TECHNICAL REPORT
R-86
May, 2022

Causal Identification under Markov equivalence: Calculus, Algorithm, and Completeness

Amin Jaber
Purdue University
jaber0@purdue.edu

Adèle H. Ribeiro
Columbia University
adele@cs.columbia.edu

Jiji Zhang
Hong Kong Baptist University
zhangjiji@hkbu.edu.hk

Elias Bareinboim
Columbia University
eb@cs.columbia.edu

Abstract

One common task in many data sciences applications is to answer questions about the effect of new interventions, like: 'what would happen to Y if we make X equal to x while observing covariates $Z = z$?'. Formally, this is known as *conditional effect identification*, where the goal is to determine whether a post-interventional distribution is computable from the combination of an observational distribution and assumptions about the underlying domain represented by a causal diagram. A plethora of methods was developed for solving this problem, including the celebrated do-calculus [Pearl, 1995]. In practice, these results are not always applicable since they require a fully specified causal diagram as input, which is usually not available. In this paper, we assume as the input of the task a less informative structure known as a partial ancestral graph (PAG), which represents a Markov equivalence class of causal diagrams, learnable from observational data. We make the following contributions under this relaxed setting. First, we introduce a new causal calculus, which subsumes the current state-of-the-art, PAG-calculus. Second, we develop an algorithm for conditional effect identification given a PAG and prove it to be both sound and complete. In words, failure of the algorithm to identify a certain effect implies that this effect is not identifiable by any method. Third, we prove the proposed calculus to be complete for the same task.

1 Introduction

Despite the recent advances in AI and machine learning, the current generation of intelligent systems still lacks the pivotal ability to represent, learn, and reason with cause and effect relationships. The discipline of causal inference aims to 'algorithmitize' causal reasoning capabilities towards producing human-like machine intelligence and rational decision-making [Pearl and Mackenzie, 2018, Pearl, 2019, Bareinboim and Pearl, 2016]. One fundamental type of inference in this setting is concerned with the effect of new interventions, e.g., 'what would happen to outcome Y if X were set to x?' More generally, we may be interested in Y's distribution in a sub-population picked out by the value of some covariates $Z = z$'. For example, a legislator might be interested in the impact that increasing the minimum wage ($X = x$) has on profits (Y) in small businesses ($Z = z$), which is written in causal language as the interventional distribution $P(y|do(x), z)$, or $P_x(y|z)$. One method capable of answering such questions is through controlled experimentation [Fisher, 1951].

In many practical settings found throughout the empirical sciences, AI, and machine learning, it is not always possible to perform a controlled experiment due to ethical, financial, and technical considerations. This motivates the study of a problem known as *causal effect identification* [Pearl,

Preprint. Under review.

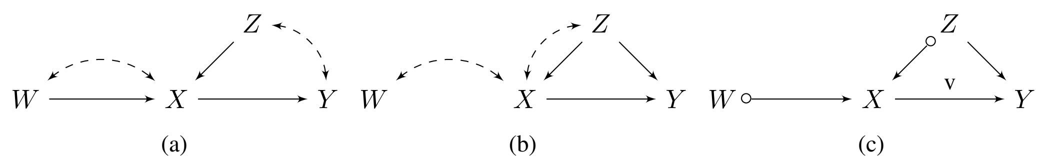

Figure 1: Sample causal diagrams (a,b) and the corresponding inferred PAG (c).

2000, Ch. 3]. The idea is to use the observational distribution $P(\mathbf{V})$ along with assumptions about the underlying domain, articulated in the form of a causal diagram $\mathcal{D}$, to infer the interventional distribution $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ when possible. For instance, Fig. 1a represents a causal diagram in which nodes correspond to measured variables, directed edges represent direct causal relations, and bidirected-dashed edges encode spurious associations due to unmeasured confounders. A plethora of methods have been developed to address the identification task including the celebrated causal calculus proposed by Pearl [1995] as well as complete algorithms [Tian, 2004, Shpitser and Pearl, 2006, Huang and Valtorta, 2006]. For instance, given the causal diagram in Fig. 1a and the query $P_x(y|z)$, the calculus sanctions the identity $P_x(y|z) = P(y|z, x)$. In words, the interventional distribution on the l.h.s. equates to the observational distribution on the right, which is available as input. Despite the power of these results, requiring the diagram as the input of the task is an Achilles heel for those methods, since background knowledge is usually not sufficient to pin down the single, true diagram.

To circumvent these challenges, a growing literature develops data-driven methods that attempt to learn the causal diagram from data first, and then perform identification from there. In practice, however, only an *equivalence class* (EC) of diagrams can be inferred from observational data without making substantial assumptions about the causal mechanisms [Verma, 1993, Spirtes et al., 2001, Pearl, 2000]. A prominent representation of this class is known as *partial ancestral graphs* (PAGs) [Zhang, 2008b]. Fig. 1c illustrates the PAG learned from observational data consistent with both causal diagrams in Figs. 1a and 1b since they are in the same Markov equivalence class. The directed edges in a PAG encode ancestral relations, not necessarily direct, and the circle marks stand for structural uncertainty. Directed edges labeled with v signify the absence of unmeasured confounders.

Causal effect identification in a PAG is usually more challenging than from a single diagram due to the structural uncertainties and the infeasibility of enumerating each member of the EC in most cases. The do-calculus was extended for PAGs to account for the inherent structure uncertainties without the need for enumeration [Zhang, 2007]. Still, the calculus falls short of capturing all identifiable effects as we will see in Sec. 3. On the other hand, it is computationally hard to decide whether there exists (and, if so, to find) a sequence of derivations in the generalized calculus to identify an effect of interest. In a more systematic manner, a complete algorithm has been developed to identify marginal effects (i.e., $P_{\mathbf{x}}(\mathbf{y})$) given a PAG [Jaber et al., 2019a]. This algorithm can be used to identify conditional effects whenever the joint distribution $P_{\mathbf{x}}(\mathbf{y} \cup \mathbf{z})$ is identifiable. Still, many conditional effects are identifiable even if the corresponding joint effect is not (Sec. 4.2). Finally, an algorithm to identify conditional effects has been proposed in [Jaber et al., 2019b], but not proven to be complete.[1]

In this paper, we pursue a data-driven formulation for the task of identification of any conditional causal effect from a combination of an observational distribution and the corresponding PAG (instead of a fully specified causal diagram). Accordingly, we makes the following contributions:

1. We propose a causal calculus for PAGs that subsumes the stat-of-the-art calculus introduced in [Zhang, 2007]. We prove the rules are *atomic complete*, i.e., a rule is not applicable in some causal diagram in the underlying EC whenever it is not applicable given the PAG.

2. Building on these results, we develop an algorithm for the identification of conditional causal effects given a PAG. We prove the algorithm is *complete*, i.e., the effect is not identifiable in some causal diagram in the equivalence class whenever the algorithm fails.

3. Finally, we prove the calculus is *complete* for the task of identifying conditional effects.

[1]Another 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.

2 Preliminaries

In this section, we introduce the basic setup and notations. Boldface capital letters denote sets of variables, while boldface lowercase letters stand for value assignments to those variables.[2]

Structural Causal Models. We use Structural Causal Models (SCMs) as our basic semantical framework [Pearl, 2000]. Formally, an SCM M is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U})\rangle$, where $\mathbf{U}$ is a set of exogenous (unmeasured) variables and $\mathbf{V}$ is a set of endogenous (measured) variables. $\mathbf{F}$ represents a collection of functions such that each endogenous variable $V_i \in \mathbf{V}$ is determined by a function $f_i \in \mathbf{F}$. Finally, $P(\mathbf{U})$ encodes the uncertainty over the exogenous variables. Every SCM is associated with one causal diagram where every variable in $\mathbf{V} \cup \mathbf{U}$ is a node, and arrows are drawn between nodes in accordance with the functions in $\mathbf{F}$. Following standard practice, we omit the exogenous nodes and add a bidirected dashed arc between two endogenous nodes if they share an exogenous parent. We only consider recursive systems, thus the corresponding diagram is acyclic. The marginal distribution induced over the endogenous variables $P(\mathbf{V})$ is called observational. The *d-separation* criterion captures the conditional independence relations entailed by a causal diagram in $P(\mathbf{V})$. For $\mathbf{C} \subseteq \mathbf{V}$, $Q[\mathbf{C}]$ denotes the post-intervention distribution of $\mathbf{C}$ under an intervention on $\mathbf{V} \setminus \mathbf{C}$, i.e. $P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{c})$.

