal{G}_1$ with $B \neq W$, with $\{A, C\} \neq \{W, Y_k\}$ for every $Y_k \in \mathbf{Y}$, and with $B \notin SepSet_{AC}$. Since Phase 1 recovered the true skeleton, A and C are non-adjacent in $\mathcal{G}^*$ while A and B, and B and C, are adjacent in $\mathcal{G}^*$. If B were a noncollider on the path $A - B - C$ in $\mathcal{G}^*$, then every separating set for A and C would have to contain B, because otherwise the length-two path $A - B - C$ would remain open. This contradicts $B \notin SepSet_{AC}$. Hence B is a collider in $\mathcal{G}^*$, so the orientation $A \to B \leftarrow C$ is correct.

The structural-orientation step is also correct. By Assumption 3(i), any surviving edge between $X_i \in \mathbf{X}$ and W must be oriented as $X_i \to W$, any surviving edge between a variable in $\mathbf{X} \cup \{W\}$ and a short-term variable $S_j \in \mathbf{S}$ must be oriented into S_j, and any surviving edge between a variable in $\mathbf{V} \setminus \mathbf{Y}$ and a long-term outcome $Y_k \in \mathbf{Y}$ must be oriented into Y_k. Assumption 3(ii) has already ruled out the exceptional direct edge $W \to Y_k$. Finally, Meek's rules are sound: starting from a partially directed graph whose directed edges agree with $\mathcal{G}^*$, they orient only edges that are compelled by the current skeleton and unshielded colliders, without introducing new unshielded colliders or directed cycles. Therefore every directed edge in $\mathcal{G}_2$ is oriented as in $\mathcal{G}^*$, and $\mathcal{G}^*$ is a consistent extension of $\mathcal{G}_2$. Since Step 4 stores all DAGs consistent with $\mathcal{G}_2$, we have $\mathcal{G}^* \in \mathbf{G}$.

It remains to show that every $\mathcal{G} \in \mathbf{G}$ is Markov equivalent to $\mathcal{G}^*$. Every graph in $\mathbf{G}$ is a consistent extension of $\mathcal{G}_2$, so every such graph has the same skeleton as $\mathcal{G}_2$, hence the same skeleton as $\mathcal{G}^*$. We also show that the unshielded colliders coincide. Every unshielded collider in $\mathcal{G}_2$ is an unshielded collider in $\mathcal{G}^*$ because the skeletons are the same and every arrowhead in $\mathcal{G}_2$ agrees with $\mathcal{G}^*$. Conversely, let $A \to B \leftarrow C$ be any unshielded collider in $\mathcal{G}^*$. The endpoints A and C are non-adjacent in $\mathcal{G}^*$ and therefore non-adjacent in $\mathcal{G}_1$. Moreover, the structurally excluded pair $\{W, Y_k\}$ cannot be the endpoint pair of a true unshielded collider: if the middle node is in $\mathbf{X}$ or $\mathbf{S}$, Assumption 3(i) prevents an arrow out of Y_k into that middle node, and if the middle node is in $\mathbf{Y}$, Assumption 3(ii) prevents the edge $W \to \mathbf{Y}$. Thus $SepSet_{AC}$ was learned in Phase 1. Since B is a collider on the length-two path $A - B - C$, no separating set for A and C can contain B; conditioning on B would open that path and there is no other node on the path that could block it. Hence $B \notin SepSet_{AC}$, so the collider-orientation step orients $A \to B \leftarrow C$ in $\mathcal{G}_2$. Therefore $\mathcal{G}_2$ and $\mathcal{G}^*$ have the same unshielded colliders. By the definition of consistent extension, every $\mathcal{G} \in \mathbf{G}$ has the same unshielded colliders as $\mathcal{G}_2$, and hence as $\mathcal{G}^*$.

