TECHNICAL REPORT
R-35
August 2018

Causal Identification under Markov Equivalence

Amin Jaber
Computer Science Department
Purdue University, IN, USA
jaber0@purdue.edu

Jiji Zhang
Philosophy Department
Lingnan University, NT, HK
jijizhang@ln.edu.hk

Elias Bareinboim
Computer Science Department
Purdue University, IN, USA
eb@purdue.edu

Abstract

Assessing the magnitude of cause-and-effect relations is one of the central challenges found throughout the empirical sciences. The problem of identification of causal effects is concerned with determining whether a causal effect can be computed from a combination of observational data and substantive knowledge about the domain under investigation, which is formally expressed in the form of a causal graph. In many practical settings, however, the knowledge available for the researcher is not strong enough so as to specify a unique causal graph. Another line of investigation attempts to use observational data to learn a qualitative description of the domain called a Markov equivalence class, which is the collection of causal graphs that share the same set of observed features. In this paper, we marry both approaches and study the problem of causal identification from an equivalence class, represented by a partial ancestral graph (PAG). We start by deriving a set of graphical properties of PAGs that are carried over to its induced subgraphs. We then develop an algorithm to compute the effect of an arbitrary set of variables on an arbitrary outcome set. We show that the algorithm is strictly more powerful than the current state of the art found in the literature.

1 INTRODUCTION

Science is about explaining the mechanisms underlying a phenomenon that is being investigated. One of the marks imprinted by these mechanisms in reality is cause and effect relationships. Systematically discovering the existence, and magnitude, of causal relations constitutes, therefore, a central task in scientific domains. The value of inferring causal relationships is also tremendous in other, more practical domains, including, for example, engineering and business, where it is often crucial to understand how to bring about a specific change when a constrained amount of controllability is in place. If our goal is to build AI systems that can act and learn autonomously, formalizing the principles behind causal inference, so that these systems can leverage them, is a fundamental requirement (Pearl and Mackenzie, 2018).

One prominent approach to infer causal relations leverages a combination of substantive knowledge about the domain under investigation, usually encoded in the form of a causal graph, with observational (non-experimental) data (Pearl, 2000; Spirtes et al., 2001; Bareinboim and Pearl, 2016). A sample causal graph is shown in Fig. 1a such that the nodes represent variables, directed edges represent direct causal relation from tails to heads, and bi-directed arcs represent the presence of unobserved (latent) variables that generate a spurious association between the variables, also known as *confounding bias* (Pearl, 1993). The task of determining whether an interventional (experimental) distribution can be computed from a combination of observational and experimental data together with the causal graph is known as the problem of identification of causal effects (identification, for short). For instance, a possible task in this case is to identify the effect of $do(X{=}x)$ on $V_4{=}v_4$, i.e. $P_x(v_4)$, given the causal graph in Fig. 1a and data from the observational distribution $P(x, v_1, ..., v_4)$.

The problem of identification has been extensively studied in the literature, and a number of criteria have been established (Pearl, 1993; Galles and Pearl, 1995; Kuroki and Miyakawa, 1999; Tian and Pearl, 2002; Huang and Valtorta, 2006; Shpitser and Pearl, 2006; Bareinboim and Pearl, 2012), which include the celebrated back-door criterion and the do-calculus (Pearl, 1995). Despite their power, these techniques require a fully specified causal graph, which is not always available in practical settings.

Another line of investigation attempts to learn a qualitative description of the system, which in the ideal case would lead to the "true" data-generating model, the blueprint underlying the phenomenon being investigated. These efforts could certainly be deemed more "data-driven" and aligned with the zeitgeist in machine learning. In practice, however, it is common that only an equivalence class of causal models can be consistently inferred from observational data (Verma, 1993; Spirtes et al., 2001; Zhang, 2008b). One useful characterization of such an equivalence class comes under the rubric of *partial ancestral graphs (PAGs)*, which will be critical to our work. Fig. 1 shows the PAG (right) that can be inferred from observational data that is consistent with the true causal model (left). The directed edges in a PAG signify ancestral relations (not necessarily direct) and circle marks stand for structural uncertainty.

In this paper, we analyze the marriage of these two lines of investigation, where the structural invariance learned in the equivalence class will be used as input to identify the strength of causal effect relationships, if possible. Identification from an equivalence class is considerably more challenging than from a single diagram due to the structural uncertainty regarding both the direct causal relations among the variables and the presence of latent variables that confounds causal relations between observed variables. Still, there is a growing interest in identifiability results in this setting (Maathuis et al., 2010). Zhang (2007) extended the do-calculus to PAGs. In practice, however, it is in general computationally hard to decide whether there exists (and, if so, find) a sequence of applications of the rules of the generalized calculus to identify the interventional distribution. Perković et al. (2015) generalized the back-door criterion to PAGs, and provided a sound and complete algorithm to find a back-door admissible set, should such a set exist. However, in practice, the back-door criterion is not as powerful as the do-calculus, since no adjustment set exists for many identifiable causal effects. Jaber et al. (2018b) generalized the work of (Tian and Pearl, 2002) and devised a graphical criterion to identify causal effects with singleton interventions in PAGs.[1]

Building on this work, we develop here a decomposition strategy akin to the one introduced in (Tian, 2002) to identify causal effects given a PAG. Our proposed approach is computationally more attractive than the do-calculus as it provides a systematic procedure to identify

[1] Another possible approach is based on SAT (boolean constraint satisfaction) solvers (Hyttinen et al., 2015). Given its somewhat distinct nature, a closer comparison lies outside the scope of this paper. We note, however, that an open research direction would be to translate our systematic approach into logical rules so as to help improving the solver's scalability.

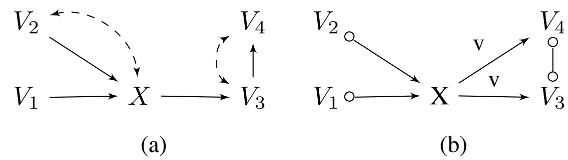

Figure 1: A causal model (left) and the inferred PAG (right).

a causal effect, if identifiable. It is also more powerful than the generalized adjustment criterion, as we show later. More specifically, our main contributions are:

1. We study some critical properties of PAGs and show that they also hold in induced subgraphs of a PAG over an arbitrary subset of nodes. We further study Tian's c-component decomposition and relax it to PAGs (when only partial knowledge about the ancestral relations and c-components is available).
2. We formulate a systematic procedure to compute the effect of an arbitrary set of intervention variables on an arbitrary outcome set from a PAG and observational data. We show that this algorithm is strictly more powerful than the adjustment criterion.

2 PRELIMINARIES

In this section, we introduce the basic notation and machinery used throughout the paper. Bold capital letters denote sets of variables, while bold lowercase letters stand for particular assignments to those variables.

Structural Causal Models. We use the language of Structural Causal Models (SCM) (Pearl, 2000, pp. 204-207) as our basic semantic framework. Formally, an SCM M is a 4-tuple $\langle U, V, F, P(u)\rangle$, where U is a set of exogenous (latent) variables and V is a set of endogenous (measured) variables. F represents a collection of functions $F = \{f_i\}$ such that each endogenous variable $V_i \in V$ is determined by a function $f_i \in F$, where f_i is a mapping from the respective domain of $U_i \cup Pa_i$ to V_i, $U_i \subseteq U$, $Pa_i \subseteq V \setminus V_i$. The uncertainty is encoded through a probability distribution over the exogenous variables, $P(u)$. A causal diagram associated with an SCM encodes the structural relations among $V \cup U$, in which an arrow is drawn from each member of $U_i \cup Pa_i$ to V_i. We constraint our results to recursive systems, which means that the corresponding diagram will be acyclic. The marginal distribution over the endogenous variables $P(v)$ is called observational, and factorizes according to the causal diagram, i.e.:

$$P(v) = \sum_u \prod_i P(v_i|pa_i, u_i)P(u)$$

Within the structural semantics, performing an action $X{=}x$ is represented through the do-operator, *do*$(X=x)$, which encodes the operation of replacing the original equation for X by the constant x and induces a submodel M_x. The resulting distribution is denoted by P_x, which is the main target for identification in this paper. For details on structural models, we refer readers to (Pearl, 2000).

Ancestral Graphs. We now introduce a graphical representation of equivalence classes of causal diagrams. A *mixed* graph can contain directed ($\rightarrow$) and bi-directed edges ($\leftrightarrow$). A is a *spouse* of B if $A \leftrightarrow B$ is present. An *almost directed cycle* happens when A is both a spouse and an ancestor of B. An *inducing path relative to* $\mathbf{L}$ is a path on which every node $V \notin \mathbf{L}$ (except for the endpoints) is a collider on the path (i.e., both edges incident to V are into V) and every collider is an ancestor of an endpoint of the path. A mixed graph is *ancestral* if it doesn't contain a directed or almost directed cycle. It is *maximal* if there is no inducing path (relative to the empty set) between any two non-adjacent nodes. A *Maximal Ancestral Graph* (MAG) is a graph that is both ancestral and maximal. MAG models are closed under marginalization (Richardson and Spirtes, 2002).

In general, a causal MAG represents a set of causal models with the same set of observed variables that entail the same independence and ancestral relations among the observed variables. Different MAGs may be Markov equivalent in that they entail the exact same independence model. A partial ancestral graph (PAG) represents an equivalence class of MAGs $[\mathcal{M}]$, which shares the same adjacencies as every MAG in $[\mathcal{M}]$ and displays all and only the invariant edge marks.

Definition 1 (PAG)**.** *Let $[\mathcal{M}]$ be the Markov equivalence class of an arbitrary MAG $\mathcal{M}$. The PAG for $[\mathcal{M}]$, $\mathcal{P}$, is a partial mixed graph such that:*

i. $\mathcal{P}$ has the same adjacencies as $\mathcal{M}$ (and hence any member of $[\mathcal{M}]$) does.

ii. An arrowhead is in $\mathcal{P}$ iff shared by all MAGs in $[\mathcal{M}]$.

iii. A tail is in $\mathcal{P}$ iff shared by all MAGs in $[\mathcal{M}]$.

iv. A mark that is neither an arrowhead nor a tail is recorded as a circle.

A PAG is learnable from the conditional independence and dependence relations among the observed variables and the FCI algorithm is a standard method to learn such an object (Zhang, 2008b). In short, a PAG represents an equivalence class of causal models with the same observed variables and independence model.

Graphical Notions. Given a DAG, MAG, or PAG, a path between X and Y is *potentially directed (causal)* from X to Y if there is no arrowhead on the path pointing towards X. Y is called a *possible descendant* of X and X a *possible ancestor* of Y if there is a potentially directed path from X to Y. A set $\mathbf{A}$ is (*descendant*) *ancestral* if no node outside $\mathbf{A}$ is a possible (descendant) ancestor of any node in $\mathbf{A}$. Y is called a *possible child* of X, i.e. $Y \in \texttt{Ch}(X)$, and X a *possible parent* of Y, i.e. $X \in \texttt{Pa}(Y)$, if they are adjacent and the edge is not into X. For a set of nodes $\mathbf{X}$, we have $\texttt{Pa}(\mathbf{X}) = \cup_{X\in\mathbf{X}}\texttt{Pa}(X)$ and $\texttt{Ch}(\mathbf{X}) = \cup_{X\in\mathbf{X}}\texttt{Ch}(X)$. Given two sets of nodes $\mathbf{X}$ and $\mathbf{Y}$, a path between them is called *proper* if one of the endpoints is in $\mathbf{X}$ and the other is in $\mathbf{Y}$, and no other node on the path is in $\mathbf{X}$ or $\mathbf{Y}$. For convenience, we use an asterisk (*) to denote any possible mark of a PAG ($\circ, >, -$) or a MAG ($>, -$). If the edge marks on a path between X and Y are all circles, we call the path a *circle path*.

A directed edge $X \rightarrow Y$ in a MAG or PAG is *visible* if there exists no DAG $\mathcal{D}(\mathbf{V}, \mathbf{L})$ in the corresponding equivalence class where there is an inducing path between X and Y that is into X relative to $\mathbf{L}$. This implies that a visible edge is not confounded ($X \leftarrow U_i \rightarrow Y$ doesn't exist). Which directed edges are visible is easily decidable by a graphical condition (Zhang, 2008a), so we simply mark visible edges by v. For brevity, we refer to any edge that is not a visible directed edge as *invisible*.

Identification Given a Causal DAG. Tian and Pearl (2002) presented an identification algorithm based on a decomposition strategy of the DAG into a set of so-called *c-components* (confounded components).

Definition 2 (C-Component)**.** *In a causal DAG, two observed variables are said to be in the same c-component if and only if they are connected by a bi-directed path, i.e. a path composed solely of such bi-directed treks as $V_i \leftarrow U_{ij} \rightarrow V_j$, where U_{ij} is an exogenous variable.*

For convenience, we often refer to a bi-directed trek like $V_i \leftarrow U_{ij} \rightarrow V_j$ as a bi-directed edge between V_i and V_j (and U_{ij} is often left implicit). For any set $\mathbf{C} \subseteq \mathbf{V}$, we define the quantity $Q[\mathbf{C}]$ to denote the post-intervention distribution of $\mathbf{C}$ under an intervention on $\mathbf{V} \setminus \mathbf{C}$:

$$Q[\mathbf{C}] = P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{c}) = \sum_u \prod_{\{i|V_i\in\mathbf{C}\}} P(v_i|pa_i, u_i)P(u)$$

The significance of c-components and their decomposition is evident from (Tian, 2002, Lemmas 10, 11), which are the basis of Tian's identification algorithm.

3 REVISIT IDENTIFICATION IN DAGS

We revisit the identification results in DAGs, focusing on Tian's algorithm (Tian, 2002). Our goal here is to have an amenable algorithm that allows the incorporation of the structural uncertainties arising in the equivalence class.

Let $\mathcal{D}_{\mathbf{A}}$ denote the (induced) subgraph of a DAG $\mathcal{D}(\mathbf{V},\mathbf{L})$ over $\mathbf{A} \subseteq \mathbf{V}$ and the latent parents of $\mathbf{A}$ (i.e. $\texttt{Pa}(\mathbf{A}) \cap \mathbf{L}$). The original algorithm (Alg. 5 in (Tian, 2002)) alternately applies Lemmas 10 and 11 in (Tian, 2002) until a solution is derived or a failure condition is triggered. We rewrite this algorithm with a more local, atomic criterion based on the following results.

Definition 3 (Composite C-Component)**.** *Given a DAG that decomposes into c-components $S_1, \dots, S_k$, $k \geq 1$, a composite c-component is the union of one or more of these c-components.*

Lemma 1. *Given a DAG $\mathcal{D}(\mathbf{V},\mathbf{L})$, $\mathbf{X} \subset \mathbf{T} \subseteq \mathbf{V}$, and $P_{\mathbf{v}\setminus\mathbf{t}}$ the interventional distribution of $\mathbf{V} \setminus \mathbf{T}$ on $\mathbf{T}$. Let $S^{\mathbf{X}}$ denote a composite c-component containing $\mathbf{X}$ in $\mathcal{D}_{\mathbf{T}}$. If $\mathbf{X}$ is a descendant set in $\mathcal{D}_{S^{\mathbf{X}}}$, then $Q[\mathbf{T} \setminus \mathbf{X}]$ is identifiable and given by*

$$Q[\mathbf{T} \setminus \mathbf{X}] = \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{Q[S^{\mathbf{X}}]} \times \sum_{\mathbf{x}} Q[S^{\mathbf{X}}] \tag{1}$$

Proof. By (Tian, 2002, Lemma 11), $Q[\mathbf{T}]$ decomposes as follows.

$$Q[\mathbf{T}] = Q[\mathbf{T} \setminus S^{\mathbf{X}}] \times Q[S^{\mathbf{X}}] = \frac{Q[\mathbf{T}]}{Q[S^{\mathbf{X}}]} \times Q[S^{\mathbf{X}}]$$

$Q[S^{\mathbf{X}}]$ is computable from $P_{\mathbf{v}\setminus\mathbf{t}}$ using Lemma 11 in (Tian, 2002), and $Q[S^{\mathbf{X}} \setminus \mathbf{X}]$ is computable from $Q[S^{\mathbf{X}}]$ using (Tian, 2002, Lemma 10) as $\mathbf{X}$ is a descendant set in $\mathcal{D}_{S^{\mathbf{X}}}$. Therefore,

$$Q[\mathbf{T} \setminus \mathbf{X}] = \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{Q[S^{\mathbf{X}}]} \cdot Q[S^{\mathbf{X}} \setminus \mathbf{X}] = \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{Q[S^{\mathbf{X}}]} \cdot \sum_{\mathbf{x}} Q[S^{\mathbf{X}}]$$

□

The next result follows directly when $\mathbf{X}$ is a singleton.