Ancestral Graphs. We now introduce a graphical representation of equivalence classes of causal diagrams. A MAG represents a set of causal diagrams with the same set of observed variables that entail the same conditional independence and ancestral relations among the observed variables [Richardson and Spirtes, 2002]. *M-separation* extends d-separation to MAGs such that d-separation in a causal diagram corresponds to m-separation in its unique MAG over the observed variables, and vice versa.

Definition 1 (m-separation). *A path p between X and Y is active (or m-connecting) relative to $\mathbf{Z}$ ($X, Y \notin \mathbf{Z}$) if every non-collider on p is not in $\mathbf{Z}$, and every collider on p is an ancestor of some $Z \in \mathbf{Z}$. X and Y are m-separated by $\mathbf{Z}$ if there is no active path between X and Y relative to $\mathbf{Z}$.*

Different MAGs entail the same independence model and hence are Markov equivalent. A 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. A circle indicates an edge mark that is not invariant. A PAG is learnable from the independence model over the observed variables, and the FCI algorithm is a standard method to learn such an object [Zhang, 2008b]. In this work, an oracle for conditional independences is assumed to be available, which leads to the true PAG.

Graphical Notions. Given a 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. For a set of nodes $\mathbf{X}$, let $\texttt{An}(\mathbf{X})$ ($\texttt{De}(\mathbf{X})$) denote the union of $\mathbf{X}$ and the set of possible ancestors (descendants) of $\mathbf{X}$. 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 $\mathbf{Y}$. Let $\langle A, B, C\rangle$ be any consecutive triple along a path p. B is a collider on p if both edges are into B. B is a (definite) non-collider on p if one of the edges is out of B, or both edges have circle marks at B and there is no edge between A and C. A path is *definite status* if every non-endpoint node along it is either a collider or a non-collider. If the edge marks on a path between X and Y are all circles, we call the path a *circle path*. We refer to the closure of nodes connected with circle paths as a *bucket*.

A directed edge $X \to Y$ in a PAG is *visible* if there exists no causal diagram in the corresponding equivalence class where the relation between X and Y is confounded. Which directed edges are visible is easily decidable by a graphical condition [Zhang, 2008a], so we mark visible edges by v.

Manipulations in PAGs. Let $\mathcal{P}$ denote a PAG over $\mathbf{V}$ and $\mathbf{X} \subseteq \mathbf{V}$. $\mathcal{P}_{\mathbf{X}}$ denotes the induced subgraph of $\mathcal{P}$ over $\mathbf{X}$. The $\mathbf{X}$-lower-manipulation of $\mathcal{P}$ deletes all those edges that are visible in $\mathcal{P}$ and are out of variables in $\mathbf{X}$, replaces all those edges that are out of variables in $\mathbf{X}$ but are invisible in $\mathcal{P}$ with bi-directed edges, and otherwise keeps $\mathcal{P}$ as it is. The resulting graph is denoted as $\mathcal{P}_{\underline{\mathbf{X}}}$. The $\mathbf{X}$-upper-manipulation of $\mathcal{P}$ deletes all those edges in $\mathcal{P}$ that are into variables in $\mathbf{X}$, and otherwise keeps $\mathcal{P}$ as it is. The resulting graph is denoted as $\mathcal{P}_{\overline{\mathbf{X}}}$.

[2]A more comprehensive discussion about the background is provided in Appendix A.

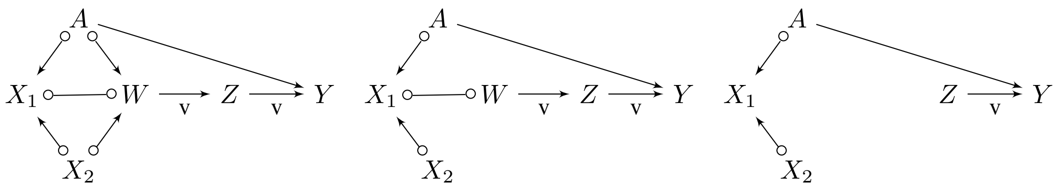

(a) PAG $\mathcal{P}$. (b) $\mathcal{P}_{\overline{W}}$. (c) $\mathcal{P}_{\mathbf{V}\setminus\{W\}}$.

Figure 3: Alternative methods to read ancestral relations under interventions from PAGs.

3 Causal Calculus for PAGs

The causal calculus introduced in [Pearl, 1995] is a seminal work that has been instrumental for understanding and eventually solving the task of effect identification from causal diagrams. Zhang [2007] generalized this result to the context of ancestral graphs, where a PAG is taken as the input of the task, instead of the specific causal diagram. In Sec. 3.1, we discuss Zhang's rules and try to understand the reasons they are insufficient to solve the identification problem in full generality. Further, in Sec.3.2, we introduce another generalization of the original calculus and prove that it is complete for atomic identification. This result will be further strengthened in subsequent sections.

3.1 Zhang's Calculus

An obvious extension of the m-separation criterion shown in Def. 1 to PAGs blocks all *possibly m-connecting paths*, as defined next.

Definition 2 (Possibly m-connecting path)**.** *In a PAG, a path p between X and Y is a possibly m-connecting path relative to a (possibly empty) set of nodes $\mathbf{Z}$ ($X, Y \notin \mathbf{Z}$) if every definite non-collider on p is not a member of $\mathbf{Z}$, and every collider on p is a possible ancestor of some member of $\mathbf{Z}$. X and Y are $\hat{m}$-separated by $\mathbf{Z}$ if there is no possibly m-connecting path between them relative to $\mathbf{Z}$.*

Using this notion of separation, Zhang [2007] proposed a calculus given a PAG as shown next.

Proposition 1 (Zhang's Calculus)**.** *Let $\mathcal{P}$ be the PAG over $\mathbf{V}$, and $\mathbf{X}, \mathbf{Y}, \mathbf{W}, \mathbf{Z}$ be disjoint subsets of $\mathbf{V}$. The following rules are valid, in the sense that if the antecedent of the rule holds, then the consequent holds in every MAG and consequently every causal diagram represented by $\mathcal{P}$.*

1. $P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are $\hat{m}$-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$.*
2. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are $\hat{m}$-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}, \underline{\mathbf{X}}}$.*
3. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are $\hat{m}$-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$.* *where* $\mathbf{X}(\mathbf{Z}) := \mathbf{X} \setminus \mathtt{PossAn}(\mathbf{Z})_{\mathcal{P}_{\overline{\mathbf{W}}}}$.

In words, rule 1 generalizes m-separation to interventional settings. Further, rule 2 licenses alternating a subset $\mathbf{X}$ between intervention and conditioning. Finally, rule 3 allows the adding/removal of an intervention $\mathbf{do}(\mathbf{X} = \mathbf{x})$. The next two examples illustrate the shortcomings of this result, where the first reveals the drawback of using Def. 2 to establish graphical separation and the second inspects evaluating $\mathbf{X}(\mathbf{Z})$ in rule 3 (where the notion of possible ancestors are evoked).

Example 1. *Consider the PAG $\mathcal{P}$ shown in Fig. 2. Since X and Y are not adjacent in $\mathcal{P}$, it is easy to show that X and Y are separable given $\{Z_1, Z_2\}$ in every causal diagram in the equivalence class. If rule 3 of Pearl's calculus is used in each diagram, then $P_x(y|z_1, z_2) = P(y|z_1, z_2)$. Further, applying rule 2 of do-calculus in each diagram, it's also the case that $P_x(y|z_1, z_2) = P(y|z_1, z_2, x)$. However, due to the possibly m-connecting path $\langle X, Z_1, Z_2, Y\rangle$, rules 3 and 2 in Prop. 1 are not applicable to $\mathcal{P}$. In other words, even though Pearl's calculus rules 2 and 3 are applicable to each diagram in the equivalence class, the same results cannot be established by Zhang's calculus.*

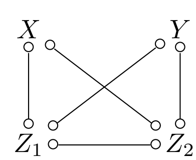

Figure 2: Sample PAG

Example 2. *Consider the PAG in Fig. 3a, and the evaluation on whether the equality $P_{w,x_1,x_2}(y|z) = P_w(y|z)$ holds. In order to apply rule 3 of Prop. 1, we need to evaluate whether $\{X_1, X_2\}$ is separated from $\{Y\}$ given $\{W, Z\}$ in the manipulated graph in Fig. 3b, which is not true in this case. However, the rule can be improved to be applicable in this case, as we will show later on (Sec. 3.2). The critical step will be the evaluation of the set $\mathbf{X}(\mathbf{Z})$ from $\mathcal{P}_{\overline{\mathbf{W}}}$.*

3.2 A New Calculus

Building on the analysis of the calculus proposed in [Zhang, 2007], we introduce next a set of rules centered around blocking *definite m-connecting paths*, as defined next.