Thus every $\mathcal{G} \in \mathbf{G}$ has the same skeleton and the same unshielded colliders as $\mathcal{G}^*$. By Lemma 2, every $\mathcal{G} \in \mathbf{G}$ is Markov equivalent to $\mathcal{G}^*$. Since we have also shown $\mathcal{G}^* \in \mathbf{G}$, the theorem follows. □

Proof of Proposition 1. Note that this proof utilizes the rules of do-calculus as outlined by Pearl (2000). For detailed information on these rules, please refer to Appendix D. We will use the following notation in this proof: $\mathcal{G}_{\overline{X}}$ denotes the graph obtained by deleting all incoming arrows to nodes in $\mathbf{X}$ from $\mathcal{G}$. Similarly, $\mathcal{G}_{\underline{\mathbf{X}}}$ denotes the graph obtained by deleting all outgoing arrows from nodes in X from $\mathcal{G}$. Fix $w \in \{0, 1\}$ and $Y_k \in \mathbf{Y}$. To simplify notation, write

$$\mathbf{S} = \mathbf{S}_{\mathcal{G},Y_k}, \qquad \mathbf{Z} = \mathbf{Z}_{\mathcal{G},Y_k}, \qquad \mathcal{S} = \mathcal{S}_{\mathcal{G},Y_k}, \qquad \mathcal{Z} = \mathcal{Z}_{\mathcal{G},Y_k}.$$

Also write $P_E(\cdot) = P(\cdot \mid D = E)$ and $P_O(\cdot) = P(\cdot \mid D = O)$. We state the proof for discrete variables; the continuous case follows by replacing sums with integrals.

By the law of total probability,

$$P_E(Y_k = y \mid \mathrm{do}(W = w)) = \sum_{\mathbf{x}} \sum_{\mathbf{z} \in \mathscr{Z}} \sum_{\mathbf{s} \in \mathscr{S}} P_E(Y_k = y \mid \mathrm{do}(W = w), \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}) \times P_E(\mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s} \mid \mathrm{do}(W = w), \mathbf{X} = \mathbf{x}) \times P_E(\mathbf{X} = \mathbf{x} \mid \mathrm{do}(W = w)).$$

By Assumption 3(i), the variables in $\mathbf{X}$ are pre-treatment relative to W. Hence intervening on W does not change their distribution:

$$P_E(\mathbf{X} = \mathbf{x} \mid \mathrm{do}(W = w)) = P_E(\mathbf{X} = \mathbf{x}).$$

Next, by Assumption 2, treatment is randomized conditional on $\mathbf{X}$ in the experimental sample. Therefore, for all cells satisfying the usual treatment-overlap condition,

$$P_E(\mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s} \mid \mathrm{do}(W = w), \mathbf{X} = \mathbf{x}) = P_E(\mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s} \mid W = w, \mathbf{X} = \mathbf{x}).$$

It remains to simplify the conditional distribution of Y_k. Since $\mathbf{Z}$ is a valid surrogate–outcome adjustment set for $(\mathbf{S}, Y_k)$, Definition 2 implies that every path from W to Y_k is blocked by

$$\mathbf{X} \cup \mathbf{Z} \cup \mathbf{S}.$$

Thus,

$$(Y_k \perp\!\!\!\perp W \mid \mathbf{X}, \mathbf{Z}, \mathbf{S})_{\mathcal{G}}.$$

Under Assumptions 4 and 6, this graphical separation implies

$$P_E(Y_k = y \mid W = w, \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}) = P_E(Y_k = y \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}).$$

Moreover, deleting outgoing arrows from W cannot open any path from W to Y_k. Hence the same conditioning set also separates W and Y_k in $\mathcal{G}_{\underline{W}}$. By the action–observation exchange rule of do-calculus,

$$P_E(Y_k = y \mid \mathrm{do}(W = w), \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}) = P_E(Y_k = y \mid W = w, \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}).$$

Combining the last two displays gives

$$P_E(Y_k = y \mid \mathrm{do}(W = w), \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}) = P_E(Y_k = y \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}).$$

Substituting these equalities into the law-of-total-probability expansion yields

$$P_E(Y_k = y \mid \mathrm{do}(W = w)) = \sum_{\mathbf{x}} \sum_{\mathbf{z} \in \mathscr{Z}} \sum_{\mathbf{s} \in \mathscr{S}} P_E(Y_k = y \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}) \times P_E(\mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s} \mid W = w, \mathbf{X} = \mathbf{x}) \times P_E(\mathbf{X} = \mathbf{x}).$$