Corollary 1. *Given a DAG $\mathcal{D}(\mathbf{V},\mathbf{L})$, $X \in \mathbf{T} \subseteq \mathbf{V}$, and $P_{\mathbf{v}\setminus\mathbf{t}}$. If X is not in the same c-component with a child in $\mathcal{D}_{\mathbf{T}}$, then $Q[\mathbf{T} \setminus \{X\}]$ is identifiable and given by*

$$Q[\mathbf{T} \setminus \{X\}] = \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{Q[S^{X}]} \times \sum_{x} Q[S^{X}] \tag{2}$$

where S^X is the c-component of X in $\mathcal{D}_{\mathbf{T}}$.

The significance of Corol. 1 stems from the fact that it can be used to rewrite the identification algorithm in a step-wise fashion, which is shown in Algorithm 1. The same is equivalent to the original algorithm since neither one of Lemmas 10 nor 11 in (Tian, 2002) is applicable whenever Corol. 1 is not applicable, which is shown by Lemmas 2 and 3. This result may not be surprising since Corol. 1 follows from the application of these lemmas.

Algorithm 1: ID$(\mathbf{x},\mathbf{y})$ given DAG $\mathcal{G}$

input : two disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$
output: Expression for $P_{\mathbf{x}}(\mathbf{y})$ or `FAIL`

1. Let $\mathbf{D} = \texttt{An}(\mathbf{Y})_{\mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}}$
2. Let the c-components of $\mathcal{G}_{\mathbf{D}}$ be $\mathbf{D}_i$, $i = 1, \dots, k$
3. $P_{\mathbf{x}}(\mathbf{y}) = \sum_{\mathbf{d}\setminus\mathbf{y}} \prod_i \texttt{Identify}(\mathbf{D}_i, \mathbf{V}, P)$

Function `Identify`$(\mathbf{C}, \mathbf{T}, Q = Q[\mathbf{T}])$**:**

```
if C = T then
    return Q[T];
end
/* Let S^B be the c-component of {B} in G_T */
if ∃B ∈ T \ C such that S^B ∩ Ch(B) = ∅ then
    Compute Q[T \ {B}] from Q; // Corollary 1
    return Identify(C, T \ {B}, Q[T \ {B}]);
else
    throw FAIL;
end
```

Lemma 2. *Given a DAG $\mathcal{D}(\mathbf{V},\mathbf{L})$, $\mathbf{C} \subset \mathbf{T} \subseteq \mathbf{V}$. If $\mathbf{A} = An(\mathbf{C})_{\mathcal{D}_{\mathbf{T}}} \neq \mathbf{T}$, then there exist some node $X \in \mathbf{T} \setminus \mathbf{A}$ such that X is not in the same c-component with any child in $\mathcal{D}_{\mathbf{T}}$.*

Proof. If $\mathbf{A} \neq \mathbf{T}$, then $\mathbf{T} \setminus \mathbf{A}$ is a non-empty set where none of the nodes is an ancestor of $\mathbf{A}$. Since the graph is acyclic, then at least one node of $\mathbf{T} \setminus \mathbf{A}$ is with no children. Hence, the above conclusion follows. □

Lemma 3. *Given a DAG $\mathcal{D}(\mathbf{V},\mathbf{L})$, $\mathbf{C} \subset \mathbf{T} \subseteq \mathbf{V}$, and assume $\mathcal{D}_{\mathbf{C}}$ is a single c-component. If $\mathcal{D}_{\mathbf{T}}$ partitions into c-components $\mathbf{S}_1 \dots \mathbf{S}_k$, where $k > 1$, then there exists some node $X \in \mathbf{S}_i$ such that $\mathbf{C} \not\subseteq \mathbf{S}_i$ and X is not in the same c-component with any child in $\mathcal{D}_{\mathbf{T}}$.*

Proof. Subgraph $\mathcal{D}_{\mathbf{S}_i}$ is acyclic, so there must exist some node (X) that doesn't have any children in $\mathcal{D}_{\mathbf{S}_i}$. Since $\mathbf{S}_i$ is one of the c-components in $\mathcal{D}_{\mathbf{T}}$, then X is not in the same c-component with any of its children in $\mathcal{D}_{\mathbf{T}}$. □

The revised algorithm requires checking an atomic criterion at every instance of the recursive routine `Identify`. This might not be crucial when the precise causal diagram is known and induced subgraphs preserve a complete graphical characterization of the c-components and the ancestral relations between the nodes. The latter, unfortunately, doesn't hold when the model is an equivalence class represented by a PAG.[2]

[2]We thank a reviewer for bringing to our attention a similar formulation of Alg. 1 (Richardson et al., 2017, Thm. 60).

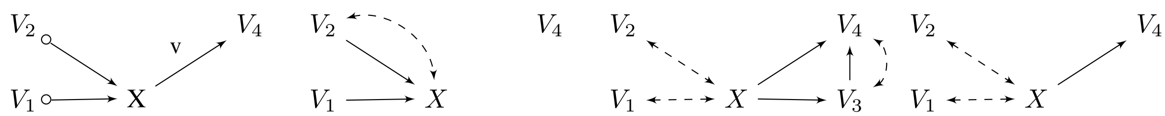

(a) Subgraph of Fig. 1b. (b) Subgraph of Fig. 1a. (c) Equivalent DAG to Fig. 1a. (d) Subgraph of Fig. 2c

Figure 2: Example for properties discussed in Section 4

4 PAG-SUBGRAPH PROPERTIES

Evidently, induced subgraphs of the original causal model play a critical role in identification (cf Alg. 1). It is natural to expect that in the generalized setting we study here, induced subgraphs of the given PAG will also play an important role. An immediate challenge, however, is that a subgraph of a PAG $\mathcal{P}$ over $\mathbf{V}$ induced by $\mathbf{A} \subseteq \mathbf{V}$ is, in general, not a PAG that represents a full Markov equivalence class. In particular, if $\mathcal{D}(\mathbf{V}, \mathbf{L})$ is a DAG in the equivalence class represented by $\mathcal{P}$, $\mathcal{P}_{\mathbf{A}}$ is in general not the PAG that represents the equivalence class of $\mathcal{D}_{\mathbf{A}}$. To witness, let $\mathcal{D}$ and $\mathcal{P}$ denote the DAG and the corresponding PAG in Figure 1, respectively, and let $\mathbf{A} = \{V_1, V_2, X, V_4\}$. The induced subgraph of $\mathcal{P}$ over $\mathbf{A}$ (Fig. 2a) does not represent the equivalence class of the corresponding induced subgraph of $\mathcal{D}$ (Fig. 2b). Despite this subtlety, we establish a few facts below showing that for any $\mathbf{A} \subseteq \mathbf{V}$ and any DAG $\mathcal{D}$ in the equivalence class represented by $\mathcal{P}$, some information about $\mathcal{D}_{\mathbf{A}}$, which is particularly relevant to identification, can be read off from $\mathcal{P}_{\mathbf{A}}$.

Proposition 1. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{D}(\mathbf{V}, \mathbf{L})$ be any DAG in the equivalence class represented by $\mathcal{P}$. Let $X \neq Y$ be two nodes in $\mathbf{A} \subseteq \mathbf{V}$. If X is an ancestor of Y in $\mathcal{D}_{\mathbf{A}}$, then X is a possible ancestor of Y in $\mathcal{P}_{\mathbf{A}}$.*

Proof. If X is an ancestor of Y in $\mathcal{D}_{\mathbf{A}}$, then there is a path p in $\mathcal{D}_{\mathbf{A}}$ composed of nodes $\langle X = V_0, \dots, Y = V_m \rangle$, $m \geq 1$ such that $V_i \in \mathbf{A}$ and $V_i \rightarrow V_{i+1}$, $0 \leq i < m$. Path p is obviously also present in $\mathcal{D}$, and consequently the corresponding MAG $\mathcal{M}$. Hence, p corresponds to a possibly directed path in $\mathcal{P}$. Since all the nodes along p are in $\mathbf{A}$, then p is present in $\mathcal{P}_{\mathbf{A}}$ and so X is a possible ancestor of Y in $\mathcal{P}_{\mathbf{A}}$. □

This simple proposition guarantees that possible-ancestral relationship in $\mathcal{P}_{\mathbf{A}}$ subsumes ancestral relationship in $\mathcal{D}_{\mathbf{A}}$ for every $\mathcal{D}$ in the class represented by $\mathcal{P}$. This is illustrated by $\mathcal{D}_{\mathbf{A}}$ and $\mathcal{P}_{\mathbf{A}}$ in Figures 2a and 2b.

Given an induced subgraph of a PAG, $\mathcal{P}_{\mathbf{A}}$, a directed edge $X \rightarrow Y$ in $\mathcal{P}_{\mathbf{A}}$ is said to be *visible* if for every DAG $\mathcal{D}$ in the class represented by $\mathcal{P}$, there is no inducing path in $\mathcal{D}_{\mathbf{A}}$ between X and Y relative to the latent nodes in $\mathcal{D}_{\mathbf{A}}$ that is into X.

Lemma 4. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{P}_{\mathbf{A}}$ be an induced subgraph of $\mathcal{P}$ over $\mathbf{A} \subseteq \mathbf{V}$. For every $X \rightarrow Y$ in $\mathcal{P}_{\mathbf{A}}$, if it is visible in $\mathcal{P}$, then it remains visible in $\mathcal{P}_{\mathbf{A}}$.*

Proof. Let $\mathcal{D}(\mathbf{V}, \mathbf{L})$ be any causal model in the equivalence class represented by $\mathcal{P}$, and let $X \rightarrow Y$ be a visible edge in $\mathcal{P}$, $X, Y \in \mathbf{A}$. Then, there is no inducing path between X and Y relative to $\mathbf{L}$ that is into X in $\mathcal{D}$. It follows that no such inducing path (relative to the latent nodes in $\mathcal{D}_{\mathbf{A}}$) exists in the subgraph $\mathcal{D}_{\mathbf{A}}$. □

Visibility is relevant for identification because it implies absence of confounding, which is the major obstacle to identification. Lemma 4 shows that an edge in an induced subgraph that is visible in the original PAG also implies absence of confounding in the induced subgraphs. Interestingly, note that a directed edge $X \rightarrow Y$ in $\mathcal{P}_{\mathbf{A}}$, visible or not, does not imply that X is an ancestor of Y in $\mathcal{D}_{\mathbf{A}}$ for every $\mathcal{D}$ in the class represented by $\mathcal{P}$. For example, X is not an ancestor of V_4 in Fig. 2b, even though $X \rightarrow V_4$ is a visible edge in Fig. 2a.

Definition 4 (PC-Component)**.** *In a MAG, a PAG, or any of its induced subgraphs, two nodes X and Y are in the same possible c-component (pc-component) if there is a path between the two nodes such that (1) all non-endpoint nodes along the path are colliders, and (2) none of the edges is visible.*

As alluded earlier, a c-component in a causal graph plays a central role in identification. The following proposition establishes a graphical condition in an induced subgraph $\mathcal{P}_{\mathbf{A}}$ that is necessary for two nodes being in the same c-component in $\mathcal{D}_{\mathbf{A}}$ for some DAG $\mathcal{D}$ represented by $\mathcal{P}$.

Proposition 2. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{D}(\mathbf{V}, \mathbf{L})$ be any DAG in the equivalence class represented by $\mathcal{P}$. For any $X, Y \in \mathbf{A} \subseteq \mathbf{V}$, if X and Y are in the same c-component in $\mathcal{D}_{\mathbf{A}}$, then X and Y are in the same pc-component in $\mathcal{P}_{\mathbf{A}}$.*

Proof Sketch. If X and Y are in the same c-component in $\mathcal{D}_{\mathbf{A}}$, then there is a path p in $\mathcal{D}_{\mathbf{A}}$ composed of nodes $\langle X = V_0, \dots, Y = V_m \rangle$, $m \geq 1$, such that $V_i \in \mathbf{A}$ and $V_i \leftarrow L_{i,i+1} \rightarrow V_{i+1}$, $0 \leq i < m$. We prove that X and Y are in the same pc-component in $\mathcal{M}$, the MAG of $\mathcal{D}$ over $\mathbf{V}$, due to a path p' over a subsequence of p. We then show that X and Y are in the same pc-component in $\mathcal{P}$, the PAG of $\mathcal{M}$, due to a path p^* over a subsequence of p'. Since all the nodes along p^* are in $\mathbf{A}$, then p^* is present in $\mathcal{P}_{\mathbf{A}}$, and so X and Y are in the same pc-component in $\mathcal{P}_{\mathbf{A}}$. Due to space constraints, the complete proofs are provided in (Jaber et al., 2018a). $\square$

This result provides a sufficient condition for *not* belonging to the same c-component in any of the relevant causal graphs. In Fig. 2a, for example, V_1 and V_4 or X and V_4 are not in the same pc-component, which implies by Prop. 2 that they are not in the same c-component in $\mathcal{D}_{\mathbf{A}}$ for any DAG $\mathcal{D}$ in the equivalence class represented by the PAG in Fig. 1b.

As a special case of Def. 4, we define the following notion, which will prove useful later on.

Definition 5 (DC-Component)**.** *In a MAG, a PAG, or any of its induced subgraphs, two nodes X and Y are in the same definite c-component (dc-component) if they are connected with a bi-directed path, i.e. a path composed solely of bi-directed edges.*

One challenge with the notion of *pc-component* is that it is not transitive as *c-component* is. Consider the PAG $V_1 \circ\!-\!\circ V_2 \circ\!-\!\circ V_3$. Here, V_1 and V_2 are in the same pc-component, V_2 and V_3 are in the same pc-component, however, V_1 and V_3 are not in the same pc-component. Hence, we define a notion that is a transitive closure of the notion of pc-component, which will prove instrumental to our goal.

Definition 6 (CPC-Component)**.** *Let $\mathcal{P}$ denote a PAG or a corresponding induced subgraph. Nodes X and Y are in the same composite pc-component in $\mathcal{P}$, denoted cpc-component, if there exist a sequence of nodes $\langle X = V_0, \ldots, Y = V_m \rangle$, $m \geq 1$, such that V_i and V_{i+1} are in the same pc-component, $0 \leq i < m$.*

It follows from the above definition that a PAG or an induced subgraph $\mathcal{P}$ can be decomposed into unique sets of cpc-components. For instance, the cpc-components in Fig. 2a are $\mathbf{S}_1 = \{V_1, V_2, X\}$ and $\mathbf{S}_2 = \{V_4\}$. The significance of a *cpc-component* is that it corresponds to a composite c-component in the relevant causal graphs as shown in the following proposition.

Proposition 3. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, $\mathcal{D}(\mathbf{V}, \mathbf{L})$ be any DAG in the equivalence class represented by $\mathcal{P}$, and $\mathbf{A} \subseteq \mathbf{V}$. If $\mathbf{C} \subseteq \mathbf{A}$ is a cpc-component in $\mathcal{P}_{\mathbf{A}}$, then $\mathbf{C}$ is a composite c-component in $\mathcal{D}_{\mathbf{A}}$.*

Proof. According to Definition 6, $\mathbf{C}$ includes all the nodes that are in the same pc-component with some node in $\mathbf{C}$ in $\mathcal{P}_{\mathbf{A}}$. If follows from the contrapositive of Prop. 2 that no node outside $\mathbf{C}$ is in the same c-component with any node in $\mathbf{C}$ in $\mathcal{D}_{\mathbf{A}}$. Hence, set $\mathbf{C}$ represents a composite c-component in $\mathcal{D}_{\mathbf{A}}$ by Definition 3. $\square$

Algorithm 2: PTO Algorithm

---

input : PAG $\mathcal{P}$ over $\mathbf{V}$
output: PTO over $\mathcal{P}$

1- Create singleton buckets $\mathbf{B_i}$ each containing $V_i \in \mathbf{V}$.

2- Merge buckets $\mathbf{B_i}$ and $\mathbf{B_j}$ if there is a circle edge between them ($\mathbf{B_i} \ni X \circ\!-\!\circ Y \in \mathbf{B_j}$).

3- **while** *set of buckets ($\mathbf{B}$) is not empty* **do**
  (i) Extract $\mathbf{B_i}$ with only arrowheads incident on it.
  (ii) Remove edges between $\mathbf{B_i}$ and other buckets.
end

4- The partial order is $\mathbf{B_1} < \mathbf{B_2} < \cdots < \mathbf{B_m}$ in reverse order of the bucket extraction. Hence, $\mathbf{B_1}$ is the last bucket extracted and $\mathbf{B_m}$ is the first bucket extracted.