Definition 3 (Definite m-connecting path)**.** *In a PAG, a path p between X and Y is a definite m-connecting path relative to a (possibly empty) set $\mathbf{Z}$ ($X, Y \notin \mathbf{Z}$) if p is definite status, every definite non-collider on p is not a member of $\mathbf{Z}$, and every collider on p is an ancestor of some member of $\mathbf{Z}$. X and Y are m-separated by $\mathbf{Z}$ if there is no definite m-connecting path between them relative to $\mathbf{Z}$.*

It is easy to see that every definite m-connecting path is a possibly m-connecting path, according to Def. 2; however, the converse is not true. For example, given the PAG in Figure 2, we have two definite status paths between X and Y. The first is $X \circ\!\!-\!\!\circ Z_1 \circ\!\!-\!\!\circ Y$ and the second is $X \circ\!\!-\!\!\circ Z_2 \circ\!\!-\!\!\circ Y$ where Z_1 and Z_2 are definite non-colliders. Given set $\mathbf{Z} = \{Z_1, Z_2\}$, $\mathbf{Z}$ blocks all definite status paths between X and Y. Alternatively, the path $X \circ\!\!-\!\!\circ Z_1 \circ\!\!-\!\!\circ Z_2 \circ\!\!-\!\!\circ Y$ is not definite status since Z_1, Z_2 are not colliders or non-colliders on this path. Hence, the path is a possibly m-connecting path relative to $\mathbf{Z}$ by Def. 2 but not a definite m-connecting path by Def. 3.

We are now ready to use this new definition and formulate a more powerful calculus.

Theorem 1. *Let $\mathcal{P}$ be the PAG over $\mathbf{V}$, and $\mathbf{X}, \mathbf{Y}, \mathbf{W}, \mathbf{Z}$ be disjoint subsets of $\mathbf{V}$. The following rules are valid, in the sense that if the antecedent of the rule holds, then the consequent holds in every MAG and consequently every causal diagram represented by $\mathcal{P}$.*[3]

1. $P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$.*
2. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$.*
3. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. where $\mathbf{X}(\mathbf{Z}) := \mathbf{X} \setminus \mathit{PossAn}(\mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{W}}}$.*

A few observations are important at this point. Despite the visual similarity to Prop. 1, there are two pivotal differences between these calculi. First, Thm. 1 only requires blocking the definite status paths, hence the use of 'm-separation' in Thm. 1 instead of '$\hat{m}$-separation'. Consider the PAG $\mathcal{P}$ in Fig. 2. We want to evaluate whether $P_x(y|z_1, z_2) = P(y|x, z_1, z_2)$ by applying Rule 2 in Theorem 1. Since all the edges in the PAG are circle edges, then $\mathcal{P}_{\underline{X}} = \mathcal{P}$. As discussed earlier, the set $\mathbf{Z} = \{Z_1, Z_2\}$ blocks all the definite status path between X and Y. Hence, X and Y are m-separated by $\mathbf{Z}$ and $P_x(y|z_1, z_2) = P(y|x, z_1, z_2)$ holds true by Rule 2.

Second, Thm. 1 defines $\mathbf{X}(\mathbf{Z})$ as the subset of $\mathbf{X}$ that is not in the possible ancestors of $\mathbf{Z}$ in $\mathcal{P}_{\mathbf{V}\setminus\mathbf{W}}$, as opposed to $\mathcal{P}_{\overline{\mathbf{W}}}$ in Prop. 1. We revisit the query in Ex. 2 to clarify this subtle but significant difference. Given the PAG in Fig. 3a, we want to evaluate whether $P_{w,x_1,x_2}(y|z) = P_w(y|z)$ by applying Rule 3 from Thm. 1 instead of Prop. 1. Fig. 3c shows $\mathcal{P}_{\mathbf{V}\setminus\{W\}}$ where $\mathbf{X} = \{X_1, X_2\}$ are not possible ancestors of Z. Therefore, $\mathbf{X}(Z) = \mathbf{X}$, the edges into X_1 are cut in $\mathcal{P}_{\overline{W},\overline{\mathbf{X}(Z)}}$, and $\mathbf{X}$ and Y are m-separated therein.

Third, the proof of the Theorem 1 is provided in the appendix, but it follows from the relationship between m-connecting path in a manipulated MAG to a definite m-connecting path in the corresponding manipulated PAG. It was conjectured in [Zhang, 2008a, Footnote 15] that, for $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, if there is an m-connecting path in $\mathcal{M}_{\overline{\mathbf{Y}},\underline{\mathbf{X}}}$, then there is a definite m-connecting path in $\mathcal{P}_{\overline{\mathbf{Y}},\underline{\mathbf{X}}}$. In this work, we prove that conjecture to be true for the special class of manipulations required in the rules of the calculus. Finally, the next proposition establishes the necessity of the antecedents in Thm. 1 in order to apply the corresponding rule given every diagram in the equivalence class.

[3] All the proofs can be found in the appendix.

Theorem 2 (Atomic Completeness)**.** *The calculus in Theorem 1 is atomically complete; meaning, whenever a rule is not applicable given a PAG, then the corresponding rule in Pearl's calculus is not applicable given some causal diagram in the Markov equivalence class.*

For instance, considering PAG $\mathcal{P}$ in Fig. 1c, we note that $(Z \not\perp\!\!\!\perp Y | X)_{\mathcal{P}_{\overline{X},\underline{Z}}}$, which means rule 2 is not applicable. Clearly, the diagram in Fig. 1a is in the equivalence class of $\mathcal{P}$ and the corresponding rule of Pearl's calculus is not applicable due to the latent confounder between Z and Y.

4 Effect Identification: A Complete Algorithm

It is challenging to use the calculus rules in Thm. 1 to identify causal effects since it is computationally hard to decide whether there exists (and, if so, to find) a sequence of derivations in the generalized calculus to identify an effect of interest. The goal of this section is to formulate an algorithm to identify conditional causal effects. The next definition formalizes the notion of identifiability from a PAG, generalizing the causal-diagram-specific notion introduced in [Pearl, 2000, Tian, 2004].

Definition 4 (Causal-Effect Identifiability)**.** *Let* $\mathbf{X}, \mathbf{Y}, \mathbf{Y}$ *be disjoint sets of endogenous variables,* $\mathbf{V}$*. The causal effect of* $\mathbf{X}$ *on* $\mathbf{Y}$ *conditioned on* $\mathbf{Z}$ *is said to be identifiable from a PAG* $\mathcal{P}$ *if the quantity* $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ *can be computed uniquely from the observational distribution* $P(\mathbf{V})$ *given every causal diagram* $\mathcal{D}$ *in the Markov equivalence class represented by* $\mathcal{P}$*.*

The remainder of the section is organized as follows. Sec. 4.1 introduces a version of the **IDP** algorithm [Jaber et al., 2019a] to identify marginal causal effects. The attractiveness of this version is that it yields simpler expressions whenever the effect is identifiable while preserving the same expressive power, i.e., completeness for marginal identification. Sec. 4.2 utilizes the new algorithm along with the calculus in Thm. 1 to formulate a complete algorithm for conditional identification.

4.1 Marginal Effect Identification

We introduce the notion of pc-component next, which generalizes the notion of c-component that is instrumental to solve identificaiton problems in a causal diagram [Tian and Pearl, 2002].

Definition 5 (PC-Component)**.** *In a PAG, or any induced subgraph thereof, two nodes are in the same possible c-component (pc-component) if there is a path between them such that (1) all non-endpoint nodes along the path are colliders, and (2) none of the edges is visible.*

Following Def. 5, e.g., W and Z in Fig. 1c are in the same pc-component due to $W {\circ\!\!\rightarrow} X \leftarrow\!\!\circ Z$. By contrast, X, Y are not in the same pc-component since the direct edge between them is visible and Z along $\langle X, Z, Y\rangle$ is not a collider. Building on pc-components, we define the key notion of regions.