Finally, by Assumption 8, for every $(\mathbf{x}, \mathbf{z}, \mathbf{s})$ that receives positive probability in the experimental sample under treatment arm w,

$$P_E(Y_k = y \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}) = P_O(Y_k = y \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{S} = \mathbf{s}).$$

The support component of Assumption 8 ensures that the right-hand side is defined for every such cell. Therefore,

$$\begin{aligned} P(Y_k = y \mid D = E, \operatorname{do}(W = w)) = \sum_{\mathbf{x}} \sum_{\substack{\mathbf{z} \in \mathcal{Z}_{\mathcal{G},Y_k} \\ \mathbf{s} \in \mathcal{S}_{\mathcal{G},Y_k}}} & P(Y_k = y \mid D = O, \mathbf{X} = \mathbf{x}, \mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s}) \\ & \times P(\mathbf{Z}_{\mathcal{G},Y_k} = \mathbf{z}, \mathbf{S}_{\mathcal{G},Y_k} = \mathbf{s} \mid D = E, W = w, \mathbf{X} = \mathbf{x}) \\ & \times P(\mathbf{X} = \mathbf{x} \mid D = E), \end{aligned}$$

which is the desired formula. □

Proof of Theorem 2. Fix $Y_k \in \mathbf{Y}$. Under Assumptions 1, 2, 3, 4, 5, 6, and 7, Theorem 1 implies that the true graph $\mathcal{G}^*$ is contained in the set of graphs $\mathbf{G}$ returned by the COMB-PC algorithm:

$$\mathcal{G}^* \in \mathbf{G}.$$

By Proposition 1, together with Assumption 8, the graph-specific quantity $\tau_{\mathcal{G}^*,Y_k}$ is the true average treatment effect of W on Y_k in the experimental population. Since $\mathcal{G}^* \in \mathbf{G}$, this true effect is one element of the candidate set

$$\left\{\tau_{\mathcal{G},Y_k} : \mathcal{G} \in \mathbf{G}\right\}.$$

Therefore, because the minimum and maximum of a finite set bound every element of that set,

$$\min_{\mathcal{G} \in \mathbf{G}} \tau_{\mathcal{G},Y_k} \leq \tau_{\mathcal{G}^*,Y_k} \leq \max_{\mathcal{G} \in \mathbf{G}} \tau_{\mathcal{G},Y_k}.$$

This proves the claim. □

Proof of Proposition 2. It suffices to construct one admissible instance. Let $\mathbf{X} = \emptyset$, let $\mathbf{S} = S_1, S_2$, and let $\mathbf{Y} = Y_k$. Consider two DAGs on W, S_1, S_2, Y_k. The first graph $\mathcal{G}$ has edges $W \to S_1$, $W \to S_2$, $S_1 \to S_2$, $S_1 \to Y_k$, and $S_2 \to Y_k$. The second graph $\mathcal{G}'$ has the same edges except that the edge between S_1 and S_2 is reversed: $W \to S_1$, $W \to S_2$, $S_2 \to S_1$, $S_1 \to Y_k$, and $S_2 \to Y_k$. The two graphs have the same skeleton. They also have the same unshielded colliders: the triples involving S_1 and S_2 are shielded because S_1 and S_2 are adjacent, and the remaining unshielded triples are noncolliders. Hence, by Lemma 2, $\mathcal{G}$ and $\mathcal{G}'$ are Markov equivalent.