Recall that the algorithm for identification given a DAG uses a topological order over the nodes. Similarly, the algorithm we design for PAGs will depend on some (partial) topological order. Thanks to the possible presence of circle edges ($\circ\!-\!\circ$) in a PAG, in general, there may be no complete topological order that is valid for all DAGs in the equivalence class. Algorithm 2 presents a procedure to derive a *partial topological order* over the nodes in a PAG, using buckets of nodes that are connected with circle paths (Jaber et al., 2018b). This algorithm remains valid over an induced subgraph of a PAG. To show this, the following lemma is crucial:

Lemma 5. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{P}_{\mathbf{A}}$ be the induced subgraph over $\mathbf{A} \subseteq \mathbf{V}$. For any three nodes A, B, C, if $A {*\!\!\rightarrow} B \circ\!\!-\!* C$, then there is an edge between A and C with an arrowhead at C, namely, $A {*\!\!\rightarrow} C$. Furthermore, if the edge between A and B is $A \rightarrow B$, then the edge between A and C is either $A \rightarrow C$ or $A \circ\!\!\rightarrow C$ (i.e., it is not $A \leftrightarrow C$).*

Proof. Lemma 3.3.1 of (Zhang, 2006) establishes the above property for every PAG. By the definition of an induced subgraph, the property is preserved in $\mathcal{P}_{\mathbf{A}}$. $\square$

Thus, a characteristic feature of PAGs carries over to their induced subgraphs. It follows that Algorithm 2 is sound for induced subgraphs as well.

Proposition 4. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and let $\mathcal{P}_{\mathbf{A}}$ be the subgraph of $\mathcal{P}$ induced by $\mathbf{A} \subseteq \mathbf{V}$. Then, Algorithm 2 is sound over $\mathcal{P}_{\mathbf{A}}$, in the sense that the partial order is valid with respect to $\mathcal{D}_{\mathbf{A}}$, for every DAG $\mathcal{D}$ in the equivalence class represented by $\mathcal{P}$.*

Proof. Let D be any DAG in the equivalence class represented by $\mathcal{P}$. By Prop. 1, the possible-ancestral relations in $\mathcal{P}_\mathbf{A}$ subsume those present in $\mathcal{D}_\mathbf{A}$. Hence, a partial topological order that is valid with respect to $\mathcal{P}_\mathbf{A}$ is valid with respect to $\mathcal{D}_\mathbf{A}$. The correctness of Alg. 2 with respect to a PAG in (Jaber et al., 2018b) depends only on the property in Lemma 5, a proof of which is given in the Supplementary Materials for completeness. Therefore, thanks to Lemma 5, the algorithm is also sound with respect to an induced subgraph $\mathcal{P}_\mathbf{A}$. □

For example, for $\mathcal{P}_\mathbf{A}$ in Fig. 2a, a partial topological order over the nodes is $V_1 < V_2 < X < V_4$, which is valid for all the relevant DAGs.

With these results about induced subgraphs of a PAG, we are ready to develop a recursive approach for identification given a PAG, to which we now turn.

5 IDENTIFICATION IN PAGS

We start by formally defining the notion of identification given a PAG, which generalizes the model-specific notion (Pearl, 2000, pp. 70).

Definition 7. *Given a PAG $\mathcal{P}$ over $\mathbf{V}$ and a query $P_\mathbf{x}(\mathbf{y})$ where $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, $P_\mathbf{x}(\mathbf{y})$ is identifiable given $\mathcal{P}$ if and only if $P_\mathbf{x}(\mathbf{y})$ is identifiable given every DAG $\mathcal{D}(\mathbf{V}, \mathbf{L})$ in the Markov equivalence class represented by $\mathcal{P}$, and with the same expression.*

We first derive an atomic identification criterion analogous to Corollary 1. As seen in the algorithm for constructing a partial order (Alg. 2), a bucket or circle component in a PAG is for our purpose analogous to a single node in a DAG. Therefore, the following criterion targets a bucket $\mathbf{X}$ rather than a single node.

Theorem 1. *Given a PAG $\mathcal{P}$ over $\mathbf{V}$, a partial topological order $\mathbf{B}_1 < \cdots < \mathbf{B}_m$ with respect to $\mathcal{P}$, a bucket $\mathbf{X}{=}\mathbf{B}_j \subset \mathbf{T} \subseteq \mathbf{V}$, for some $1{\le}j{\le}m$ and where $\mathbf{T}$ is a subset of the buckets in $\mathcal{P}$, and $P_{\mathbf{v}\setminus\mathbf{t}}$ (i.e. $Q[\mathbf{T}]$), $Q[\mathbf{T} \setminus \mathbf{X}]$ is identifiable if and only if there does not exist $X \in \mathbf{X}$ such that X has a possible child $C \notin \mathbf{X}$ that is in the same pc-component as X in $\mathcal{P}_\mathbf{T}$. If identifiable, then the expression is given by*

$$Q[\mathbf{T} \setminus \mathbf{X}] = \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{\prod_{\{i|\mathbf{B_i} \subseteq S^\mathbf{X}\}} P_{\mathbf{v}\setminus\mathbf{t}}(\mathbf{B_i}|\mathbf{B^{(i-1)}})} \times \sum_{\mathbf{x}} \prod_{\{i|\mathbf{B_i} \subseteq S^\mathbf{X}\}} P_{\mathbf{v}\setminus\mathbf{t}}(\mathbf{B_i}|\mathbf{B^{(i-1)}}), \quad (3)$$

where $S^\mathbf{X} = \bigcup_{X \in \mathbf{X}} S^X$, S^X being the dc-component of X in $\mathcal{P}_\mathbf{T}$, and $\mathbf{B}^{(i-1)}$ denoting the set of nodes preceding bucket $\mathbf{B}_i$ in the partial order.

Proof Sketch. **(if)** Let $\mathcal{D}$ be any DAG in the equivalence class represented by $\mathcal{P}$, $\mathcal{D}_\mathbf{T}$ be the induced subgraph over $\mathbf{T}$, and S' be the smallest composite c-component containing $\mathbf{X}$ in $\mathcal{D}_\mathbf{T}$. We show that $\mathbf{X}$ is a descendant set in $\mathcal{D}_{S'}$. Suppose otherwise for the sake of contradiction. Then, there is a node $C \in S' \setminus \mathbf{X}$ such that C is a child of X_i and is in the same c-component with X_j, where $X_i, X_j \in \mathbf{X}$ and possibly $i = j$. By Prop. 2, X_j is in the same pc-component with C in $\mathcal{P}_\mathbf{T}$. Let T_i be the node closest to X_j along the collider path in $\mathcal{P}_\mathbf{T}$ between X_j and C consistent with Def. 4. If the edge between X_j and T_i in $\mathcal{P}_\mathbf{T}$ is not into X_j, then X_j is in the same pc-component with a possible child as the edge is not visible. This violates the criterion stated in the theorem. Otherwise, the edge is $X_j \leftrightarrow T_i$ and there exist a bi-directed edge between T_i and every node in $\mathbf{X}$ (which follows from Lemma 5). Hence, X_i is in the same pc-component with a possible child C in $\mathcal{P}_\mathbf{T}$ (Prop. 1), and the criterion stated in the theorem is violated again. Therefore, $\mathbf{X}$ is a descendant set in $\mathcal{D}_{S'}$ and $Q[\mathbf{T} \setminus \mathbf{X}]$ is identifiable from $Q[\mathbf{T}]$ by Lemma 1. It remains to show that Eq. 3 is equivalent to Eq. 1 for $\mathcal{D}$. The details for this step are left to the Supplementary Material.

(only if) Suppose the criterion in question is not satisfied. Then some $X_i \in \mathbf{X}$ is in the pc-component with a possible child $C \notin \mathbf{X}$ in $\mathcal{P}_\mathbf{T}$. The edge between X_i and C is $X_i {*\!\!\rightarrow} C$ as C is outside of $\mathbf{X}$. If the edge is not visible in $\mathcal{P}_\mathbf{T}$, then this edge is not visible in $\mathcal{P}$ (Lemma 4). Hence, we can construct a DAG $\mathcal{D}$ in the equivalence class of $\mathcal{P}$ where C is a child of X_i and the two nodes share a latent variable. The pair of sets $\mathbf{F} = \{X_i, C\}$ and $\mathbf{F}' = \{C\}$ form a so-called hedge for $Q[\mathbf{T} \setminus \mathbf{X}]$ and the effect is not identifiable in $\mathcal{D}$ (Shpitser and Pearl, 2006, Theorem 4), and hence not identifiable given $\mathcal{P}$.

Otherwise, $X_i \rightarrow C$ is visible in $\mathcal{P}_\mathbf{T}$. So, there is a collider path between X_i and C consistent with Def. 4 such that the two nodes are in the same pc-component. Let $p{=}\langle X_i{=}T_0, T_1, \ldots, T_m{=}C\rangle$ denote the shortest such path in $\mathcal{P}_\mathbf{T}$. If the edge between X_i and T_1 is not into X_i, then T_1 is a child of X_i and the proof follows as in the previous case. Otherwise, we have $X_i \leftrightarrow T_1$ and we can show that X_i is the only node along p that belongs to $\mathbf{X}$ (details in the Supplementary Material). In $\mathcal{P}$, path p is present with $X_i \rightarrow C$ visible. Hence, we can construct a DAG $\mathcal{D}$ in the equivalence class of $\mathcal{P}$ such that C is a child of X_i and both are in the same c-component through a sequence of bi-directed edges along the corresponding nodes of p. The pair of sets $\mathbf{F} = \{X_i, T_1, \ldots, T_m = C\}$ and $\mathbf{F}' = \{T_1, \ldots, T_m = C\}$ form a hedge for $Q[\mathbf{T} \setminus \mathbf{X}]$ and the effect is not identifiable in $\mathcal{D}$, and hence it is not identifiable given $\mathcal{P}$. □

Note that the above result simplifies into computing the

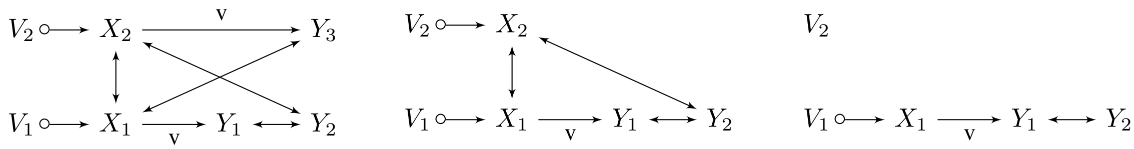

(a) (b) (c)

Figure 3: Sample PAG $\mathcal{P}$ (left) and induced subgraphs used to identify $Q[\{Y_1, Y_2\}]$.

interventional distribution $P_{\mathbf{x}}$ whenever the input distribution is the observational distribution, i.e. $\mathbf{T} = \mathbf{V}$. Consider the query $P_x(\mathbf{v} \setminus \{x\})$ over the PAG in Fig. 1b. The intervention node X is not in the same pc-component with any of its possible children (V_3 and V_4), hence the effect is identifiable and given by

$$\begin{aligned} P_x(\mathbf{v} \setminus \{x\}) &= \frac{P(\mathbf{v})}{P(x|v_1, v_2)} \times \sum_{x'} P(x'|v_1, v_2) \\ &= P(v_1, v_2) P(v_4, v_5|v_1, v_2, x) \end{aligned}$$

Putting these observations together leads to the procedure we call **IDP**, which is shown in Alg. 3. In words, the main idea of **IDP** goes as follows. After receiving the sets $\mathbf{X}, \mathbf{Y}$, and a PAG $\mathcal{P}$, the algorithm starts the pre-processing steps: First, it computes $\mathbf{D}$, the set of possible ancestors of $\mathbf{Y}$ in $\mathcal{P}_{\mathbf{V} \setminus \mathbf{X}}$. Second, it uses $\mathcal{P}_{\mathbf{D}}$ to partition set $\mathbf{D}$ into cpc-components. Following the pre-processing stage, the procedure calls the subroutine `Identify` over each cpc-component $\mathbf{D}_i$ to compute $Q[\mathbf{D}_i]$ from the observational distribution $P(\mathbf{V})$. The recursive routine basically checks for the presence of a bucket $\mathbf{B}$ in $\mathcal{P}_{\mathbf{T}}$ that is a subset of the intervention nodes, i.e. $\mathbf{B} \subseteq \mathbf{T} \setminus \mathbf{C}$, and satisfies the conditions of Thm. 1. If found, it is able to successfully compute $Q[\mathbf{T} \setminus \mathbf{B}]$ using Eq. 3, and proceed with a recursive call. Alternatively, if such a bucket doesn't exist in $\mathcal{P}_{\mathbf{T}}$, then **IDP** throws a failure condition, since it's unable to identify the query. We show next that this procedure is, indeed, correct.

Theorem 2. *Algorithm* **IDP** *(Alg.3) is sound.*

Proof. Let $\mathcal{G}(\mathbf{V}, \mathbf{L})$ be any causal graph in the equivalence class of PAG $\mathcal{P}$ over $\mathbf{V}$, and let $\mathbf{V}' = \mathbf{V} \setminus \mathbf{X}$. We have

$$P_{\mathbf{x}}(\mathbf{y}) = \sum_{\mathbf{v}' \setminus \mathbf{y}} P_{\mathbf{x}}(\mathbf{v}') = \sum_{\mathbf{v}' \setminus \mathbf{y}} Q[\mathbf{V}'] = \sum_{\mathbf{v}' \setminus \mathbf{d}} \sum_{\mathbf{d} \setminus \mathbf{y}} Q[\mathbf{V}']$$

By definition, $\mathbf{D}$ is an ancestral set in $\mathcal{P}_{\mathbf{V}'}$, and hence it is ancestral in $\mathcal{G}_{\mathbf{V}'}$ by Prop. 1. So, we have the following by (Tian, 2002, Lemma 10):

$$P_{\mathbf{x}}(\mathbf{y}) = \sum_{\mathbf{d} \setminus \mathbf{y}} \sum_{\mathbf{v}' \setminus \mathbf{d}} Q[\mathbf{V}'] = \sum_{\mathbf{d} \setminus \mathbf{y}} Q[\mathbf{D}] \tag{4}$$

Algorithm 3: IDP$(\mathbf{x}, \mathbf{y})$ given PAG $\mathcal{P}$

input : two disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$
output: Expression for $P_{\mathbf{x}}(\mathbf{y})$ or `FAIL`

1. Let $\mathbf{D} = \mathtt{An}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V} \setminus \mathbf{X}}}$
2. Let the cpc-components of $\mathcal{P}_{\mathbf{D}}$ be $\mathbf{D}_i, i = 1, \ldots, k$
3. $P_{\mathbf{x}}(\mathbf{y}) = \sum_{\mathbf{d} \setminus \mathbf{y}} \prod_i \mathtt{Identify}(\mathbf{D}_i, \mathbf{V}, P)$

Function `Identify` ($\mathbf{C}$, $\mathbf{T}$, $Q = Q[\mathbf{T}]$) **:**
  if $\mathbf{C} = \mathbf{T}$ **then**
    return $Q[\mathbf{T}]$;
  end
  `/* In` $\mathcal{P}_{\mathbf{T}}$`, let` $\mathbf{B}$ `be a bucket, and` $C^{\mathbf{B}}$ `be the pc-component of` $\mathbf{B}$ `*/`
  if $\exists \mathbf{B} \subseteq \mathbf{T} \setminus \mathbf{C}$ *such that* $C^{\mathbf{B}} \cap \mathit{Ch}(\mathbf{B}) \subseteq \mathbf{B}$ **then**
    Compute $Q[\mathbf{T} \setminus \mathbf{B}]$ from Q; `// Theorem 1`
    return `Identify`$(\mathbf{C}, \mathbf{T} \setminus \mathbf{B}, Q[\mathbf{T} \setminus \mathbf{B}])$;
  else
    throw `FAIL`;
  end

Using Prop. 3, each cpc-component in $\mathcal{P}_{\mathbf{D}}$ corresponds to a composite c-component in $\mathcal{G}_{\mathbf{D}}$. Hence, Eq. 4 can be decomposed as follows by (Tian, 2002, Lemma 11).

$$P_{\mathbf{x}}(\mathbf{y}) = \sum_{\mathbf{d} \setminus \mathbf{y}} Q[\mathbf{D}] = \sum_{\mathbf{d} \setminus \mathbf{y}} \prod_i Q[\mathbf{D}_i] \tag{5}$$

Eq. 5 is equivalent to the decomposition we have in step 3 of Alg. 3, where we attempt to compute each $Q[\mathbf{D}_i]$ from P. Finally, the correctness of the recursive routine `Identify` follows from that of Theorem 1. □

5.1 ILLUSTRATIVE EXAMPLE

Consider the query $P_{x_1, x_2}(y_1, y_2, y_3)$ given $\mathcal{P}$ in Fig. 3a. We have $\mathbf{D} = \{Y_1, Y_2, Y_3\}$, and the cpc-components in $\mathcal{P}_{\mathbf{D}}$ are $\mathbf{D}_1 = \{Y_1, Y_2\}$ and $\mathbf{D}_2 = \{Y_3\}$. Hence, the problem reduces to computing $Q[\{Y_1, Y_2\}] \cdot Q[\{Y_3\}]$.