Definition 6 (Region $\mathcal{R}^{\mathbf{C}}_{\mathbf{A}}$)**.** *Given PAG* $\mathcal{P}$ *over* $\mathbf{V}$*, and* $\mathbf{A} \subseteq \mathbf{C} \subseteq \mathbf{V}$*. The region of* $\mathbf{A}$ *with respect to* $\mathbf{C}$*, denoted* $\mathcal{R}^{\mathbf{C}}_{\mathbf{A}}$*, is the union of the buckets that contain nodes in the pc-component of* $\mathbf{A}$ *in* $\mathcal{P}_{\mathbf{C}}$*.*

A region expands a pc-component and will prove to be useful in the identification algorithm. For example, the pc-component of X in Fig. 2 is $\{X, Z_1, Z_2\}$ and the region $\mathcal{R}^{\mathbf{V}}_{X} = \{X, Z_1, Z_2, Y\}$. Building further on these definitions and the new calculus, we derive a new identification criterion.

Proposition 2. *Let* $\mathcal{P}$ *denote a PAG over* $\mathbf{V}$*,* $\mathbf{T}$ *be a union of a subset of the buckets in* $\mathcal{P}$*, and* $\mathbf{X} \subset \mathbf{T}$ *be a bucket. Given* $P_{\mathbf{v}\setminus\mathbf{t}}$ *(i.e., an observational expression for* $Q[\mathbf{T}]$*),* $Q[\mathbf{T} \setminus \mathbf{X}]$ *is identifiable by the following expression if, in* $\mathcal{P}_{\mathbf{T}}$*,* $C^{\mathbf{X}} \cap \texttt{PossDe}(\mathbf{X}) \subseteq \mathbf{X}$*, where* $C^{\mathbf{X}}$ *is the pc-component of* $\mathbf{X}$*.*

$$Q[\mathbf{T} \setminus \mathbf{X}] = \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{P_{\mathbf{v}\setminus\mathbf{t}}(\mathbf{X}|\mathbf{T} \setminus \texttt{PossDe}(\mathbf{X}))} \tag{1}$$

Note the interventions are over buckets which may or may not be single nodes. Since there is little to no causal information inside a bucket, marginal effects of interventions over subsets of buckets are not identifiable. Also, the input distribution is possibly interventional which licenses recursive applications of the criterion. The next example illustrates the power of the new criterion.

Example 3. *Consider PAG* $\mathcal{P}$ *in Fig. 3a and the query* $P_{x_1,x_2,w}(y, z, a)$*. Starting with the observational distribution* $P(\mathbf{V})$ *as input, let* $\mathbf{T} = \mathbf{V}$ *and* $\mathbf{X} = \{X_1, W\}$*. We have* $C^{\mathbf{X}} = \{X_1, W, A, X_2\}$*,*

Algorithm 1 $\mathbf{IDP}(\mathcal{P}, \mathbf{x}, \mathbf{y})$

Input: PAG $\mathcal{P}$ and two disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$
Output: Expression for $P_{\mathbf{x}}(\mathbf{y})$ or `FAIL`

1: Let $\mathbf{D} = \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$
2: **return** $\sum_{\mathbf{d}\setminus\mathbf{y}}$ IDENTIFY($\mathbf{D}, \mathbf{V}, P$)

3: **function** IDENTIFY($\mathbf{C}$, $\mathbf{T}$, $Q = Q[\mathbf{T}]$)
4:     **if** $\mathbf{C} = \emptyset$ **then return** 1
5:     **if** $\mathbf{C} = \mathbf{T}$ **then return** Q
/* *In* $\mathcal{P}_{\mathbf{T}}$*, let* $\mathbf{B}$ *denote a bucket, and let* $C^{\mathbf{B}}$ *denote the pc-component of* $\mathbf{B}$ */
6:     **if** $\exists \mathbf{B} \subset \mathbf{T} \setminus \mathbf{C}$ such that $C^{\mathbf{B}} \cap \texttt{PossDe}(\mathbf{B})_{\mathcal{P}_{\mathbf{T}}} \subseteq \mathbf{B}$ **then**
7:         Compute $Q[\mathbf{T} \setminus \mathbf{B}]$ from Q; ▷ via Prop. 2
8:         **return** IDENTIFY($\mathbf{C}$,$\mathbf{T} \setminus \mathbf{B}$,$Q[\mathbf{T} \setminus \mathbf{B}]$)
9:     **else if** $\exists \mathbf{B} \subset \mathbf{C}$ such that $\mathcal{R}_{\mathbf{B}} \neq \mathbf{C}$ **then** ▷ $\mathcal{R}_{\mathbf{B}}$ is equivalent to $\mathcal{R}^{\mathbf{C}}_{\mathbf{B}}$
10:         **return** $\dfrac{\text{IDENTIFY}(\mathcal{R}_{\mathbf{B}},\mathbf{T},Q) \times \text{IDENTIFY}(\mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_{\mathbf{B}}},\mathbf{T},Q)}{\text{IDENTIFY}(\mathcal{R}_{\mathbf{B}} \cap \mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_{\mathbf{B}}},\mathbf{T},Q)}$
11:     **else throw** `FAIL`

$\texttt{PossDe}(\mathbf{X}) = \{X_1, W, Z, Y\}$, and $C^{\mathbf{X}} \cap \texttt{PossDe}(\mathbf{X}) = \mathbf{X}$. Hence, the criterion in Prop. 2 is applicable and we have $P_{\mathbf{x}}(y, z, a, x_2) = \frac{P(\mathbf{V})}{P(x_1, w|a, x_2)} = P(a, x_2) \times P(y, z|a, w)$ after simplification. Next, we consider intervening on X_2 given $P_{x_1,w}(y, z, a, x_2)$. Notice X_2 is disconnected from the other nodes in $\mathcal{P}_{\mathbf{V}\setminus\{X_1,W\}}$ and it trivially satisfies the criterion in Prop. 2. Therefore, we get the expression $P_{x_1,x_2,w}(y, z, a) = \frac{P_{x_1,w}(y,z,a,x_2)}{P_{x_1,w}(x_2|y,z,a)} = P(a) \times P(y, z|w, a)$ after simplification.

A more general criterion was introduced in [Jaber et al., 2018a, Thm. 1] based on the *possible children* of the intervention bucket $\mathbf{X}$ instead of the possible descendants. However, the corresponding expression is convoluted and usually large, which could be intractable even if the effect is identifiable. Alg. 1 shows the proposed version of **IDP**, which builds on the new criterion (Prop. 2). Specifically, the key difference between this algorithm and the one proposed in [Jaber et al., 2019a] is in Lines 6-7, where the criterion in Prop. 2 is used as opposed to that in [Jaber et al., 2018a, Thm. 1]. Interestingly enough, the new criterion is "just right," namely, it is also sufficient to obtain a complete algorithm for marginal effect identification, as shown in the next result.[4]

Theorem 3 (completeness)**.** *Alg. 1 is complete for identifying marginal effects $P_{\mathbf{x}}(\mathbf{y})$. Moreover, the calculus in Thm. 1, together with standard probability manipulations are complete for the same task.*

4.2 Conditional Effect Identification

We start by making a couple of observations, and then build on those observations to formulate an algorithm to identify conditional causal effects. The proposed algorithm leverages the calculus in Thm. 1 and the **IDP** algorithm in Alg. 1. Obs. 1 notes that a conditional effect $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ can be rewritten as $\frac{P_{\mathbf{x}}(\mathbf{y},\mathbf{z})}{\sum_{\mathbf{y}} P_{\mathbf{x}}(\mathbf{y},\mathbf{z})}$, and hence it is identifiable if $P_{\mathbf{x}}(\mathbf{y}, \mathbf{z})$ is identifiable by Alg. 1.

Observation 1 (Marginal Effect)**.** *Consider PAG $\mathcal{P}_1$ in Fig. 4a where the goal is to identify the causal effect $P_b(a, c|d)$. We notice that the effect $P_b(a, c, d)$ is identifiable using the **IDP** algorithm. Let $E \coloneqq P(a, d) \times P(c|b, d)$ denote the expression for the marginal effect $P_b(a, c, d)$ which can be obtained from **IDP**. Consequently, the target effect can be computed using the expression $E/\sum_{a,c} E$.*

Whenever the marginal effect $P_{\mathbf{x}}(\mathbf{y}, \mathbf{z})$ is not identifiable using Alg. 1, Observations 2 and 3 propose techniques to identify the conditional effect using the calculus in Thm. 1, namely rule 2. Obs. 2 uses rule 2 of Thm. 1, when applicable, to move variables from the conditioning to the intervention set. The marginal effect of the resulting conditional query turns out to be identifiable, and consequently does the conditional effect. We note that the work in [Shpitser and Pearl, 2006] uses the same trick to formulate an algorithm for conditional effect identification given a causal diagram.

[4] A more detailed comparison of the two algorithms along with illustrative examples is provided in Appx. F.