Now take $\mathcal{G}$ to be the true graph $\mathcal{G}^*$. For example, let the true structural equations be $W = \varepsilon_W$, $S_1 = aW + \varepsilon_1$, $S_2 = bW + cS_1 + \varepsilon_2$, and $Y_k = dS_1 + eS_2 + \varepsilon_Y$, where the error terms are jointly independent and the coefficients are chosen generically with $c \neq 0$ and $e \neq 0$. This data-generating process satisfies the

required Markov, faithfulness, sufficiency, randomization, and structural ordering assumptions. Since $\mathcal{G}$ and $\mathcal{G}'$ are Markov equivalent, the conditional independence information available to COMB-PC does not distinguish the orientation of the edge between S_1 and S_2. Thus, under the oracle independence-test assumption, COMB-PC returns an equivalence class containing both $\mathcal{G}$ and $\mathcal{G}'$.

We now compare the gain from the same program modification under the two graphs. Consider a modification $r^h = (\mathcal{A}^h, \boldsymbol{\delta}^h)$ that targets only S_1, so that $\mathcal{A}^h = 1$ and $\delta_1^h > 0$. Under $\mathcal{G}$, there are two directed paths from S_1 to Y_k: $S_1 \to Y_k$ and $S_1 \to S_2 \to Y_k$. Therefore, $\Gamma_{S_1,Y_k,\mathcal{G}} = \beta^{\mathcal{G}}_{S_1,Y_k} + \beta^{\mathcal{G}}_{S_1,S_2}\beta^{\mathcal{G}}_{S_2,Y_k} = d + ce$. Hence the gain from the modification under $\mathcal{G}$ is $\Delta_{Y_k,\mathcal{G}}(r^h) = \delta_1^h(d + ce)$.

Under $\mathcal{G}'$, the edge between the two short-term variables is reversed: $S_2 \to S_1$. Thus S_1 no longer has a directed path to Y_k through S_2. The only directed path from S_1 to Y_k is $S_1 \to Y_k$. Since the parent set of Y_k is S_1, S_2 in both graphs, the coefficient on S_1 in the outcome equation is the same direct coefficient d. Therefore, $\Gamma_{S_1,Y_k,\mathcal{G}'} = \beta^{\mathcal{G}'} * S_1, Y_k = d$, and the gain under $\mathcal{G}'$ is $\Delta * Y_k, \mathcal{G}'(r^h) = \delta_1^h d$.

Since $\delta_1^h > 0$, $c \neq 0$, and $e \neq 0$, $\Delta_{Y_k,\mathcal{G}}(r^h) - \Delta_{Y_k,\mathcal{G}'}(r^h) = \delta_1^h ce \neq 0$. Thus the same program modification can imply different policy gains under two graphs returned by COMB-PC. This proves the claim. □

Proof of Theorem 3. By Theorem 1, the assumptions imply that the true graph $\mathcal{G}^*$ is contained in the set of graphs returned by COMB-PC, so $\mathcal{G}^* \in \mathbf{G}$. Fix any program modification r^h. Since $\mathcal{G}^* \in \mathbf{G}$, the true-graph gain $\Delta_{\mathcal{G}^*}(r^h)$ is one element of the collection $\Delta_{\mathcal{G}}(r^h) : \mathcal{G} \in \mathbf{G}$. Therefore, it must lie between the minimum and maximum values of that collection: $\min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h) \leq \Delta_{\mathcal{G}^*}(r^h) \leq \max_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h)$. This proves the first claim.

Now define $L_h := \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^h)$, and let $h^\star \in \arg\max_{h\in 0,1,\ldots,H} L_h$. Because the status quo option r^0 is included in the menu and satisfies $\Delta_{\mathcal{G}}(r^0) = 0$ for every $\mathcal{G} \in \mathbf{G}$, we have $L_0 = \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^0) = 0$. Since $h^\star$ maximizes L_h over a set that includes $h = 0$, it follows that $L_{h^\star} \geq L_0 = 0$. Equivalently, $\min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^{h^\star}) \geq 0$. Finally, applying the first part of the theorem to the selected modification $r^{h^\star}$ gives $\Delta_{\mathcal{G}^*}(r^{h^\star}) \geq \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^{h^\star})$. Combining the last two inequalities yields $\Delta_{\mathcal{G}^*}(r^{h^\star}) \geq \min_{\mathcal{G}\in\mathbf{G}} \Delta_{\mathcal{G}}(r^{h^\star}) \geq 0$. □