We start with the call `Identify`$(\mathbf{D}_1, \mathbf{V}, P)$. Consider the singleton bucket Y_3 the pc-component of which includes all the nodes in $\mathcal{P}$. This node satisfies the condition in `Identify` as it has no children, and we compute

$Q[\mathbf{V} \setminus \{Y_3\}]$ using Theorem 1.

$$\begin{aligned} Q[\mathbf{V} \setminus \{Y_3\}] &= \frac{P(\mathbf{v})}{P(y_1, y_2, y_3, x_1, x_2 | v_1, v_2)} \times \\ &\quad \sum_{y_3} P(y_1, y_2, y_3, x_1, x_2 | v_1, v_2) \\ &= P(v_1, v_2) \cdot P(y_1, y_2, x_1, x_2 | v_1, v_2) \\ &= P(y_1, y_2, x_1, x_2, v_1, v_2) \end{aligned} \tag{6}$$

In the next recursive call, $\mathbf{T}_1 = \mathbf{V} \setminus \{Y_3\}$, P_{y_3} corresponds to Eq. 6, and the induced subgraph $\mathcal{P}_{\mathbf{T}_1}$ is shown in Fig. 3b. Now, X_2 satisfies the criterion and we can compute $Q[\mathbf{T}_1 \setminus \{X_2\}]$ from $P_{y_3} = Q[\mathbf{T}_1]$, i.e.,

$$\begin{aligned} Q[\mathbf{T}_1 \setminus \{X_2\}] &= \frac{P_{y_3}}{P_{y_3}(y_1, y_2, x_1, x_2 | v_1, v_2)} \times \\ &\quad \sum_{x_2} P_{y_3}(y_1, y_2, x_1, x_2 | v_1, v_2) \\ &= P(y_1, y_2, x_1, v_1, v_2) \end{aligned} \tag{7}$$

Let $\mathbf{T}_2 = \mathbf{T}_1 \setminus \{X_2\}$, where the induced subgraph $\mathcal{P}_{\mathbf{T}_2}$ is shown in Fig. 3c. Now, X_1 satisfies the criterion and we can compute $Q[\mathbf{T}_2 \setminus \{X_1\}]$ from Eq. 7,

$$\begin{aligned} Q[\mathbf{T}_2 \setminus \{X_1\}] &= \frac{P_{y_3,x_2}}{P_{y_3,x_2}(x_1 | v_1, v_2)} \times \\ &\quad \sum_{x_1} P_{y_3,x_2}(x_1 | v_1, v_2) \\ &= \frac{P(v_1, v_2) \cdot P(y_1, y_2, x_1, v_1, v_2)}{P(x_1, v_1, v_2)} \\ &= P(v_1, v_2) \cdot P(y_1, y_2 | x_1, v_1, v_2) \end{aligned}$$

Choosing V_1 and V_2 in the next two recursive calls, we finally obtain the simplified expression:

$$Q[\{Y_1, Y_2\}] = P(y_1, y_2 | x_1)$$

Next, we solve for $Q[\mathbf{D}_2]$ and we get an expression analogous to that of $Q[\mathbf{D}_1]$. Hence, the final solution is:

$$P_{x_1,x_2}(y_1, y_2, y_3) = P(y_1, y_2 | x_1) \times P(y_3 | x_2)$$

5.2 COMPARISON TO STATE OF THE ART

In the previous section, we formulated an identification algorithm in PAGs for causal queries of the form $P_{\mathbf{x}}(\mathbf{y})$, $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$. A natural question arises about the expressiveness of the **IDP** in comparison with the state-of-the-art methods. One of the well established results in the literature is the adjustment method (Perković et al., 2015), which is complete whenever an adjustment set exists.

In the sequel, we formally show that the proposed algorithm subsumes the adjustment method.

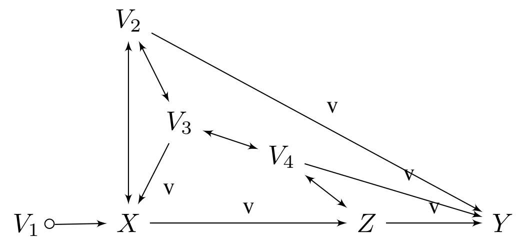

Figure 4: Query $P_x(y)$ is identifiable by **IDP**.

Theorem 3. *Let $\mathcal{P}$ be a PAG over set $\mathbf{V}$ and let $P_{\mathbf{x}}(\mathbf{y})$ be a causal query where $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$. If the distribution $P_{\mathbf{x}}(\mathbf{y})$ is not identifiable using* **IDP** *(Alg. 3), then the effect is not identifiable using the generalized adjustment criterion in (Perković et al., 2015).*

Proof Sketch. Whenever **IDP** fails to identify some query, it is due to one of the recursive calls to `Identify`. We use the failing condition inside this call to systematically identify a proper definite status non-causal path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{P}$ that is m-connecting given set Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) (Perković et al., 2016, Def. 4.1). As this set fails to satisfy the adjustment criterion, then there exist no adjustment set relative to the pair ($\mathbf{X}$,$\mathbf{Y}$) in $\mathcal{P}$ (Perković et al., 2016, Cor. 4.4). The details of the proof are left to the Supplementary Material. □

Based on this result, one may wonder whether these algorithms are, after all, just equivalent. In reality, **IDP** captures strictly more identifiable effects than the adjustment criterion. To witness, consider the PAG in Fig. 4 and note that the causal distribution $P_x(y)$ is identifiable by **IDP** but not by adjustment in this case.

6 CONCLUSION

We studied the problem of identification of interventional distributions in Markov equivalence classes represented by PAGs. We first investigated graphical properties for induced subgraphs of PAGs over an arbitrary subset of nodes with respect to induced subgraphs of DAGs that are in the equivalence class. We believe that these results can be useful to general tasks related to causal inference from equivalence classes. We further developed an identification algorithm in PAGs and proved it to subsume the state-of-the-art adjustment method.

Acknowledgments

We thank Sanghack Lee and the reviewers for all the feedback provided. Bareinboim and Jaber are supported in parts by grants from NSF IIS-1704352 and IIS-1750807 (CAREER). Zhang is supported in part by the Research Grants Council of Hong Kong under the General Research Fund LU13600715.

References

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

Elias Bareinboim and Judea Pearl. Causal inference and the data-fusion problem. *Proc. Natl. Acad. Sci.*, 113: 7345–7352, 2016.

D. Galles and J. Pearl. Testing identifiability of causal effects. In P. Besnard and S. Hanks, editors, *Uncertainty in Artificial Intelligence 11*, pages 185–195. Morgan Kaufmann, San Francisco, 1995.

Yimin Huang and Marco Valtorta. Pearl's calculus of intervention is complete. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, UAI'06, pages 217–224. AUAI Press, 2006.

Antti Hyttinen, Frederick Eberhardt, and Matti Järvisalo. Do-calculus when the true graph is unknown. In *UAI*, pages 395–404, 2015.

Amin Jaber, Jiji Zhang, and Elias Bareinboim. Causal identification under Markov equivalence. Technical report, R-35, Purdue AI Lab, Department of Computer Science, Purdue University, 2018a.

Amin Jaber, Jiji Zhang, and Elias Bareinboim. A graphical criterion for effect identification in equivalence classes of causal diagrams. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI'18, 2018b.

Manabu Kuroki and Masami Miyakawa. Identifiability criteria for causal effects of joint interventions. *Journal of the Japan Statistical Society*, 29(2):105–117, 1999.

Marloes H. Maathuis, Diego Colombo, Markus Kalisch, and Peter Bühlmann. Predicting causal effects in large-scale systems from observational data. *Nature Methods*, 7(4):247–248, 2010.

J. Pearl. Aspects of graphical models connected with causality. In *Proceedings of the 49th Session of the International Statistical Institute*, pages 391–401, Tome LV, Book 1, Florence, Italy, 1993.

J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2000. 2nd edition, 2009.

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

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

Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. A complete generalized adjustment criterion. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 682–691, 2015.

Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. Complete graphical characterization and construction of adjustment sets in Markov equivalence classes of ancestral graphs. *arXiv preprint arXiv:1606.06903*, 2016.

Thomas Richardson and Peter Spirtes. Ancestral graph Markov models. *Annals of Statistics*, pages 962–1030, 2002.

Thomas S Richardson, Robin J Evans, James M Robins, and Ilya Shpitser. Nested Markov properties for acyclic directed mixed graphs. *arXiv preprint arXiv:1701.06686*, 2017.

Ilya Shpitser and Judea Pearl. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, page 1219. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006.

Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*, volume 81. MIT press, 2001.

Jin Tian. *Studies in causal reasoning and learning*. PhD thesis, University of California, Los Angeles, 2002.

Jin Tian and Judea Pearl. A general identification condition for causal effects. In *AAAI/IAAI*, pages 567–573, 2002.

TS Verma. Graphical aspects of causal models. *Technical R eport R-191, UCLA*, 1993.

Jiji Zhang. *Causal inference and reasoning in causally insufficient systems*. PhD thesis, Carnegie Mellon University, 2006.

Jiji Zhang. Generalized do-calculus with testable causal assumptions. In *International Conference on Artificial Intelligence and Statistics*, pages 667–674, 2007.

Jiji Zhang. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research*, 9(Jul):1437–1474, 2008a.

Jiji Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16): 1873–1896, 2008b.

“Causal Identification under Markov Equivalence” Supplemental Material

In this document, we present the full proofs for Proposition 2, Theorem 1, and Theorem 3. We also reproduce the main two lemmas from (Tian, 2002), the soundness proof for the PTO algorithm for full PAGs, and some relevant results from the literature that are used in the proofs. We start by introducing some definitions that are used in the proofs.

A Relevant Definitions

Definition 8 (Definite Status Node)**.** *A non-endpoint node X along a path is called a definite collider (DC) if both incident edges are into X. A non-endpoint node X along a path is called a definite non-collider (DNC) if one of the incident edges is out of X or it is $A \ast\!\!-\!\!\circ\, X \circ\!\!-\!\!\ast B$ such that A and B are not adjacent.*

Definition 9 (Definite Status Path)**.** *A path is of definite status if every non-endpoint node along the path is of definite status.*

Definition 10 (Discriminating Path)**.** *Given a MAG $\mathcal{M}$ or a PAG $\mathcal{P}$, a path between D and C, $u = \langle D, \ldots, A, B, C\rangle$, is said to be a discriminating path for B if:*

1. *u includes at least three edges.*
2. *B is a non-endpoint node along the path, and is adjacent to C on u.*
3. *D is not adjacent to C, and every node between D and B is a (definite) collider on u and is a parent of C.*

B Full Proofs

B.1 Proposition 2

Proof of Proposition 2. If X and Y are in the same c-component in $\mathcal{D}_{\mathbf{A}}$, then there is a path p in $\mathcal{D}_{\mathbf{A}}$ composed of nodes $\langle X = V_0, \ldots, Y = V_m\rangle$, $m \geq 1$, such that $V_i \in \mathbf{A}$ and $V_i \leftarrow L_{i,i+1} \rightarrow V_{i+1}$, $0 \leq i < m$.

By lemma 6, X and Y are in the same pc-component in $\mathcal{M}$, the MAG of $\mathcal{D}$ over $\mathbf{V}$, due to a path p' over a subsequence of p. Also, by lemma 7, X and Y are in the same pc-component in $\mathcal{P}$, the PAG of $\mathcal{M}$, due to a path p^* over a subsequence of p'. Since all the nodes along p^* are in $\mathbf{A}$, then p^* is present in $\mathcal{P}_{\mathbf{A}}$, hence X and Y are in the same pc-component in $\mathcal{P}_{\mathbf{A}}$. ☐

Lemma 6. *Let $\mathcal{G}(\mathbf{V}, \mathbf{L})$ be a DAG, and $\mathcal{M}$ be the MAG over $\mathbf{V}$ that represents the DAG. For any $X, Y \in \mathbf{V}$, if there is a bi-directed path p between X and Y in $\mathcal{G}$, then there is a path p' between X and Y in $\mathcal{M}$ over a subsequence of p such that (1) all the non-endpoint nodes are colliders, and (2) all directed edges on p' are not visible.*

Proof. We prove this lemma by induction on the size (number of bi-directed edges) of the bi-directed path between X and Y in $\mathcal{G}$. Before we start with the proof, note that the only possible directed edges along the path p' are the first and last edges since all non-endpoint nodes are colliders. Also, we will prove the following stronger statement in order to prove our main lemma.

> Let $\mathcal{G}(\mathbf{V}, \mathbf{L})$ be a DAG, and $\mathcal{M}$ be the corresponding MAG over $\mathbf{V}$. For any $X, Y \in \mathbf{V}$, if there is a bi-directed path p between X and Y in $\mathcal{G}$, then there is a path p' between X and Y over a subsequence of p in $\mathcal{M}$ such that:
>
> 1. all the non-endpoint nodes are colliders; and
> 2. all directed edges $A \rightarrow B$ on p', have an inducing path between A and B into A in $\mathcal{G}$.

The above statement implies the main lemma as the only difference is the second point regarding the directed edges along path p' in $\mathcal{M}$. It is evident that a directed edge in MAG $\mathcal{M}$ is not visible if it has an inducing path into its source in DAG $\mathcal{G}$ where $\mathcal{G}$ is in the equivalence class of $\mathcal{M}$.

Consider as a base case path p of size 1 between X and Y in $\mathcal{G}$. The bi-directed edge between X and Y ($X \leftarrow L_{XY} \rightarrow Y$) is an inducing path between X and Y relative to the latent variable L_{XY} in $\mathcal{G}$. Hence, X and Y are adjacent in MAG $\mathcal{M}$. If neither X nor Y is an ancestor of the other in $\mathcal{G}$, then the edge between X and Y in $\mathcal{M}$ is bi-directed and the lemma holds. If X is an ancestor of Y or the opposite in $\mathcal{G}$, then the edge in $\mathcal{M}$ is directed ($X \rightarrow Y$ or $X \leftarrow Y$). In both cases, the latent variable L_{XY} in $\mathcal{G}$ provides an inducing path between X and Y that is into both ends. Hence, the lemma holds.