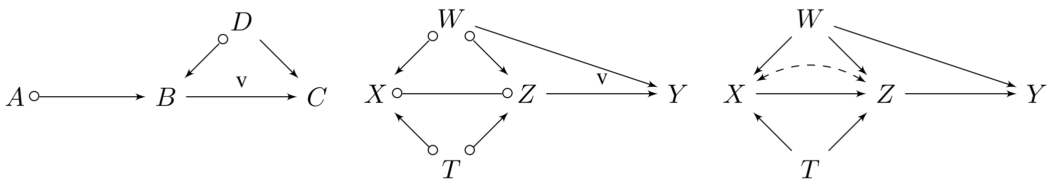

(a) PAG $\mathcal{P}_1$. (b) PAG $\mathcal{P}_2$ (c) Causal Diagram $\mathcal{G}$.

Figure 4: (a,b) Two sample PAGs, and (c) a causal diagram in the equivalence class of (b).

Algorithm 2 **CIDP**$(\mathcal{P}, \mathbf{x}, \mathbf{y}, \mathbf{z})$

Input: PAG $\mathcal{P}$ and three disjoint sets $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \subset \mathbf{V}$
Output: Expression for $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ or `FAIL`

1: $\mathbf{D} \leftarrow \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$
/* *Let* $\mathbf{B}_1, \ldots, \mathbf{B}_m$ *denote the buckets in* $\mathcal{P}$ */
2: **while** $\exists \mathbf{B}_i$ s.t. $\mathbf{B}_i \cap \mathbf{D} \neq \emptyset \wedge \mathbf{B}_i \not\subseteq \mathbf{D}$ **do**
3: $\mathbf{X}' \leftarrow \mathbf{B}_i \cap \mathbf{X}$
4: **if** $(\mathbf{X}' \perp\!\!\!\perp \mathbf{Y} | (\mathbf{X} \setminus \mathbf{X}') \cup \mathbf{Z})_{\mathcal{P}_{\overline{\mathbf{X}\setminus\mathbf{X}'},\underline{\mathbf{X}'}}}$ **then**
5: $\mathbf{x} \leftarrow \mathbf{x} \setminus \mathbf{x}'$; $\mathbf{z} \leftarrow \mathbf{z} \cup \mathbf{x}'$ ▷ Apply rule 2 of Thm. 1
6: $\mathbf{D} \leftarrow \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$
7: **else throw** `FAIL`
/* *Let* $\mathbf{Z}_1, \ldots, \mathbf{Z}_m$ *partition* $\mathbf{Z}$ *such that* $\mathbf{Z}_i := \mathbf{Z} \cap \mathbf{B}_i$ */
8: **while** $\exists \mathbf{Z}_i$ s.t. $(\mathbf{Z}_i \perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup (\mathbf{Z} \setminus \mathbf{Z}_i))_{\mathcal{P}_{\overline{\mathbf{X}},\underline{\mathbf{Z}_i}}}$ **do**
9: $\mathbf{x} \leftarrow \mathbf{x} \cup \mathbf{z}_i$; $\mathbf{z} \leftarrow \mathbf{z} \setminus \mathbf{z}_i$ ▷ Apply rule 2 of Thm. 1
10: $E \leftarrow$ **IDP**$(\mathcal{P}, \mathbf{x}, \mathbf{y} \cup \mathbf{z})$
11: **return** $E / \sum_{\mathbf{y}} E$

Observation 2 (Flip Observations to Interventions)**.** *Consider PAG $\mathcal{P}_1$ in Fig. 4a and the causal query $P_a(c|b, d)$. Unlike the case in Obs. 1, the marginal effect $P_a(b, c, d)$ is not identifiable by the **IDP** algorithm. Using rule 2 of Thm. 1, we have $(B \perp\!\!\!\perp C|D)_{\mathcal{P}_{\overline{A},\underline{B}}}$ and we move B from conditioning to intervention, i.e., $P_a(c|b, d) = P_{a,b}(c|d)$. The marginal effect $P_{a,b}(c, d)$ is identifiable by **IDP** and we get the expression $E := P(d) \times P(c|b, d)$. Hence, we have $P_a(c|b, d) = P_{a,b}(c|d) = E / \sum_c E$.*

Finally, Obs. 3 comes as a surprise since it requires flipping interventions to observations, contrary to Obs. 2. A key graphical structure in the PAG that requires such a treatment is the presence of a proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y} \cup \mathbf{Z}$ that starts with a circle edge.

Observation 3 (Flip Interventions to Observations)**.** *We revisit the query in Example. 1. First, the marginal effect $P_x(y, z_1, z_2)$ is not identifiable by the **IDP** algorithm. Also, we cannot use rule 2 in Thm. 1 to flip Z_1 or Z_2 into interventions since they are both adjacent to Y with a circle edge (∘—∘). However, we can use rule 2 to flip X to the conditioning set since there are no definite m-connecting paths between X and Y given $\{Z_1, Z_2\}$ in the PAG. Hence, we obtain $P_x(y|z_1, z_2) = P(y|z_1, z_2, x)$. Alternatively, consider PAG $\mathcal{P}_2$ in Fig. 4b with the causal query $P_x(y|z)$. We cannot use rule 2 to flip X to an observation since $\langle X, W, Y \rangle$ is active given Z in $\mathcal{P}_{2\underline{X}}$. In fact, the causal diagram $\mathcal{G}$ in Fig. 4c is in the equivalence class of $\mathcal{P}_2$ and such that $P_x(y|z)$ is provably not identifiable [Shpitser and Pearl, 2006, Corol. 2]. Hence, the effect is not identifiable given $\mathcal{P}_2$ according to Def. 4.*

Putting these observations together, we formulate the **CIDP** algorithm (Alg. 2) for identifying conditional causal effects given a PAG. The algorithm is divided into three phases. In Phase I (lines 1-7), Obs. 3 is used to check for proper possibly directed paths from $\mathbf{X}$ to $\mathbf{Y} \cup \mathbf{Z}$ that start with a circle edge. This is checked algorithmically by computing $\mathbf{D} = \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$, iteratively, and checking if some bucket $\mathbf{B}_i$ in $\mathcal{P}$ intersects with, but is not a subset of, $\mathbf{D}$. If such a bucket exists, **CIDP** flips $\mathbf{B}_i \cap \mathbf{X}$ from interventions to observations using rule 2, when applicable, else the algorithm throws a fail and the effect is not computable. In Phase II (lines 8-9), Obs. 2 is used to flip the subset of observations in each bucket into interventions by applying rule 2 of Thm. 1,

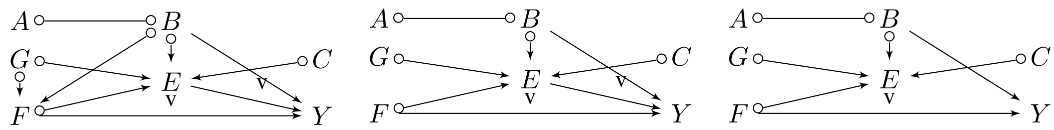

(a) PAG $\mathcal{P}$. (b) $\mathcal{P}_{\overline{F},\underline{A}}$ (c) $\mathcal{P}_{\overline{F},\underline{E}}$

Figure 5: Illustrating example for Alg. 2.

whenever applicable. Finally, in Phase III (line 10), the marginal effect $P_{\mathbf{x}}(\mathbf{y} \cup \mathbf{z})$ is computed from the modified sets $\mathbf{X}$ and $\mathbf{Z}$, using the **IDP** algorithm in Alg. 1. If the call is successful, an expression for the conditional effect is returned at line 11. The example below illustrates **CIDP** in action. An empirical evaluation of **CIDP** is provided in App. G; the R package will be made available.