In the induction step, we assume that the lemma holds for bi-directed paths between X and Y in $\mathcal{G}$ of length $\leq n$ and prove it for bi-directed paths of length $n+1$. Consider the bi-directed path $p = \langle X, O_1, \dots, O_n, Y \rangle$ of length n+1 in $\mathcal{G}$. We split the path into two paths $p_1 = \langle X, O_1 \rangle$ and $p_2 = \langle O_1, \dots, Y \rangle$. By the assumption of the induction step, the lemma holds for each path and there are collider paths p'_1 and p'_2 in MAG $\mathcal{M}$ corresponding to p_1 and p_2 in $\mathcal{G}$. Consider the options of path p'_1:

case 1: O_1 is not an ancestor of X in $\mathcal{G}$. Hence, the edge between X and O_1 in $\mathcal{M}$ is into O_1 ($X \rightarrow O_1$ or $X \leftrightarrow O_1$). If the edge is directed ($X \rightarrow O_1$), there is an inducing path into X in $\mathcal{G}$ ($X \leftarrow L_{XO_1} \rightarrow O_1$). Now consider the first edge in p'_2 which, according to the lemma, could be one of the following two options:

case 1-1: The edge is $O_1 \leftrightarrow O_i$. Then, edge $X \rightarrow O_1$ concatenated with path p'_2 create a path between X and Y in $\mathcal{M}$ which is consistent with the conditions in the lemma. Hence, the lemma holds.

case 1-2: The edge is $O_1 \rightarrow O_i$. By the induction step, there is an inducing path between O_1 and O_i into O_1 in $\mathcal{G}$ so the edge $O_1 \rightarrow O_i$ is not visible in $\mathcal{M}$. Hence, X and O_i must be adjacent in $\mathcal{M}$ for otherwise edge $O_1 \rightarrow O_i$ would be visible. Also, the edge between X and O_i is directed into O_i, otherwise the MAG is not ancestral. Moreover, there is an inducing path between X and O_1 that is into X and there is an inducing path between O_1 and O_i that is into O_1. Node O_1 is a collider between those two inducing paths and an ancestor of O_i. So, there is an inducing path between X and O_i that is into X in $\mathcal{G}$. Hence, there is a path between X and Y in $\mathcal{M}$ composed of edge $X \rightarrow O_i$ and the subpath of p'_2 $\langle O_i, \dots, Y \rangle$ which is consistent with the conditions of the lemma.

case 2: O_1 is an ancestor of X in $\mathcal{G}$. Hence, the edge between X and O_1 in $\mathcal{M}$ is $X \leftarrow O_1$. Also, there is an inducing path between X and O_1 into O_1 in $\mathcal{G}$ ($X \leftarrow L_{XO_1} \rightarrow O_1$) so edge $X \leftarrow O_1$ in $\mathcal{M}$ is not visible. Next, we consider the first edge in path p'_2.

case 2-1: The edge is $O_1 \leftrightarrow O_i$. Then, path p'_2 is a collider path into O_1 in $\mathcal{M}$. We show, by induction, that X is a child of each node along p'_2 in $\mathcal{M}$ starting from O_1 until there exist a bi-directed edge between X and some node along p'_2. For the base case, O_i is adjacent to X for otherwise edge $X \leftarrow O_1$ is visible. Also, the edge is into X, otherwise the MAG is not ancestral. If the edge is bi-directed, then we are done and the statement holds, otherwise the edge is out of O_i. In the induction step, assume all nodes along path p'_2 starting with O_i until O_j are parents of X. We have a collider path between O_{j+1}, which could be Y, and O_1 that is into O_1 and each collider is a parent of X. Hence, node O_{j+1} must be adjacent to X for otherwise $X \leftarrow O_1$ is visible. Also, the edge between X and O_{j+1} is into X for otherwise $\mathcal{M}$ is not ancestral due to $O_j \rightarrow X \rightarrow O_{j+1}$ and $O_{j+1} {*\!\!\rightarrow} O_j$. Hence, the statement holds.
If there exist a bi-directed edge between X and some node O_k along p'_2, then there exist a path between X and Y composed of $X \leftrightarrow O_k$ and the subpath of p'_2 $\langle O_k, \dots, Y \rangle$ that is consistent with the lemma. Otherwise, X and Y are adjacent and Y is a parent of X by the above induction. In the latter case, consider the concatenated path of p'_1 and p'_2 ($X \leftarrow O_1 \leftrightarrow O_i \longleftarrow\!\!-\!\!-\!\!\longrightarrow O_j \leftarrow\!\!* Y$). There exist an inducing path between every two consecutive nodes and every non-endpoint node is an ancestor of X. Also, by the induction assumption, the first and last edges along the path have inducing paths into the source node (O_1, and Y if $O_j \leftarrow Y$). Hence, there exist an inducing path between X and Y that is into Y in $\mathcal{G}$ and the lemma holds.

case 2-2: The edge is $O_1 \rightarrow O_i$. Both edges $X \leftarrow O_1$ and $O_1 \rightarrow O_i$ have inducing paths into O_1 in $\mathcal{G}$ by the induction assumption. Hence, X and O_i are adjacent in $\mathcal{M}$. We consider all the orientations of the edge between X and O_i in $\mathcal{M}$. If the edge is bi-directed, then there is a path between X and Y in $\mathcal{M}$ composed of $X \leftrightarrow O_i$

and the subpath of p_2' $\langle O_i, \ldots, Y\rangle$ consistent with the conditions of the lemma. If the edge is out of X, then we have the path $\langle X, O_i, \ldots, Y\rangle$. Moreover, the edge $X \to O_i$ has an inducing path into X in $\mathcal{G}$. The reason is that O_1 is a collider in $\mathcal{G}$ between the two inducing paths corresponding to edges $X \leftarrow O_1$ and $O_1 \to O_i$ and the node is an ancestor of X and O_i. Hence, the lemma holds too. The last case is when the edge is out of O_i, hence we have the path $X \leftarrow O_i, \ldots, Y$. Similar to the earlier case, the edge $X \leftarrow O_i$ has an inducing path into O_i in $\mathcal{G}$. The argument for this case follows same as that of case 2-1.

This exhausts all the options in the induction step and the lemma holds. □

Lemma 7. *Let $\mathcal{M}$ be a MAG over $\mathbf{V}$, and $\mathcal{P}$ be the PAG that represents the equivalence class of $\mathcal{M}$. For any $X, Y \in \mathbf{V}$, if there is a path p between X and Y in $\mathcal{M}$ such that (1) all non-endpoint nodes are colliders and (2) all directed edges, if any, are not visible, then there is a path p^* between X and Y in $\mathcal{P}$ over a subsequence of p such that (1) all non-endpoint nodes along the path are definite colliders, and (2) none of the edges are visible.*

Proof. We use the following lemma from (Zhang, 2006, pp. 208) several times throughout the proof.

> [**Lemma I**:] If a path $\langle U, \ldots, X, Y, Z\rangle$ is a discriminating path for Y in $\mathcal{M}$, and the corresponding subpath between U and Y in $\mathcal{P}$ is (also) a collider path, then the path is also a discriminating path for Y in $\mathcal{P}$.

In order to prove the main lemma, we start with a simpler lemma.

> [**lemma II**:] For any $X, Y \in \mathbf{V}$, if X and Y are not adjacent and there is a collider path between X and Y in $\mathcal{M}$, denoted $p = \langle X = O_0, \ldots, O_n = Y\rangle$, and p is the shortest collider path between X and Y over any subsequence of $O_1, \ldots, O_{n-1}$ in $\mathcal{M}$, then path p is of definite status in $\mathcal{P}$.

proof: We prove this lemma by contradiction. For that purpose, we first establish the following claim.

> [**claim**$_{II}$:] For every $1 \leq j \leq n-1$, if O_j is not of a definite status on p in PAG $\mathcal{P}$, then O_{j+1} is a parent of O_{j-1} in $\mathcal{M}$.

The claim holds if all the non-endpoint nodes are of definite status. However, suppose there exist at least one non-endpoint node that is not of definite status. We prove the claim by induction. In the base case, let O_j where $1 \leq j \leq n-1$ be the first non-definite status node along p that is closest to X. Then, nodes O_{j-1} and O_{j+1} must be adjacent for otherwise we are able to detect the collider at O_j. The edge between O_{j-1} and O_{j+1} can't be bi-directed in $\mathcal{M}$ since this creates a shorter collider path over a subsequence of $O_1, \ldots, O_{n-1}$ and violates our choice of path p in lemma II. Suppose O_{j-1} is a parent of O_{j+1} in $\mathcal{M}$. The collider path between X and O_j is of definite status in $\mathcal{P}$ as O_j is the first node that is not of definite status along the path. If O_{j-2} is not adjacent to O_{j+1} then $\langle O_{j-2}, O_{j-1}, O_j, O_{j+1}\rangle$ is a discriminating path for O_j in $\mathcal{M}$ and the subpath $\langle O_{j-2}, O_{j-1}, O_j\rangle$ is a collider path in $\mathcal{P}$. Hence, O_j is oriented as a collider along the path in $\mathcal{P}$ by lemma I and that contradicts our choice of O_j which is not of definite status. Hence, O_{j-2} is adjacent to O_{j+1} and the edge is out of O_{j-2}. If the edge is bi-directed, we violate our choice of p in lemma II as shortest over $O_1, \ldots, O_{n-1}$ and if the edge is out of O_{j+1}, we violate the ancestral property in $\mathcal{M}$. We apply the previous argument by induction on all the nodes starting with O_{j-2} back to O_0. Hence, each node O_k where $0 \leq k \leq j-1$ is adjacent to O_{j+1} and the edge is out of O_k in $\mathcal{M}$. Now, X has an edge out of it and into O_{j+1} which violates our choice of p in lemma II as the shortest path. Therefore, the initial assumption about the orientation of the edge between O_{j-1} and O_{j+1} is invalid and the edge is out of O_{j+1} in $\mathcal{M}$.

In the induction step, we assume that the lemma holds for some non-definite status node O_r where $1 \leq r \leq n-1$ and we prove that the lemma holds for the next non-definite status node O_{r+l} along the path. Similar to the base case argument, the nodes O_{r+l-1} and O_{r+l+1} must be adjacent and the edge can't be bi-directed in $\mathcal{M}$. Assume that the edge is out of O_{r+l-1} in $\mathcal{M}$, then we show by induction, similar to the base case, that each node O_j where $r \leq j \leq r+l-1$ is adjacent to O_{r+l+1} and the edge is out of O_j. Now, we consider node O_r which is not of definite status and note that O_{r+1} is a parent of O_{r-1} in $\mathcal{M}$ by the induction step. We use the same argument by induction as before to show that each node O_k where $r+1 \leq k \leq r+l$ is adjacent to O_{r-1} and the edge is out of O_k. This creates an inducing path between O_{r-1} and O_{r+l+1} hence those two nodes must be adjacent in $\mathcal{M}$, otherwise we violate the maximality property. The edge between O_{r-1} and O_{r+l+1} can't be bi-directed in $\mathcal{M}$ as it violates the choice of the path in lemma II. If the edge is out of O_{r-1}, then we violate the ancestral property of $\mathcal{M}$ due to the structure $O_{r+l} \rightarrow O_{r-1} \rightarrow O_{r+l+1}$ and $O_{r+l} \leftarrow\!* O_{r+l+1}$. Similarly, the edge can't be out of O_{r+l+1}. Hence, the initial assumption that the edge between O_{r+l-1} and O_{r+l+1} is out of O_{r+l-1} is invalid and the edge is out of O_{r+l+1} in $\mathcal{M}$. Hence , the claim holds for O_{r+l}.

With the previous claim proven, we finish our proof for lemma II by establishing the following claim.

> [**claim**$^*_{II}$:] For every $1 \leq j \leq n-1$, if O_j is not of a definite status on p in $\mathcal{P}$, then O_{j-1} is a parent of O_{j+1} in $\mathcal{M}$.

claim$^*_{II}$ is symmetric to **claim**$_{II}$. Hence, its proof is carried out same as the proof of the previous claim with the difference of starting with the first non-definite status node that is closest to $Y(O_n)$ instead of $X(O_0)$. However, both claims being valid is contradicting as long as there exist at least one non-definite status node O_j where $1 \leq j \leq n-1$. Therefore, all the nodes along path p are of definite status. This concludes our proof of lemma II.

Now, we are ready to prove our main lemma. We choose the shortest subsequence of p between X and Y in $\mathcal{M}$, denoted p^*, such that (1) all none endpoint nodes along the path are colliders, and (2) none of the directed edges, possibly the first and last edges along p^*, are visible. We show that path p^* is of definite status and none of the directed edges along the path are visible in $\mathcal{P}$.

If p^* is a single edge between X and Y, then the edge is not visible in $\mathcal{M}$ and hence it is not visible in $\mathcal{P}$ and the lemma holds. Otherwise, path p^* is a collider path with at least one non-endpoint node. It is evident that if the first or last edge along p^* is directed and not visible in $\mathcal{M}$, then those edges will not be visible in $\mathcal{P}$ if directed. We use a proof similar to that of lemma II to prove that p^* is of definite status in $\mathcal{P}$ and start with the following claim.

> [**claim**$_I$:] For every $1 \leq j \leq n-1$, if O_j is not of a definite status on p^* in $\mathcal{P}$, then O_{j+1} is a parent of O_{j-1} in $\mathcal{M}$.

We prove this claim by induction. In the base case, we choose the first node O_j that is not of definite status along the path in $\mathcal{P}$ and closest to X. O_{j-1} and O_{j+1} are adjacent in $\mathcal{M}$ and the edge can't be bi-directed because of the choice of p^*. First, we assume that the edge is out of O_{j-1}. Similar to the proof of **claim**$_{II}$, each node O_k where $0 \leq k \leq j-1$ must be adjacent to O_{j+1} and the edge is out of O_k in $\mathcal{M}$. Then, X and O_{j+1} are adjacent and the edge is out of X in $\mathcal{M}$. This edge has to be visible in $\mathcal{M}$ for otherwise we violate our choice of p^*. Consider the shortest path into X consistent with the graphical condition for visibility that makes edge $X \rightarrow O_{j+1}$ visible in $\mathcal{M}$. We refer to the path as p^v and denote the nodes along p^v as $\langle C_m, \ldots C_1 \rangle$ where $m \geq 1$. Note that $m = 1$ when the visibility is due to a single edge into X ($C_1 *\!\!\rightarrow X \rightarrow O_{j+1}$) and C_1 is not adjacent to O_{j+1}.

Let p' denote the concatenated path of p^v and the subpath of p^* $\langle O_0 = X, \ldots, O_{j+1} \rangle$. The following two points hold:

1. There exist at least one collider path between C_m and O_{j+1} over a subsequence of the non-endpoint nodes along p'. Note that all the non-endpoint nodes along p' are colliders in $\mathcal{M}$ except for $\langle C_1, X, O_1 \rangle$ where the edge between X and O_1 can be directed out of X. If the edge between X and O_1 is bi-directed, then we are done as p' is a collider path between C_m and O_{j+1} in $\mathcal{M}$. If the edge is directed ($X \rightarrow O_1$) in $\mathcal{M}$, we prove by induction that O_1 is the child of every node along p^v starting with C_1 until O_1 is connected with a bi-directed edge to some C_i in $\mathcal{M}$. In the base case, the edge between C_1 and X is into X so C_1 and O_1 must be adjacent for otherwise $X \rightarrow O_1$ is visible and this violates our choice of p^* where none of the edges are visible. The edge between C_1 and O_1 can't be out of O_1 for this violates the ancestral property of $\mathcal{M}$ due to $X \rightarrow O_1 \rightarrow C_1 *\!\!\rightarrow X$. If the edge is bi-directed, then we are done. Otherwise, the edge is out of C_1 and we continue with the induction step. In the induction step, we consider node C_i and assume that all the nodes C_j where $1 \leq j < i \leq m$ are parents of O_1. This creates a collider path between C_i and X that is into X and each collider along the path is a parent of O_1 hence $X \rightarrow O_1$ is visible in $\mathcal{M}$. Thus, C_i must be adjacent to O_1 and the edge can't be out of O_1 as this violates the ancestral property due to $C_{i-1} \rightarrow O_1 \rightarrow C_i *\!\!\rightarrow C_{i-1}$. Hence, the edge has to be either bi-directed or out of C_i. This concludes the proof.
2. Every collider path between C_m and O_{j+1} over a subsequence of the non-endpoint nodes along p' must go through O_j. The reason is that O_{j+1} is a child of each non-endpoint node along p' except for O_j and it is not

adjacent to C_m due to the visibility of edge $X \to O_{j+1}$. O_j is the only node among the nodes in p' such that O_{j+1} is connected to and the edge is into O_j in $\mathcal{M}$. Hence, the second point is valid as well.

By the above two points, we choose the shortest collider path in $\mathcal{M}$ between C_m and O_{j+1} over any subsequence of the nodes along p'. By lemma II, this is a definite status path in $\mathcal{P}$, and the edge between O_{j+1} and O_j is into O_j.

If the edge between O_{j-1} and O_j has a circle incident on O_j in $\mathcal{P}$, then the edge between O_{j-1} and O_{j+1} has an arrowhead incident on O_{j-1} in $\mathcal{P}$ (Lemma 15). However, this is not possible as the edge between O_{j-1} and O_{j+1} is out of O_{j-1} in $\mathcal{M}$ according to the assumption in the base case of the proof of **claim**$_I$. Since $\mathcal{P}$ is the PAG representing the equivalence class of $\mathcal{M}$, then $\mathcal{P}$ can't have an arrowhead incident on O_{j-1} from O_{j+1} as this orientation is not invariant in all the MAGs in the equivalence class. Hence, the edge between O_{j-1} and O_j can't have a circle incident on O_j and the edge will be into O_j in $\mathcal{P}$. Having O_j a definite collider along $\langle O_{j-1}, O_j, O_{j+1}\rangle$ contradicts our assumption in the base case that O_j is the first node that is not of definite status along p^* in $\mathcal{P}$. Hence, the initial assumption that the edge between O_{j-1} and O_{j+1} is out of O_{j-1} in $\mathcal{M}$ is invalid and the edge must be out of O_{j+1}.