Example 4. *Consider PAG $\mathcal{P}$ in Fig. 5a and the conditional query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) := P_{a,f}(y|b,e)$. In Phase I, we have $\mathbf{D} = \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}} = \{Y, B, E, C, G\}$, and the bucket $\{A, B\}$ satisfies the conditions at line 2 since $A \notin \mathbf{D}$. In $\mathcal{P}_{\overline{F},\underline{A}}$ (as shown in Fig. 5b), $\mathbf{X}' = \{A\}$ is m-separated from Y given $\{B, E, F\}$ which satisfies the if condition at line 4. Hence, we flip A to the conditioning set $\mathbf{Z}$ via rule 2 of Thm. 1 to obtain the updated query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) = P_f(y|a, b, e)$. In Phase II (lines 8-9), let $\mathbf{Z}_1 = \{E\}$ and $\mathbf{Z}_2 = \{A, B\}$. In $\mathcal{P}_{\overline{F},\underline{E}}$ (see Fig. 5c), we have E m-separated from Y given $\{F\} \cup \mathbf{Z_2}$ which satisfies the if condition at line 8. Hence, we flip E to the intervention set using rule 2 of Thm. 1 and we get the updated query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) = P_{e,f}(y|a, b)$. Next, we check if $\mathbf{Z}_2$ is m-separated from Y given $\mathbf{Z}_1$ in $\mathcal{P}_{\overline{\mathbf{Z}_1},\underline{\mathbf{Z}_2}}$ which does not hold due to a bidirected edge between B and E. Hence, rule 2 is not applicable and $\mathbf{Z}_2$ remains in the conditioning set. Finally, we call **IDP** in Alg. 1 to compute the marginal effect $P_{e,f}(y, a, b)$, if possible. The effect is identifiable with the simplified expression $P(y|b, e, f) \times P(a, b)$. Hence, $P_{a,f}(y|b, e) = P_{e,f}(y|a, b) = \frac{P(y|b,e,f) \times P(a,b)}{\sum_y P(y|b,e,f) \times P(a,b)} = P(y|b, e, f)$.*

The soundness of Alg. 2 follows from that of Alg. 1 and Thm. 1. Next, we turn to its completeness. According to Def. 4, whenever **CIDP** fails, we need to establish one of two conditions for completeness. Either there exist two causal diagrams in the equivalence class with different identifications, or the effect is not identifiable in some causal diagram according to the criterion in [Shpitser and Pearl, 2006, Corol. 2]. Thm. 4 establishes completeness by proving that the latter is always the case. This result along with the completeness of the calculus rules for the identification of marginal effects (see Thm. 3) implies that the rules are complete for conditional effects as well.

Theorem 4 (completeness)**.** *Alg. 2 is complete for identifying conditional effects $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$. Also, the calculus in Thm. 1, together with standard probability manipulations are complete for the same task.*

5 Conclusions

In this work, we investigate the problem of identifying conditional interventional distributions given a Markov equivalence class of causal diagrams represented by a PAG. We introduce a new generalization of the do-calculus for identification of interventional distributions in PAGs (Thm. 1) and show it to be atomically complete (Thm. 2). Building on these results, we develop the **CIDP** algorithm (Alg. 2), which is both sound and complete, i.e., it identifies any conditional effects of the form $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ that is identifiable (Thm. 4). Finally, we show that the new calculus rules, along with standard probability manipulations, are complete for the same task. These results close the problem of effect identification under Markov equivalence in that they completely delineate the theoretical boundaries of what is, in principle, computable from a certain data collection. We expect the newly introduced machinery to help data scientists to identify novel effects in real world settings.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section ...
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.

- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] The three contributions claimed in the abstract and reiterated at the end of the introduction are reflected in Section 3 for contribution 1, and in Section 4 for contributions 2 and 3.
   (b) Did you describe the limitations of your work? [Yes] See Section 2, lines 103–104.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] The assumption in lines 103-104 leads to the true PAG which is what the theoretical results results are based on.
   (b) Did you include complete proofs of all theoretical results? [Yes] The proofs can be found in the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The code will be made freely available with the camera-ready version of the paper in case of acceptance.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Extended Preliminaries

In this section, we introduce the basic setup and notations. Boldface capital letters denote sets of variables, while boldface lowercase letters stand for value assignments to those variables.

Structural Causal Models. We use Structural Causal Models (SCMs) [Pearl, 2000, pp. 204-207] as our basic semantical framework. Formally, an SCM M is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U}) \rangle$, where $\mathbf{U}$ is a set of exogenous (unmeasured) variables and $\mathbf{V}$ is a set of endogenous (measured) variables. $\mathbf{F}$ represents a collection of functions such that each endogenous variable $V_i \in \mathbf{V}$ is determined by a function $f_i \in \mathbf{F}$, where f_i is a mapping from the respective domain of $\mathbf{U}_i \cup \mathbf{Pa}_i$ to V_i, $\mathbf{U}_i \subseteq \mathbf{U}$, $\mathbf{Pa}_i \subseteq \mathbf{V} \setminus V_i$. The uncertainty is encoded through a probability distribution over the exogenous variables, $P(\mathbf{U})$. Every SCM is associated with one causal diagram where every variable in $\mathbf{V} \cup \mathbf{U}$ is a node, and an arrow is drawn from each member of $\mathbf{U}_i \cup \mathbf{Pa}_i$ to V_i. Following standard practice, when drawing a causal diagram, we omit the exogenous nodes and add a bidirected dashed arc between two endogenous nodes if they share an exogenous parent. We restrict our study to recursive systems, which means that the corresponding diagram will be acyclic. The marginal distribution induced over the endogenous variables $P(\mathbf{V})$ is called observational, and factorizes according to the causal diagram. *D-separation* criterion captures the conditional independence relations entailed by a causal diagram in $P(\mathbf{V})$. 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 any set $\mathbf{C} \subseteq \mathbf{V}$, the quantity $Q[\mathbf{C}]$ is defined to denote the post-intervention distribution of $\mathbf{C}$ under an intervention on $\mathbf{V} \setminus \mathbf{C}$, i.e. $P_{\mathbf{v} \setminus \mathbf{c}}(\mathbf{c})$.

Ancestral Graphs. We now introduce a graphical representation of equivalence classes of causal diagrams. A *mixed* graph can contain directed and bi-directed edges. A is an ancestor of B if there is a directed path from A to B. 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* is a path on which every non-endpoint node is a collider on the path (i.e., both edges incident to the node are into it) and is an ancestor of an endpoint of the path. A mixed graph is *ancestral* if it does not contain directed or almost directed cycles. It is *maximal* if there is no inducing path between any two non-adjacent nodes. A *Maximal Ancestral Graph* (MAG) is a graph that is both ancestral and maximal [Richardson and Spirtes, 2002]. In general, a MAG represents a set of causal diagrams with the same set of observed variables that entail the same conditional independence and ancestral relations among the observed variables. *M-separation* extends d-separation to MAGs such that d-separation in a causal diagram corresponds to m-separation in its unique MAG over the observed variables, and vice versa.

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 (i.e., edge marks that are shared by all members of $[\mathcal{M}]$). A circle indicates an edge mark that is not invariant. A PAG is learnable from the independence model over the observed variables, and the FCI algorithm is a standard method to learn such an object [Zhang, 2008b]. In short, a PAG represents a class of causal diagrams with the same observed variables that entail the same independence model.

Graphical Notions. Given a causal diagram, a MAG, or a 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. Y is called a *possible child* of X and X a *possible parent* of Y if they are adjacent and the edge is not into X. For a set of nodes $\mathbf{X}$, let $\texttt{An}(\mathbf{X})$ ($\texttt{De}(\mathbf{X})$) denote the union of $\mathbf{X}$ and the set of possible ancestors (descendants) of $\mathbf{X}$. 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 (*) as a wildcard to denote any possible mark of a PAG ($\circ, >, -$) or a MAG ($>, -$). Let p be any path in a PAG, and $\langle A, B, C \rangle$ be any consecutive triple along p. B is a collider on p if both edges are into B (i.e., $A {*\!\!\rightarrow} B \leftarrow\!\!* C$). B is a non-collider on p if one of the edges is out of B ($A \leftarrow B *\!\!-\!\!* C$ or $A *\!\!-\!\!* B \rightarrow C$), or both edges have a circle mark at B and there is no edge between A and B (i.e., $A \circ\!\!-\!\!\circ B \circ\!\!-\!\!\circ C$, where A and B are not adjacent). A path is *definite status* if every non-endpoint node along it is either a collider or a non-collider. If the edge marks on a path between X and Y are all circles, we call the path a *circle path*. We refer to the closure of nodes connected with circle paths as a *bucket*. Obviously, given a PAG, nodes are partitioned into a unique set of buckets.

A directed edge $X \rightarrow Y$ in a MAG or a PAG is *visible* if there exists no causal diagram in the corresponding equivalence class where there is an inducing path between X and Y that is into X. This implies that a visible edge is not confounded ($X \leftarrow\dashrightarrow Y$ does not 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*.

Manipulations in PAGs. Let $\mathcal{P}$ denote a PAG over $\mathbf{V}$ and $\mathbf{X} \subseteq \mathbf{V}$. The $\mathbf{X}$-lower-manipulation of $\mathcal{P}$ deletes all those edges that are visible in $\mathcal{P}$ and are out of variables in $\mathbf{X}$, replaces all those edges that are out of variables in $\mathbf{X}$ but are invisible in $\mathcal{P}$ with bi-directed edges, and otherwise keeps $\mathcal{P}$ as it is. The resulting graph is denoted as $\mathcal{P}_{\underline{\mathbf{X}}}$. The $\mathbf{X}$-upper-manipulation of $\mathcal{P}$ deletes all those edges in $\mathcal{P}$ that are into variables in $\mathbf{X}$, and otherwise keeps $\mathcal{P}$ as it is. The resulting graph is denoted as $\mathcal{P}_{\overline{\mathbf{X}}}$.

B Background Results

Lemma 1 (Lemma A.1 in [Zhang, 2008b]). *In a PAG, the following property holds:*

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 $A \leftrightarrow C$).

Lemma 2 (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, i.e., a path consisting of $\circ\!\!-\!\!\circ$ edges, between A and B, then:*

i. if there is an edge between A and B, the edge is not into A or B.

ii. for any other node C, $C {*\!\!\rightarrow} A$ if and only if $C {*\!\!\rightarrow} B$. Furthermore, $C \leftrightarrow A$ if and only if $C \leftrightarrow B$.

Lemma 3 (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 4 (Lemma B.1 in [Zhang, 2008b]). *If $p = \langle A, \ldots, B\rangle$ is a possibly directed path from A to B in a PAG $\mathcal{P}$, then some subsequence of p forms an uncovered possibly directed path from A to B in $\mathcal{P}$.*

Lemma 5 (cf. Lemma B.2 in [Zhang, 2008b], Lemma 7.2 in [Maathuis and Colombo, 2015]). *Let X and Y be distinct nodes in a PAG $\mathcal{P}$. If $p = \langle X = V_0, \ldots, V_k = Y\rangle$, $k \geq 2$, is an unshielded possibly directed path from X to Y in $\mathcal{P}$, and $V_{i-1} {*\!\!\rightarrow} V_i$ for some $i \in \{1, \ldots, k\}$, then $V_{j-1} \rightarrow V_j$ for all $j \in \{i+1, \ldots, k\}$.*

Lemma 6 (cf. Thm. 2 in [Zhang, 2008b], Lemma 7.6 in [Maathuis and Colombo, 2015]). *Let $\mathcal{P}$ be a PAG. Let $\mathcal{M}$ be the graph resulting from the following procedure applied to a $\mathcal{P}$:*

i. replace all partially directed edges $\circ\!\!\rightarrow$ in $\mathcal{P}$ with directed edges $\rightarrow$, and

ii. orient the subgraph of $\mathcal{P}$ consisting of all non-directed edges $\circ\!\!-\!\!\circ$ into a DAG with no unshielded colliders.

Then $\mathcal{M}$ is in the Markov equivalence class of $\mathcal{P}$. Moreover, if X is a node in $\mathcal{P}$, then one can always find an orientation of (ii) that does not create any new edges into X.

Lemma 7 (Lemma 48 in [Perković et al., 2018]). *Let X be a node in a PAG $\mathcal{P}$. Let $\mathcal{M}$ be a MAG in the equivalence class of $\mathcal{P}$ that satisfies Lemma 6. Then any edge that is either $X \circ\!\!-\!\!\circ Y$, $X \circ\!\!\rightarrow Y$ or invisible $X \rightarrow Y$ in $\mathcal{P}$ is invisible $X \rightarrow Y$ in $\mathcal{M}$.*

Lemma 8 (Proposition 1 in [Jaber et al., 2018a]). *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{D}$ be any causal diagram 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 $\mathcal{P}_{\mathbf{A}}$.*

Lemma 9 (Lemma 49 in [Perković et al., 2018]). *Let X and Y be distinct nodes in a PAG $\mathcal{P}$ such that there is a possibly directed path p^* from X to Y in $\mathcal{P}$ that does not start with a visible edge out of X. Then there is a MAG $\mathcal{M}$ in the Markov equivalence class of $\mathcal{P}$ (constructed according to*

Lemma 6) such that the path p in $\mathcal{M}$, consisting of the same sequence of nodes as p^* in $\mathcal{P}$, contains a subsequence p' that is a directed path from X to Y starting with an invisible edge in $\mathcal{M}$.

Lemma 10 (Lemma 10 in [Zhang, 2008a])**.** *Let $\mathcal{M}$ be any MAG over a set of variables $\mathbf{V}$ and edge set $\mathbf{E}$, and $A \to B$ be any directed invisible edge in $\mathcal{M}$. Let $\mathbf{E}'$ be a set of edges such that $X \to Y$ is in $\mathbf{E}'$ iff it is in $\mathbf{E}$ and $X \leftarrow\text{-}\text{-}\text{-}\rightarrow Y$ is in $\mathbf{E}'$ iff $X \leftrightarrow Y$ is in $\mathbf{E}$ for any $X, Y \in \mathbf{V}$. Finally, let $\mathcal{G}$ be a causal diagram over $\mathbf{V}$ with edge set $\mathbf{E}' \cup \{A \leftarrow\text{-}\text{-}\text{-}\rightarrow B\}$. Then, $\mathcal{M}$ is the MAG corresponding to $\mathcal{G}$.*

Lemma 11. *Let p be a path m-connecting X and Y given $\mathbf{Z}$ in MAG $\mathcal{M}$ such that no path over a subsequence of the vertices along p is m-connecting. Let p^* denote the corresponding path constituted by the same sequence of vertices in $\mathcal{P}$. Then, every non-endpoint vertex on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider.*

Proof. The exact same proof of [Zhang, 2006, Lemma 1; p. 208] applies here. □

Lemma 12. *Let p be a path m-connecting X and Y given $\mathbf{Z}$ in MAG $\mathcal{M}$ whose corresponding path in PAG $\mathcal{P}$, denoted by p^*, is not out of X with a definitely visible edge and such that no similar path over a subsequence of nodes along p is m-connecting. Then, every non-endpoint node on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider.*

Proof. The exact same proof of [Zhang, 2006, Lemma 1'; p. 217] applies here. □

Definition 7 (C-forest)**.** *Let $\mathcal{G}$ be a causal diagram, where $\mathbf{Y}$ is the root set. Then, $\mathcal{G}$ is a $\mathbf{Y}$-rooted C-forest if all nodes in $\mathcal{G}$ form a C-component, and all nodes have at most one child.*

Definition 8 (Hedge)**.** *Let $\mathbf{X}, \mathbf{Y}$ be sets of nodes in $\mathcal{G}$. Let $\mathbf{F}, \mathbf{F}'$ be $\mathbf{R}$-rooted C-forests such that $\mathbf{F} \cap \mathbf{X} \neq \emptyset$, $\mathbf{F}' \cap \mathbf{X} = \emptyset$, $\mathbf{F}' \subset \mathbf{F}$, and $\mathbf{R} \subseteq \mathit{PossAn}(\mathbf{Y})_{\mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}}$. Then, $\mathbf{F}, \mathbf{F}'$ form a hedge for $P_{\mathbf{x}}(\mathbf{y})$.*

Theorem 5 (Back-Door Hedge Criterion [Shpitser and Pearl, 2006])**.** *Let $\mathbf{W} \subseteq \mathbf{Z}$ be the unique maximal set such that $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) = P_{\mathbf{x},\mathbf{w}}(\mathbf{y}|\mathbf{z} \setminus \mathbf{w})$. Then $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ is identifiable from $P(\mathbf{V})$ if and only if there does not exist a hedge for $P_{\mathbf{x}'}(\mathbf{y}')$, for any $\mathbf{Y}' \subseteq (\mathbf{Y} \cup \mathbf{Z}) \setminus \mathbf{W}, \mathbf{X}' \subseteq \mathbf{X} \cup \mathbf{W}$.*

Definition 9 (DC-Component)**.** *In a MAG, a PAG, or any induced subgraph thereof, two nodes are in the same definite c-component (dc-component) if they are connected with a bi-directed path, i.e. a path of bi-directed edges.*

Definition 10 (DC-forest)**.** *Let $\mathcal{P}$ denote a subgraph of a PAG over $\mathbf{C}$. $\mathbf{Y}$ is a root set of $\mathcal{P}$ if $\mathbf{C} = \mathit{PossAn}(\mathbf{Y})_{\mathcal{P}}$ and it is maximal if no subset satisfies the property. Let $\mathbf{Y}$ be the maximal root set of $\mathcal{P}$. Then $\mathcal{P}$ is a $\mathbf{Y}$-rooted DC-forest if $\mathcal{P}$ is a dc-component and all nodes have at most one possible child via a directed ($\to$) or partially directed ($\circ\!\!\to$) edge.*