In the induction step, we apply the exact same argument as that applied in the induction step of the proof of **claim**$_{II}$. We omit the proof to avoid redundancy. This concludes the proof of claim$_I$. Next, we prove the following claim which is symmetric to **claim**$_I$.

> [**claim**$_I^*$:] For every $1 \le j \le n-1$, if O_j is not of a definite status on p^* in $\mathcal{P}$, then O_{j-1} is a parent of O_{j+1} in $\mathcal{M}$.

Again, the proof of **claim**$_I^*$ is exactly the same as the proof of **claim**$_I$ with the difference of starting the induction with the first node that is not of definite status along p^* and is closest to Y instead of X. However, having both claims valid is contradicting as long as there exist at least one non-definite status node along p^* in $\mathcal{P}$. Hence, path p^* is of definite status and any directed edge along the path of not visible. This concludes our proof for the main lemma. $\square$

B.2 Theorem 1

Proof of Theorem 1. **(if)** Let $\mathcal{D}$ be any DAG in the equivalence class represented by $\mathcal{P}$, $\mathcal{D}_{\mathbf{T}}$ be the induced subgraph over $\mathbf{T}$, and S' be the smallest composite c-component containing $\mathbf{X}$ in $\mathcal{D}_{\mathbf{T}}$. We show that $\mathbf{X}$ is a descendant set in $\mathcal{D}_{S'}$. Suppose otherwise for the sake of contradiction. Then, there is a node $C \in S' \backslash \mathbf{X}$ such that C is a child of X_i and is in the same c-component with X_j, where $X_i, X_j \in \mathbf{X}$ and possibly $i = j$. By Prop. 2, X_j is in the same pc-component with C in $\mathcal{P}_{\mathbf{T}}$. Let T_i be the node closest to X_j along the collider path in $\mathcal{P}_{\mathbf{T}}$ between X_j and C consistent with Def. 4. If the edge between X_j and T_i in $\mathcal{P}_{\mathbf{T}}$ is not into X_j, then X_j is in the same pc-component with a possible child as the edge is not visible. This violates the criterion stated in the theorem. Otherwise, the edge is $X_j \leftrightarrow T_i$ and there exist a bi-directed edge between T_i and every node in $\mathbf{X}$ (which follows from Lemma 5). Hence, X_i is in the same pc-component with a possible child C in $\mathcal{P}_{\mathbf{T}}$ (Prop. 1), and the criterion stated in the theorem is violated again. Therefore, $\mathbf{X}$ is a descendant set in $\mathcal{D}_{S'}$ and $Q[\mathbf{T} \setminus \mathbf{X}]$ is identifiable from $Q[\mathbf{T}]$ by Lemma 1. It remains to show that Eq. 3 is equivalent to Eq. 1 for $\mathcal{D}$. Eq. 3 differs from Eq. 1 by using:

1. product of conditional terms instead of $Q[S^{\mathbf{X}}]$,
2. conditional terms over buckets instead of single variables, and
3. dc-component instead of c-component.

In what follows, we address those differences.

(1) Consider Eq. 9 which computes $Q[\mathbf{H}_j]$ from $Q[\mathbf{H}]$. Each fraction $\frac{Q[H^{(i)}]}{Q[H^{(i-1)}]}$ for node V_{h_i} is nothing but the conditional probability $P_{\mathbf{v}\setminus\mathbf{h}}(V_{h_i}|V^{h_{i-1}})$, where $Q[\mathbf{H}]$ is identifiable and given by $P_{\mathbf{v}\setminus\mathbf{h}}$. Therefore, the product of conditional terms in Eq. 3 substitutes $Q[S^{\mathbf{X}}]$ in Eq. 2 whenever $S^{\mathbf{X}}$ corresponds to the composite c-component containing $\mathbf{X}$ in $\mathcal{D}$. Whether the latter is true or not is addressed in point 3.

(2) If some node in bucket $\mathbf{B}$ is in the dc-component of $X_i \in \mathbf{X}$, then every node in $\mathbf{B}$ is in the same dc-component with each node in $\mathbf{X}$ (Lemma 16). It follows that the dc-component of $\mathbf{X}$ is a union of one or more buckets in $\mathcal{P}_{\mathbf{T}}$. For any full topological order over $\mathbf{T}$ that is consistent with the partial order used in Eq. 3, the nodes within a bucket will be consecutive. Hence, the corresponding conditional terms can be jointly expressed by $P_{\mathbf{v}\setminus\mathbf{t}}(\mathbf{B_i}|\mathbf{B^{(i-1)}})$.

(3) The pc-component of $\mathbf{X}$ in $\mathcal{P}_{\mathbf{T}}$ subsumes the composite c-component containing $\mathbf{X}$ in $\mathcal{D}_{\mathbf{T}}$ (Prop. 2). In Eq. 3, we exclude the nodes that are in the pc-component, but are not in the dc-component of $\mathbf{X}$. In $\mathcal{P}_{\mathbf{T}}$, let C be one such node in bucket $\mathbf{B_i}$ that is in the pc-component, but not the dc-component, of $\mathbf{X}$. It follows that there is no bi-directed path between $\mathbf{X}$ and $\mathbf{B_i}$ and no node in $\mathbf{B_i}$ is in the dc-component of $\mathbf{X}$. If $\mathbf{B_i}$ comes before $\mathbf{X}$ in the partial order, then the conditional term for C in Eq. 1 doesn't condition on $\mathbf{X}$ and the term can be pushed behind the summation over $\mathbf{X}$ and cancels out with the corresponding term in the numerator. Otherwise, consider any proper definite status path between $\mathbf{X}$ and $\mathbf{B_i}$. One of the following properties must be true:

1. path contains at least one definite non-collider, or
2. path is not into $\mathbf{B_i}$, or
3. path is not into $\mathbf{X}$.

Consider a proper definite status path between $\mathbf{X}$ and $\mathbf{B_i}$ where the first two properties don't hold, i.e. $X_i{*\!\!\rightarrow} Y \longleftarrow\!\!-\!-\!\!\longrightarrow T_i \in \mathbf{B_i}$. The edge between X_i and Y must be visible, otherwise X_i is in the same pc-component with a non-intervention child (Y). So, we have Y in the dc-component of T_i, and consequently in the dc-component of every node in $\mathbf{B_i}$. Also, note that the first edge along the collider path between $\mathbf{X}$ and C consistent with Def. 4 is into $\mathbf{X}$, otherwise we violate the theorem condition. It follows that in $\mathcal{P}_{\mathbf{T}}$, C is in the dc-component of Y and C is also in the pc-component of X_i while Y is a child of X_i. Hence, we can construct a DAG $\mathcal{D}'$ in the equivalence class of $\mathcal{P}$ where X_i and Y are in the same c-component while Y is a child of X_i (Lemmas 24 and 14). In the subgraph $\mathcal{D}'_{\mathbf{T}}$, X_i remains in the same c-component with Y, hence X_i will be in the pc-component of Y (Lemma 2) and this violates the criterion of the proposition. Therefore, every proper definite status path between $\mathbf{X}$ and $\mathbf{B_i}$ in $\mathcal{P}_{\mathbf{T}}$ has one of the first two properties true.

Also, no node in $\mathbf{B_i}$ is adjacent to a node in $\mathbf{X}$ in $\mathcal{P}_{\mathbf{T}}$. Suppose for contradiction that the latter is not true. So the edge between the two buckets must be bi-directed or possibly out of $\mathbf{X}$ and into $\mathbf{B_i}$ ($X_i{*\!\!\rightarrow} T_i \in \mathbf{B_i}$) as $\mathbf{B_i}$ comes after $\mathbf{X}$ in the partial order. The first case is not possible as $\mathbf{B_i}$ would be in the same dc-component as $\mathbf{X}$. In the second case, we have a possibly directed edge in $\mathcal{P}_{\mathbf{T}}$ between X_i and every node in $\mathbf{B_i}$ including C (Lemma 16). So, X_i will be in the same pc-component with a possible child (C) which contradicts the criterion in the current proposition. By Lemmas 8, 9, and 10, we have $(\mathbf{B_i} \perp\!\!\!\perp \mathbf{X} | \mathbf{B^{(i-1)}} \setminus \mathbf{X})$ in $\mathcal{D}_{\mathbf{T}}$. This implies that the conditional term of C is independent of $\mathbf{X}$, and even if it was included in $S^{\mathbf{X}}$, the expression can be simplified to the one in Eq. 3.

This concludes the proof for when $\mathbf{T} = \mathbf{V}$ as $\mathcal{P}_{\mathbf{T}}$ will be the full PAG. A bi-directed edge in a PAG implies that the two nodes are in the same c-component in every DAG in the equivalence class, hence the dc-component of $\mathbf{X}$ in $\mathcal{P}$ is part of the composite c-component of $\mathbf{X}$ in $\mathcal{D}$. However, $\mathcal{P}_{\mathbf{T}}$ preserves the causal relations between the nodes in $\mathcal{D}_{\mathbf{T}}$, but not necessarily the spurious ones (Prop. 1). Hence, it is not necessarily true that $S^{\mathbf{X}}$ is a subset of the composite c-component of $\mathbf{X}$ in $\mathcal{D}_{\mathbf{T}}$.

Let $T_i \in \mathbf{B_i} \subset S^{\mathbf{X}}$ be a node in $\mathcal{D}_{\mathbf{T}}$ which is not in the composite c-component of $\mathbf{X}$, let $\mathcal{M}$ be the MAG corresponding to $\mathcal{D}_{\mathbf{T}}$, and consider a full topological order over $\mathcal{M}$ that is consistent with the partial order in Th. 1. If T_i comes before $\mathbf{X}$ in the full order, then $\mathbf{B_i}$ comes before $\mathbf{X}$ in the partial order over $\mathcal{P}_{\mathbf{T}}$ and the conditional term in Eq. 3 doesn't condition on $\mathbf{X}$. Hence, Eq. 3 is valid for $\mathcal{D}_{\mathbf{T}}$ as we can use the full order to decompose the conditional term of $\mathbf{B_i}$, then simplify the expression by canceling the conditional term for T_i with the corresponding one in the denominator. The other option is when T_i comes after $\mathbf{X}$ in the full order and consequently the partial order over $\mathcal{P}_{\mathbf{T}}$. Since T_i is not in the c-component of any node in $\mathbf{X}$, then there is no bi-directed path between T_i and $\mathbf{X}$ in $\mathcal{M}$ and one of the following is true for any proper path between them:

1. path contains at least one non-collider, or
2. path is out of T_i, or
3. path is out of $\mathbf{X}$.

Consider a proper path between $X_i \in \mathbf{X}$ and T_i in $\mathcal{M}$ where the first two properties don't hold, i.e. $X_i \rightarrow Y \longleftarrow\!\!-\!-\!\!\longrightarrow T_i$. Since Y and T_i share a bi-directed path in $\mathcal{M}$, then Y and T_i are in the same c-component in $\mathcal{D}_{\mathbf{T}}$. But, X_i and $T_i \in S^{\mathbf{X}}$ are in the same dc-component in $\mathcal{P}_{\mathbf{T}}$, and consequently $\mathcal{P}$. Hence, X_i and Y are in the same c-component in $\mathcal{D}$. It follows that we can construct a DAG $\mathcal{D}'$ in the equivalence class of $\mathcal{P}$ such that X_i and Y are in the same c-component in $\mathcal{D}'_{\mathbf{T}}$ while keeping Y a child of X_i. Hence, X_i will be in the same pc-component with a possible child Y in $\mathcal{P}_{\mathbf{T}}$ (Props. 1 and 2). The latter violates the criterion of the theorem at hand, hence such a path cannot exist in $\mathcal{M}$ and every proper path between $\mathbf{X}$ and T_i must have one of the first two properties true.

Also, T_i is not adjacent to any node in $\mathbf{X}$ in $\mathcal{M}$. Suppose for contradiction that the latter is not true. If the edge is bi-directed, this contradicts the assumption that T_i is not in the same c-component of any node in set $\mathbf{X}$ in $\mathcal{D}_{\mathbf{T}}$. If the edge is directed, it must be out of $\mathbf{X}$ ($X_i \rightarrow T_i$) as T_i comes after $\mathbf{X}$ in the partial order. $\mathcal{P}_{\mathbf{T}}$ preserves the ancestral relations, so T_i must be a possible child of X_i while being in the same dc-component as $\mathbf{X}$ in $\mathcal{P}_{\mathbf{T}}$. This violates the criterion of the proposition, so T_i is not adjacent to any node in $\mathbf{X}$ in $\mathcal{M}$. By an argument similar to that of Lemma 8, we have $(T_i \perp\!\!\!\perp \mathbf{X}|\mathbf{T^{(i-1)}} \setminus \mathbf{X})$ in $\mathcal{M}$ and consequently in $\mathcal{D}_{\mathbf{T}}$. This implies that the conditional term of T_i under the full order is independent of $\mathbf{X}$ in $\mathcal{D}_{\mathbf{T}}$, and even if it was included in $S^{\mathbf{X}}$, the expression specific to $\mathcal{D}_{\mathbf{T}}$ can be rewritten as the one in Eq. 3. Therefore, the expression is unanimous and this concludes the proof.

(only if) Whenever the criterion fails, we show that we can construct a DAG in the equivalence class of $\mathcal{P}$ such that the effect is not identifiable. By Def. 7, the effect is not identifiable in $\mathcal{P}$. If the criterion is not satisfied, then some node $X_i \in \mathbf{X}$ is in the pc-component with a non-intervention possible child C in $\mathcal{P}_{\mathbf{T}}$. The edge between X_i and C is $X_i {*\!\!\rightarrow} C$ as C is outside the bucket of X_i. If the edge is not visible in $\mathcal{P}_{\mathbf{T}}$, then this edge is not visible in $\mathcal{P}$ (Lemma 4). Hence, we can construct a DAG $\mathcal{D}$ in the equivalence class of $\mathcal{P}$ where C is a child of X_i and the two nodes share a latent variable (Lemmas 24 and 14). The pair of sets $\mathbf{F} = \{X_i, C\}$ and $\mathbf{F}' = \{C\}$ form a hedge for $Q[\mathbf{T} \setminus \mathbf{X}]$ and the effect is not identifiable in $\mathcal{D}$ (Shpitser and Pearl, 2006, Theorem 4).

Otherwise, the edge $X_i \rightarrow C$ is visible in $\mathcal{P}_{\mathbf{T}}$. Then, there is a collider path between X_i and C consistent with Def. 4 such that the two nodes are in the same pc-component. Let $p{=}\langle X_i{=}T_0, T_1, \ldots, T_m{=}C\rangle$ denote the shortest such path in $\mathcal{P}_{\mathbf{T}}$. If the edge between X_i and T_1 is not into X_i, then T_1 is a possible child of X_i and the proof follows as in the previous case. Otherwise, we have $X_i \leftrightarrow T_1$. In this case, X_i is the only node along the path that belongs to $\mathbf{X}$. Suppose for contradiction that there is another node along the path which belongs to $\mathbf{X}$ and let T_i be one such node that is closest to X_i. We have $T_i \leftrightarrow T_{i+1}$ along the path and $T_{i+1} \neq C$, otherwise we $T_i \leftarrow\!*$ with a possibly directed path from T_i to C through X_i which is not possible (Lemma 22). Then, there is a bi-directed edge between T_{i+1} and every node in $\mathbf{X}$ (Lemma 16). This creates a shorter path than the initial one chosen composed of the sequence $\langle X_i, T_{i+1}, \ldots, C\rangle$ which is a contradiction. In $\mathcal{P}$, path p is present with $X_i \rightarrow C$ visible. Hence, we can construct a DAG $\mathcal{D}$ in the equivalence class of $\mathcal{P}$ such that C is a child of X_i and both are in the same c-component through a sequence of bi-directed edges along $\langle X_i, T_1, \ldots, C\rangle$ (Lemmas 24 and 14). The pair of sets $\mathbf{F} = \{X_i, T_1, \ldots, C\}$ and $\mathbf{F}' = \{T_1, \ldots, C\}$ form a hedge for $Q[\mathbf{T} \setminus \mathbf{X}]$ and the effect is not identifiable in $\mathcal{D}$. This concludes the proof. □

Lemma 8. *Given a PAG $\mathcal{P}$, a corresponding partial order $\mathbf{B}_1 < \cdots < \mathbf{B}_m$, a bucket $\mathbf{B}_i$, and $\boldsymbol{\mathcal{H}}$ a subset of the buckets in $\mathbf{B}^{(i-1)}$. Then, $\mathbf{B}_i$ is independent of $\boldsymbol{\mathcal{H}}$ given $\mathbf{B}^{(i-1)} \setminus \boldsymbol{\mathcal{H}}$, i.e. ($\mathbf{B}_i \perp\!\!\!\perp \boldsymbol{\mathcal{H}}|\mathbf{B}^{(i-1)} \setminus \boldsymbol{\mathcal{H}}$), if:*

1. *$\nexists(v_i \in \mathbf{B}_i \wedge v_j \in \boldsymbol{\mathcal{H}} \wedge v_i *\!\!-\!\!* v_j)$, i.e. there is no path of size one connecting $\mathbf{B}_i$ and $\boldsymbol{\mathcal{H}}$, and*
2. *Each proper definite status path between $\mathbf{B}_i$ and $\boldsymbol{\mathcal{H}}$:*
    (a) Contains at least one definite non-collider, or
    (b) The path is not into $\mathbf{B}_i$.

Proof. The first condition is necessary as it is not possible to separate two sets of nodes in a PAG if there exist $v_i \in \mathbf{B}_i$ and $v_j \in \boldsymbol{\mathcal{H}}$ such that v_i and v_j are adjacent. By (Zhang, 2006, Lemmas 5.1.3 and 5.1.7), we need to block the definite status paths between any $v_i \in \mathbf{B_i}$ and $v_j \in \boldsymbol{\mathcal{H}}$ for $\mathbf{B_i}$ and $\boldsymbol{\mathcal{H}}$ to be independent. But, any definite status path between v_i and v_j includes one of the proper definite status paths between $\mathbf{B_i}$ and $\boldsymbol{\mathcal{H}}$ as a subpath. Hence, we only consider blocking the proper definite status paths between $\mathbf{B_i}$ and $\boldsymbol{\mathcal{H}}$ since this will be sufficient to block all the definite status paths between any $v_i \in \mathbf{B_i}$ and $v_j \in \boldsymbol{\mathcal{H}}$. For simplicity, we assume that a path starts from $\mathbf{B_i}$ to $\boldsymbol{\mathcal{H}}$. Consider any proper definite status path between $\mathbf{B_i}$ and $\boldsymbol{\mathcal{H}}$; $p = \langle O_1, \ldots, O_n\rangle$ where $n > 2$, $O_1 \in \mathbf{B_i}$, $O_n \in \boldsymbol{\mathcal{H}}$.

1. The path contains at least one definite non-collider denoted $O_l \in \mathbf{B_k}$. If $k < i$, then the path is blocked as $O_l \in \mathbf{B}^{(i-1)} \setminus \boldsymbol{\mathcal{H}}$. If $k > i$, then O_l is in a bucket that comes after $\mathbf{B_i}$ in the partial order. Without loss of

generality, let $p_1 = \langle O_l, \ldots, O_1 \rangle$ be the subpath of p that is possibly out of O_l. By the partial order, there must exist a collider along p_1, else O_l is a possible ancestor of O_1. The collider is not in $\mathbf{B}^{(i-1)} \setminus \boldsymbol{\mathcal{H}}$ nor any of its possible descendants. Hence, the collider blocks p.

2. If the path is not into $\mathbf{B_i}$, then the edge mark incident on O_1 is either a tail ($O_1 \rightarrow O_2$) or a circle ($O_1 \circ\!\!\rightarrow O_2$). If a circle is incident on O_1, then the other end of the edge has to be an arrowhead ($O_1 \circ\!\!\rightarrow O_2$) since a circle would put O_2 in the same bucket as O_1 and the path would not be proper anymore. In both cases, O_2 is in a bucket after $\mathbf{B}_i$. The argument for the existence of a collider after $\mathbf{B_i}$ along path p follows similar to case a. Hence, path p is blocked as well.

Therefore, the set $\mathbf{B}^{(i-1)} \setminus \boldsymbol{\mathcal{H}}$ blocks all the proper paths between $\mathbf{B}_i$ and $\boldsymbol{\mathcal{H}}$. This concludes the proof of the lemma. $\square$

Lemma 9. *Let $\mathcal{M}$ be a MAG over $\mathbf{V}$, and $\mathcal{P}$ be the PAG of $\mathcal{M}$. For any $X, Y \in \mathbf{A} \subseteq \mathbf{V}$ and $\mathbf{Z} \subseteq \mathbf{A} \setminus \{X, Y\}$, if there is an m-connecting path between X and Y given $\mathbf{Z}$ in $\mathcal{M}_{\mathbf{A}}$, then there is a definite m-connecting path between X and Y given $\mathbf{Z}$ in $\mathcal{P}_{\mathbf{A}}$.*

Proof. Let $[\mathcal{M}]$ denote the equivalence class of $\mathcal{M}$ represented by $\mathcal{P}$, and $[\mathcal{M}]_{\mathbf{A}}$ denote the set of subgraphs over $\mathbf{A}$ of every MAG in $[\mathcal{M}]$. Every subgraph of a MAG is a MAG as well, so $\mathcal{M}_{\mathbf{A}}$ is a MAG. Let $\mathcal{P}_{[\mathcal{M}_{\mathbf{A}}]}$ be the PAG corresponding to $[\mathcal{M}_{\mathbf{A}}]$, the equivalence class of $\mathcal{M}_{\mathbf{A}}$. We have the following correspondence between $\mathcal{P}_{[\mathcal{M}_{\mathbf{A}}]}$ and $\mathcal{P}_{\mathbf{A}}$:

1. $\mathcal{P}_{[\mathcal{M}_{\mathbf{A}}]}$ and $\mathcal{P}_{\mathbf{A}}$ have the same adjacencies, since both have the same adjacencies of $\mathcal{M}_{\mathbf{A}}$.

2. $\mathcal{P}_{\mathbf{A}}$ subsumes the edge marks in $\mathcal{P}_{[\mathcal{M}_{\mathbf{A}}]}$: $\mathcal{P}_{\mathbf{A}}$ shows all the common edge marks in $[\mathcal{M}]_{\mathbf{A}}$ since $\mathcal{P}$ shows the common edge marks in $[\mathcal{M}]$. Also, every MAG in $[\mathcal{M}]_{\mathbf{A}}$ is in the equivalence class of $\mathcal{M}_{\mathbf{A}}$, otherwise we would have a contradiction in the equivalence class of $\mathcal{M}$. Hence, we have $[\mathcal{M}]_{\mathbf{A}} \subseteq [\mathcal{M}_{\mathbf{A}}]$. Therefore, any common mark in $[\mathcal{M}_{\mathbf{A}}]$ must be common in $[\mathcal{M}]_{\mathbf{A}}$, and $\mathcal{P}_{\mathbf{A}}$ subsumes the edge marks in $\mathcal{P}_{[\mathcal{M}_{\mathbf{A}}]}$.

If there is an m-connecting path between X and Y given $\mathbf{Z}$ in $\mathcal{M}_{\mathbf{A}}$, then there is a definite m-connecting path between X and Y given $\mathbf{Z}$ in $\mathcal{P}_{[\mathcal{M}_{\mathbf{A}}]}$. By the above two points, the corresponding path in $\mathcal{P}_{\mathbf{A}}$ is also definitely m-connecting. $\square$

Lemma 10. *Let $\mathcal{D}(\mathbf{V}, \mathbf{L})$ be a DAG, and $\mathcal{M}$ be the MAG of $\mathcal{D}$ over $\mathbf{V}$. Also, let $\mathcal{M}^*$ be the MAG over $\mathcal{D}_{\mathbf{A}}$, $\mathbf{A} \subseteq \mathbf{V}$. For any $X, Y \in \mathbf{A}$ and $\mathbf{Z} \subseteq \mathbf{A} \setminus \{X, Y\}$, if there is an m-connecting path between X and Y given $\mathbf{Z}$ in $\mathcal{M}^*$, then there is an m-connecting path between X and Y given $\mathbf{Z}$ in $\mathcal{M}_{\mathbf{A}}$.*

Proof. The proof for this lemma follows that of (Zhang, 2006, Lemma 5.2.1). To prove this result, it is sufficient to show that $\mathcal{M}_{\mathbf{A}}$ is Markov equivalent to a supergraph of $\mathcal{M}^*$. Let $\mathbf{L}'$ denote the subset of $\mathbf{L}$ in $\mathcal{D}_{\mathbf{A}}$. For any $X, Y \in \mathbf{A}$, if there is an inducing path between X and Y relative to $\mathbf{L}'$ in $\mathcal{D}_{\mathbf{A}}$, then this path is also present in $\mathcal{D}$. Therefore, X and Y are adjacent in both $\mathcal{M}^*$ and $\mathcal{M}_{\mathbf{A}}$.

If the edge between X and Y in $\mathcal{M}^*$ is directed, then there is a directed path between X and Y in $\mathcal{D}_{\mathbf{A}}$ and consequently in $\mathcal{D}$. Hence, the corresponding edge between X and Y in $\mathcal{M}_{\mathbf{A}}$ is directed as well.

If the edge between X and Y in $\mathcal{M}^*$ is bi-directed, then the corresponding edge in $\mathcal{M}_{\mathbf{A}}$ is either bi-directed or an invisible directed edge. A bi-directed edge in $\mathcal{M}^*$ implies that there is an inducing path between X and Y relative to $\mathbf{L}'$ in $\mathcal{D}_{\mathbf{A}}$ that is into both X and Y, and consequently that path between X and Y is present in $\mathcal{D}$ as well. If the corresponding edge in $\mathcal{M}_{\mathbf{A}}$ is directed, then it is not visible due to the previous property. Hence, the graphical condition for visibility in (Zhang, 2006, Lemma 5.1.3) is not applicable for this edge.

At this point, (1) every directed edge in $\mathcal{M}^*$ is present in $\mathcal{M}_{\mathbf{A}}$, and (2) every bi-directed edge in $\mathcal{M}^*$ is represented in $\mathcal{M}_{\mathbf{A}}$ by a bi-directed edge or an invisible directed edge. The rest of the proof follows the same argument as that of (Zhang, 2006, Lemma 5.2.1). Basically, we are able to transform $\mathcal{M}_{\mathbf{A}}$ into a supergraph of $\mathcal{M}^*$ by replacing the conflicting invisible directed edges with bi-directed edges while preserving equivalence (Zhang and Spirtes, 2005; Tian, 2005). $\square$

B.3 Theorem 3

Proof of Theorem 3. We show that the set Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) doesn't satisfy GAC relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathcal{P}$ (Lemma 25), whenever $\mathbf{IDP}$ fails to identify $P_{\mathbf{x}}(\mathbf{y})$. First, we make the following observations whenever $\mathbf{IDP}$ fails:

1. $\texttt{Identify}(\mathbf{D}_i, \mathbf{V}, P)$ must contain at least one intervention $X_i \in \mathbf{X}$ when the recursive call returns `Fail`. The reason is that all the added interventions in $\mathbf{V} \setminus \{\mathbf{D} \cup \mathbf{X}\}$ are not possible ancestors of $\mathbf{D}$ except through $\mathbf{X}$ by construction. Hence, if `Identify` intervenes on all interventions in $\mathbf{X}$, it should intervene on all the buckets in $\mathbf{V} \setminus \mathbf{D}$. Last, `Identify` doesn't fail to identify $Q[\mathbf{D}_i]$ from $Q[\mathbf{D}]$ as $\mathbf{D}_i$ is one of the cpc-components of $\mathcal{P}_{\mathbf{D}}$.

2. All the nodes remaining in a failed call $\texttt{Identify}(\mathbf{C}, \mathbf{T}, Q[\mathbf{T}])$ are possible ancestors of $\mathbf{Y}$ in $\mathcal{P}$. Every node in $\mathbf{T} \setminus \mathbf{C}$ is an ancestor of $\mathbf{C}$, otherwise the recursive routine wouldn't fail. Also, all the nodes $\mathbf{D} \supseteq \mathbf{C}$ in $\mathbf{IDP}$ are ancestors of $\mathbf{Y}$ by construction.

3. Whenever $\texttt{Identify}(\mathbf{C}, \mathbf{T}, Q[\mathbf{T}])$ fails, there exist $X, C \in \mathbf{T}$ such that $X \in \mathbf{X}$, $C \in \mathbf{D}$ ($\mathbf{D} = \texttt{An}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$) is a possible child of X, and X is in the same pc-component with C. By point 1, there exist at least one intervention X when `Identify` fails. Let $\mathbf{B}_1 < \cdots < \mathbf{B}_m$ be a partial order over $\mathcal{P}_{\mathbf{T}\cap\{\mathbf{V}\setminus\mathbf{D}\}}$. If $\mathbf{B}_m$ is a strict subset of a bucket in $\mathcal{P}_{\mathbf{T}}$, then $X, C \in \mathbf{B}_m$ by construction of $\mathbf{D}$. Otherwise, $\mathbf{B}_m$ is a bucket in $\mathcal{P}_{\mathbf{T}}$, $\mathbf{B}_m$ is in the same pc-component with a possible child, and the possible child is in $\mathbf{D}$.

Let $\texttt{Identify}(\mathbf{C}, \mathbf{T}, Q[\mathbf{T}])$ be a recursive call that throws `Fail` for $P_{\mathbf{x}}(\mathbf{y})$. By point 3 above, we have an intervention $X \in \mathbf{X}$ such that $X \in \mathbf{T}$, $C \in \mathbf{D}$ is a possible child of X, and X is in the same pc-component with C. If the edge between X and C is not a directed visible edge, then $\mathcal{P}$ is not amenable relative to ($\mathbf{X}$,$\mathbf{Y}$) (Def. 11), and consequently no set $\mathbf{Z}$ satisfies GAC (Def. 13) relative to ($\mathbf{X}$,$\mathbf{Y}$). Otherwise, we have a visible edge $X \to C$ and X is in the same pc-component with C, i.e. $p = \langle X{=}T_0 {*\!\!\rightarrow} T_1 \leftarrow - - \rightarrow T_{m-1} \leftarrow\!* T_m{=}C \rangle$ where $T_i \in \mathbf{T}$. Let T_i be the last node along p starting from T_0 such that $T_i \in \mathbf{X}$. Also, let T_j be the first node along p starting from T_{i+1} such that $T_j \in$ Forb($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) $\cup\, \mathbf{Y}$. Note that every node in $\mathbf{T}$ is in $\mathbf{X}$, $\mathbf{Y}$, Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$), or Forb($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) by point 2 above, and C is consistent with the description of T_j as it is in Forb($\mathbf{X}, \mathbf{Y}, \mathcal{P}$) by construction except for possibly not being the first. Based on this, we have the following exhaustive options for T_j:

1. Special case: $T_i = T_0$, $T_j = T_1$, the edge between T_i and T_j is possibly out of T_i, and $T_j \in \mathbf{Y}$ or T_j is along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$. Hence, $\mathcal{P}$ is not amenable relative to ($\mathbf{X}$,$\mathbf{Y}$) due to the invisible edge from T_i to T_j.

2. $T_j \in \mathbf{Y}$ or T_j is along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{P}$. Hence, every non-endpoint node along the subpath of p, denoted $p' = \langle T_i, \ldots, T_j \rangle$, is in Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$). If $T_j \in \mathbf{Y}$, then we have an m-connecting proper definite status non-causal path from an intervention $T_i \in \mathbf{X}$ to a node in $\mathbf{Y}$ (T_j). Consequently, Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) doesn't satisfy GAC relative to ($\mathbf{X}, \mathbf{Y}$) in $\mathcal{P}$. The other option is that T_j is along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$. By definition, there is a possibly directed path from T_j to $\mathbf{Y}$ that doesn't go through any node in $\mathbf{X}$, so there is an uncovered possibly directed path between them as well which doesn't go through $\mathbf{X}$ (Lemma 19). Hence, we can construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ (Lemma 23) such that the uncovered path is directed. The corresponding path of p' in $\mathcal{M}$ concatenated with the later directed path form an m-connecting proper non-causal path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{M}$. By Lemma 26, there is a proper definite status non-causal path from $\mathbf{X}$ to $\mathbf{Y}$ that is m-connecting given Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) in $\mathcal{P}$. Therefore, Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) doesn't satisfy GAC relative to ($\mathbf{X}, \mathbf{Y}$) in $\mathcal{P}$.

3. $T_j \notin \mathbf{Y}$ and T_j is not along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$, i.e. T_j is a descendant of a node that lies along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{P}$. Let T_k be the last node along p that has the same properties as T_j, i.e. T_k is a descendant of a node that lies along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{P}$, possibly $j = k$. When `Identify` fails, all nodes in $\mathbf{T} \setminus \mathbf{C}$ are possible ancestors of $\mathbf{C}$ in $\mathcal{P}_{\mathbf{T}}$, so every possibly directed path from T_k to $\mathbf{C}$ in $\mathcal{P}_{\mathbf{T}}$ is blocked by some intervention in $\mathbf{X}$. Let X_n be one such intervention in $\mathbf{X}$. By Lemma 19, there is an uncovered possibly directed path from T_k to X_n, denoted p_{kn}. Also, let T_p be the first node along p starting from T_{k+1} such that T_p is along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{P}$. There exist at least one node along p that is consistent with the description of T_p, i.e. C. Again, there is an uncovered possibly directed path from T_p to some $\mathbf{Y}$, denoted p_{py}. At this point, we have a few important remarks:

- Both paths p_{kn} and p_{py} are included in Forb($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) as the set is a descendant set.
- Along the subpath of p between T_k and T_p, denoted p_{kp}, all the non-endpoint nodes are colliders in $\mathcal{P}$ and are included in Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$)$\cup\mathbf{Y}$ due to the selection of T_k and T_p.
- T_k and T_p are not in the same circle component in $\mathcal{P}$, or else T_k is along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ as well.

Based on the above, we construct MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ using Lemma 23. We orient all edges possibly out of T_k and T_p in $\mathcal{P}$ out of them in $\mathcal{M}$, simultaneously. The latter is possible as both nodes are in different circle components. In $\mathcal{M}$, both paths p_{kn} and p_{py} are directed and don't intersect by any node, or else T_k is along a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{P}$ which contradicts our choice of T_k. If some node along p_{kp} is in $\mathbf{Y}$ (Y), then there is a proper non-causal path from $\mathbf{X}$ to $\mathbf{Y}$ composed of p_{kn} and the subpath of p_{kp} starting with T_k until Y. This path is m-connecting given Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) in $\mathcal{M}$. If no node along p_{kp} is in $\mathbf{Y}$, then the concatenated path composed of p_{nk}, p_{kp}, and p_{py} is a proper non-causal path from $\mathbf{X}$ to $\mathbf{Y}$. In both cases, we have a proper definite status non-causal path from $\mathbf{X}$ to $\mathbf{Y}$ that is m-connecting given Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) in $\mathcal{P}$ (Lemma 26). Therefore, Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) doesn't satisfy GAC relative to $(\mathbf{X}, \mathbf{Y})$ in $\mathcal{P}$.

The above cases exhaust all the options and Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{P}$) doesn't satisfy GAC relative to $(\mathbf{X}, \mathbf{Y})$ in $\mathcal{P}$ for every case. This concludes the proof. □

C Background Results

C.1 Lemmas 10 and 11 from (Tian, 2002, pp. 110)

Lemma 11. *Given a DAG $\mathcal{D}(\mathbf{V}, \mathbf{L})$, $\mathbf{W} \subseteq \mathbf{C} \subseteq \mathbf{V}$, and $\mathbf{W}' = \mathbf{C} \setminus \mathbf{W}$. If $\mathbf{W}$ is an ancestral set in the subgraph $\mathcal{D}_{\mathbf{C}}$ ($An(\mathbf{W})_{\mathcal{D}_{\mathbf{C}}} \cap \mathbf{C} = \mathbf{W}$), then:*

$$\sum_{\mathbf{W}'} Q[\mathbf{C}] = Q[\mathbf{W}]$$

Lemma 12. *Given a DAG $\mathcal{D}(\mathbf{V}, \mathbf{L})$. Let $\mathbf{H} \subseteq \mathbf{V}$ and assume that $\mathbf{H}$ is partitioned into c-components $\mathbf{H}_1, \ldots, \mathbf{H}_l$ in the subgraph $\mathcal{D}_{\mathbf{H}}$. Then we have*

1. *$Q[\mathbf{H}]$ decomposes as:*

$$Q[\mathbf{H}] = \prod_i Q[\mathbf{H}_i] \tag{8}$$

2. *Let a topological order of the nodes in $\mathbf{H}$ be $V_{h_1}, \ldots, V_{h_k}$ in $\mathcal{D}_{\mathbf{H}}$. Let $\mathbf{H}^{(i)} = \{V_{h_1}, \ldots, V_{h_i}\}$, and $\mathbf{H}^{(0)} = \emptyset$. Then, each $Q[\mathbf{H}_j]$ is computable from $Q[\mathbf{H}]$ and is given by*

$$Q[\mathbf{H}_j] = \prod_{\{i|V_{h_i} \in \mathbf{H}_j\}} \frac{Q[\mathbf{H}^{(i)}]}{Q[\mathbf{H}^{(i-1)}]} \tag{9}$$

where each $Q[\mathbf{H}^{(i)}] = \sum_{\mathbf{h} \setminus \mathbf{h}^{(i)}} Q[\mathbf{H}]$

C.2 Proof of soundness for PTO

Lemma 13. *The* PTO *algorithm is sound, i.e. the partial order is valid for all the DAGs in the equivalence class.*

Proof. In this context, we refer to the second step of the algorithm as "Bucketing" and to the third step as "Extraction". First, note that all the edges within a bucket in the partial order, if any, are circle edges (○–○). Initially, the buckets contain single variables and there are no edges within the buckets so our claim is valid. In the Bucketing step, We merge two buckets $\mathbf{B_i}$ and $\mathbf{B_j}$ when they have a circle edge between them. This means that there is a circle path between any two variables $X \in \mathbf{B_i}$ and $Y \in \mathbf{B_j}$. By lemma 16, an edge between X and Y if it exists would only be a circle edge. Hence, all the edges in a bucket are circle edges only.

If a bucket $\mathbf{B_i}$ has no tails or circles incident on it, then none of its variables are possible ancestors of any variable in the remaining buckets. This implies that no node in bucket $\mathbf{B_i}$ is an ancestor of any node in the other buckets in all the DAGs in the equivalence class. Hence, it is valid to put bucket $\mathbf{B_i}$ in front of all the remaining buckets in the topological order. This partial order between the buckets is valid in all the DAGs in the equivalence class of PAG $\mathcal{P}$. A problem arises if we don't have any bucket with no tail or circle incident on it. Next, we prove that such a case does not exist which concludes our proof.

In the partial order, we don't have circle edges (∘–∘) between the buckets. If there is an edge from variable X in bucket $\mathbf{B_i}$ into a variable Y in another bucket $\mathbf{B_j}$, then there is an edge from X into every variable in bucket $\mathbf{B_j}$. This follows from lemma 16 (2) since there is a circle path between Y and every other variable in bucket $\mathbf{B_j}$. This is an important observation for what follows.

Consider any iteration of the Extraction step of the algorithm. Assume every remaining bucket has one or more circles or tails incident on it, then there exist a possible directed cycle structure over the buckets as shown below with $2 \leq k \leq m$ where ? stands for a circle or tail. The existence of this structure in a PAG contradicts Lemma 22. Hence, such a possible directed cycle over the buckets can't exist and the assumption is invalid. Consequently, we have at least one bucket with no tails or circles incident on it at every iteration of the Extraction step. This concludes the proof.

$$\begin{tikzcd} B_{i_1}\,? \arrow[r] & B_{i_2}\,? \arrow[r, dashed] & B_{i_k} \arrow[ll, bend right, "?"'] \end{tikzcd}$$

□

C.3 Supporting results

Lemma 14 (Lemma 5.1.2 in Zhang, 2006)**.** *Let $\mathcal{M}$ be any MAG over a set of variables O, and $A \to B$ be any directed edge in $\mathcal{M}$. If $A \to B$ is invisible in $\mathcal{M}$, then there is a DAG whose MAG is $\mathcal{M}$ in which A and B share a latent parent, i.e., there exists a latent variable L_{AB} in the DAG.*

Lemma 15 (Lemma 3.3.1 in Zhang, 2006)**.** *For any three nodes A, B, C, if $A{*}\!\!\to B \circ\!\!-\!\!* C$, then there is an edge between A and C with an arrowhead at C, namely, $A{*}\!\!\to B$. Furthermore, if the edge between A and B is $A \to B$, then the edge between A and C is either $A \to C$ or $A\circ\!\!\to C$ (i.e., it is not $A \leftrightarrow C$).*

Lemma 16 (Lemma 3.3.2 in Zhang, 2006)**.** *In a PAG $\mathcal{P}$, for any two nodes A and B, if there is a circle path between A and B, i.e. a path consisting of ∘–∘ edges, then:*

1. *if there is an edge between A and B, then the edge is not into A or B, i.e. A ∘–∘ B in the absence of selection bias.*
2. *for any other node C, $C{*}\!\!\to A$ if and only if $C{*}\!\!\to B$. Furthermore, $C \leftrightarrow A$ if and only if $C \leftrightarrow B$.*

Lemma 17 (Lemma 20 in Zhang, 2008a)**.** *Let G be any DAG over $\mathbf{V} \cup \mathbf{L}$, and M be the MAG of G over $\mathbf{V}$. For any $A, B \in \mathbf{V}$ and $\mathbf{C} \subseteq \mathbf{V}$ that does not contain A or B, there is a path d-connecting A and B given $\mathbf{C}$ in G if and only if there is a path m-connecting A and B given $\mathbf{C}$ in M.*

Lemma 18 (Lemma 26 in Zhang, 2008a)**.** *Let M be a MAG over $\mathbf{V}$, and P be the PAG that represents the equivalence class of M. For any $A, B \in \mathbf{V}$ and $\mathbf{C} \subseteq \mathbf{V}$ that does not contain A or B, if there is a path m-connecting A and B given $\mathbf{C}$ in M, then there is a path definitely m-connecting A and B given $\mathbf{C}$ in P.*

Lemma 19 (Lemma B.1 in Zhang, 2008b)**.** *If $p = \langle A, \ldots, B\rangle$ is a possibly directed path from A to B in PAG $\mathcal{P}$, then some subsequence of p forms an uncovered possibly directed path from A to B in $\mathcal{P}$.*

Lemma 20 (Lemma B.2 in Zhang, 2008b)**.** *If p is an uncovered possibly directed path from A to B in PAG $\mathcal{P}$, then*

1. *if there is an $\circ\!\!\to$ edge on p, then any ∘–∘ edge on p is before that edge, and any $\to$ edge on p is after that edge; and*

2. *there is at most one $\circ\!\!\rightarrow$ edge on p.*

Lemma 21 (Theorem 2 in Zhang, 2008b)**.** *Let $\mathcal{M}$ be the MAG resulting from the following procedure applied to PAG $\mathcal{P}$:*

1. *orient the circles on $\circ\!\!\rightarrow$ edges in $\mathcal{P}$ as tails; and*
2. *orient the circle components of $\mathcal{P}$ into a DAG with no unshielded colliders.*

Then $\mathcal{M}$ is in the equivalence class of $\mathcal{P}$.

Lemma 22 (Lemma 7.5 in Maathuis and Colombo, 2015)**.** *Let X and Y be two distinct nodes in a PAG $\mathcal{P}$. Then $\mathcal{P}$ cannot have both a possibly directed path from X to Y and an edge of the form $Y *\!\!\rightarrow X$.*

Lemma 23 (Lemma 7.6 in Maathuis and Colombo, 2015)**.** *Let X be a node in a PAG $\mathcal{P}$. Let $\mathcal{M}$ be the MAG resulting from the following procedure applied to a $\mathcal{P}$:*

1. *replace all partially directed edges ($\circ\!\!\rightarrow$) in $\mathcal{P}$ with directed edges ($\rightarrow$), and*
2. *orient the subgraph of $\mathcal{P}$ consisting of all circle edges ($\circ\!\!-\!\!\circ$) into a DAG with no unshielded colliders and no new edges into X.*

Then, $\mathcal{M}$ is in the Markov equivalence class of $\mathcal{P}$.

Lemma 24 (lemma B.1 in Perković et al., 2016)**.** *Given a PAG $\mathcal{P}$ and X a node in $\mathcal{P}$. Let MAG $\mathcal{M}$ be in the equivalence class of $\mathcal{P}$ and satisfy the construction in lemma 23. Then, the edge $X \circ\!\!-\!\!\circ Y$, $X \circ\!\!\rightarrow Y$, or invisible $X \rightarrow Y$ in $\mathcal{P}$ is invisible $X \rightarrow Y$ in $\mathcal{M}$.*

The following Definitions, Theorem, and lemma summarize the generalized adjustment criterion introduced in (Perković et al., 2016) and introduced necessary results relevant to our work.

Definition 11 (Amenability)**.** *Let $\mathbf{X}$ and $\mathbf{Y}$ be disjoint node sets in a MAG or PAG $\mathcal{G}$. Then $\mathcal{G}$ is said to be amenable relative to ($\mathbf{X}$,$\mathbf{Y}$) if every possibly directed proper path from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{G}$ starts with a visible edge out of $\mathbf{X}$.*

Definition 12 (Forbidden set; Forb($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{G}$))**.** *Let $\mathbf{X}$ and $\mathbf{Y}$ be disjoint node sets in a MAG or PAG $\mathcal{G}$. Then the forbidden set relative to ($\mathbf{X}$,$\mathbf{Y}$) is defined as the set of nodes that are possible descendants of nodes $W \notin \mathbf{X}$ that lie along proper possibly directed paths from $\mathbf{X}$ to $\mathbf{Y}$ in $\mathcal{G}$.*

Definition 13 (Generalized adjustment criterion (GAC))**.** *Let $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ be pairwise disjoint node sets in a MAG or PAG $\mathbf{G}$. Then $\mathbf{Z}$ satisfies the generalized adjustment criterion relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathbf{G}$ if the following three conditions hold:*

*(**Amenability**) $\mathbf{G}$ is adjustment amenable relative to ($\mathbf{X}$,$\mathbf{Y}$), and*

*(**Forbidden set**) $\mathbf{Z} \cap$ Forb($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{G}$) $= \phi$, and*

*(**Blocking**) all proper definite status non-causal paths from $\mathbf{X}$ to $\mathbf{Y}$ are blocked by $\mathbf{Z}$ in $\mathcal{G}$.*

Theorem 4. *Let $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ be pairwise disjoint node sets in a MAG or PAG $\mathbf{G}$. Then $\mathbf{Z}$ is an adjustment set relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathbf{G}$ if and only if $\mathbf{Z}$ satisfies the generalized adjustment criterion relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathbf{G}$.*

Definition 14 (Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{G}$))**.** *Let $\mathbf{X}$ and $\mathbf{Y}$ be disjoint node sets in a MAG or PAG $\mathcal{G}$. We define Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{G}$) to be the set of possible ancestors of $\mathbf{X}$ and $\mathbf{Y}$ excluding $\mathbf{X}$, $\mathbf{Y}$, and Forb($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{G}$).*

Lemma 25 (corollary 4.4 in Perković et al., 2016)**.** *Let $\mathbf{X}$ and $\mathbf{Y}$ be disjoint node sets in a MAG or PAG $\mathcal{G}$. There exists an adjustment set relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathbf{G}$ if and only if Adjust($\mathbf{X}$,$\mathbf{Y}$,$\mathcal{G}$) satisfies the generalized adjustment criterion relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathbf{G}$.*

Lemma 26 (Lemma B.6 in Perković et al., 2016)**.** *Let $\mathbf{X}$,$\mathbf{Y}$ and $\mathbf{Z}$ be pairwise disjoint node sets in a PAG $\mathcal{P}$ and let $\mathcal{M}$ be a MAG in the equivalence class of $\mathcal{P}$. Let $\mathbf{Z}$ satisfy the amenability condition and the forbidden set condition relative to ($\mathbf{X}$,$\mathbf{Y}$) in $\mathcal{P}$. If there is a proper non-causal paths from $\mathbf{X}$ to $\mathbf{Y}$ that is m-connecting given $\mathbf{Z}$ in $\mathcal{M}$, then there is a proper definite status non-causal path from $\mathbf{X}$ to $\mathbf{Y}$ that is m-connecting given $\mathbf{Z}$ in $\mathcal{P}$.*

References

[Supp1] Marloes H. Maathuis and Diego Colombo. A generalized back-door criterion. *The Annals of Statistics*, 43(3): 1060–1088, 2015.

[Supp2] Jin Tian. Generating Markov equivalent maximal ancestral graphs by single edge replacement. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*, UAI'05, pages 591–598. AUAI Press, 2005.

[Supp3] Jiji Zhang and Peter Spirtes. A transformational characterization of Markov equivalence for directed acyclic graphs with latent variables. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*, UAI'05, pages 667–674. AUAI Press, 2005.