Definition 11 ($\mathcal{P}$-Hedge)**.** *Let $\mathbf{X}$, $\mathbf{Y}$ disjoint sets of nodes in PAG $\mathcal{P}$. Let $\mathbf{F}$, $\mathbf{F}'$ be $\mathbf{R}$-rooted DC-forests such that $\mathbf{F} \cap \mathbf{X} \neq \emptyset$, $\mathbf{F}' \cap \mathbf{X} = \emptyset$, $\mathbf{F}' \subset \mathbf{F}$, $\mathbf{R} \subseteq \mathit{An}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$. Then $\mathbf{F}$ and $\mathbf{F}'$ form a $\mathcal{P}$-hedge for $P_{\mathbf{x}}(\mathbf{Y})$ in $\mathcal{P}$.*

Theorem 6 (Non-Identifiability Criterion for PAGs [Jaber et al., 2019a])**.** *Given a PAG $\mathcal{P}$, $P_{\mathbf{x}}(\mathbf{y})$ is non-identifiable in $\mathcal{P}$ if and only if there exist:*

1. *proper possibly directed path from $\mathbf{X}$ to $\mathbf{Y}$ that starts with an invisible edge; or*
2. *dc-forests $\mathbf{F}$, $\mathbf{F}'$ forming a $\mathcal{P}$-hedge for $P_{\mathbf{x}}(\mathbf{y})$.*

Theorem 7 (Corol. 2 in [Jaber et al., 2019a])**. IDP** *(Alg. 1) is complete for marginal effect identification.*

Lemma 13 (Lemma 4 in [Jaber et al., 2019a])**.** *Given a PAG $\mathcal{P}$ over $\mathbf{V}$, $\mathbf{A} \subset \mathbf{C} \subseteq \mathbf{V}$, there doesn't exist a node $Z \in \mathbf{C}$ such that $Z \notin \mathcal{R}_{\mathbf{A}}^{\mathbf{C}} \wedge Y \in \mathcal{R}_{\mathbf{A}}^{\mathbf{C}} \wedge$ invisible $Z{*}\!\!\to Y$.*

Theorem 8 (Do-Calculus for MAGs; Thm. 1 in [Zhang, 2007])**.** *Let $\mathcal{M}$ be the MAG over $\mathbf{V}$, and $\mathbf{X}, \mathbf{Y}, \mathbf{W}, \mathbf{Z}$ be disjoint subsets of $\mathbf{V}$. The following rules are valid, in the sense that if the antecedent of the rule holds, then the consequent holds in every causal diagram represented by $\mathcal{M}$.*

1. $P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$.*
2. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$.*
3. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$.*

   where $\mathbf{X}(\mathbf{Z}) := \mathbf{X} \setminus \mathit{An}(\mathbf{Z})_{\mathcal{M}_{\mathbf{V}\setminus\mathbf{W}}}$.

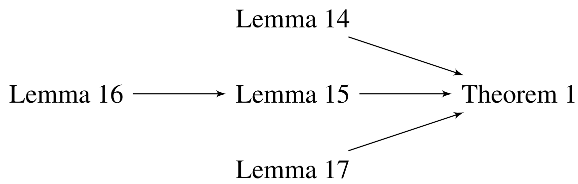

Figure 6: Roadmap to proving the do-calculus in Thm. 1.

Theorem 9 (Do-Calculus for Causal Diagrams; Thm. 3.4.1 in [Pearl, 2000])**.** *Let $\mathcal{G}$ be a causal diagram compatible with a structural causal model M, with endogenous variables $\mathbf{V}$. For any disjoint $\mathbf{X}, \mathbf{Y}, \mathbf{W}, \mathbf{Z} \subseteq \mathbf{V}$, the following rules are valid.*

1. $P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are d-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{W}}}$.*
2. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{x}, \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are d-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{W}}, \underline{\mathbf{X}}}$.*
3. $P(\mathbf{y}|do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y}|do(\mathbf{w}), \mathbf{z})$, *if $\mathbf{X}$ and $\mathbf{Y}$ are d-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$. where $\mathbf{X}(\mathbf{Z}) := \mathbf{X} \setminus An(\mathbf{Z})_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{W}}}$.*

C Proofs of Section 3.2

Figure 6 shows the results involved in proving the calculus in Thm. 1.

Proof of Theorem 1. **Rule 1.** By Lemma 14, if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$, then $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$. Hence, the consequent follows by rule 1 of Theorem 8.

Rule 2. By Lemma 15, if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}, \underline{\mathbf{X}}}$, then $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}, \underline{\mathbf{X}}}$. Hence, the consequent follows by rule 2 of Theorem 8.

Rule 3. By Lemma 17, if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$, $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$. Hence, the consequent follows by rule 3 of Theorem 8. □

Lemma 14. *Let $\mathcal{M}$ be a MAG over $\mathbf{V}$, and let $\mathcal{P}$ be the PAG that represents the Markov equivalence class of $\mathcal{M}$, i.e., $[\mathcal{M}]$. For $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, and disjoint sets $\mathbf{Z}, \mathbf{W} \subseteq \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y})$, if there is a proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$, then there is a proper path definitely m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$.*

Proof. We derive the following weaker claim first.

lemma I: Let p be a shortest proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$. Let p^* denote the corresponding path constituted by the same sequence of variables in $\mathcal{P}_{\overline{\mathbf{W}}}$. Then every non-endpoint node on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider.

Proof. First, if p is an active path in $\mathcal{M}_{\overline{\mathbf{W}}}$, then it is also active in $\mathcal{M}$. If p satisfies the if condition of Lemma 11, then p is a definite status path in $\mathcal{P}$ and consequently in $\mathcal{P}_{\overline{\mathbf{W}}}$. Otherwise, there is a subsequence of the nodes along p that constitutes an active path given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}$ but not in $\mathcal{M}_{\overline{\mathbf{W}}}$. In what follows, we argue that such a path does not exist which concludes the proof.

Assume for the sake of contradiction that there is a subsequence of the nodes along p that constitutes an active path between $X \in \mathbf{X}$ and $Y \in \mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}$, but not in $\mathcal{M}_{\overline{\mathbf{W}}}$, and let $p' = \langle X = O_0, \ldots, O_m = Y \rangle$ denote any such path. We have two cases to consider.

In the first case, p' is not defined in $\mathcal{M}_{\overline{\mathbf{W}}}$ which means that at least one of the edges along p' is into a node, denoted O_k, and $O_k \in \mathbf{W}$. Hence, O_k along p has two edges out of it, otherwise p is not defined in $\mathcal{M}_{\overline{\mathbf{W}}}$ which is a contradiction. However, this also leads to a contradiction since p is not active in $\mathcal{M}_{\overline{\mathbf{W}}}$ given $\mathbf{W} \cup \mathbf{Z}$ since $O_k \in \mathbf{W}$ and O_k is a non-collider.

In the second case, p' is defined in $\mathcal{M}_{\overline{\mathbf{W}}}$, but it is not m-connecting given $\mathbf{Z} \cup \mathbf{W}$. Let O_k denote any inactive node along p' in $\mathcal{M}_{\overline{\mathbf{W}}}$, and recall that O_k is active along p given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ since p is active in $\mathcal{M}_{\overline{\mathbf{W}}}$ by definition. Hence, O_k must have different status (collider or non-collider) along p and p'. If O_k is a non-collider along p' and it is active given $\mathbf{Z} \cup \mathbf{W}$ in $\mathcal{M}$ by definition, then it must be also active in $\mathcal{M}_{\overline{\mathbf{W}}}$ which is a contradiction. Hence O_k is a collider along p' and a non-collider along p. Let O_p denote the node that is adjacent to O_k along p' and closest to X, and let p'' denote the subpath of p between O_k and O_p. Since O_k is a non-collider along p, one of the edges incident on O_k must be out of it. Assume without loss of generality that p'' is out of O_k. We have two cases to consider here both of which lead to contradictions. If p'' is a directed path from O_k to O_p then we have a cycle or almost directed cycle due to $O_k \dashrightarrow O_p$ and one of $O_p \rightarrow O_k$ or $O_p \leftrightarrow O_k$. This violates the ancestral property in MAGs and hence it is not possible. Otherwise, let O_l denote the first collider along p'' starting from O_k. Recall O_l is active along p given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ by definition, then O_l has a descendant in $\mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$. Therefore, in $\mathcal{M}_{\overline{\mathbf{W}}}$, O_k has a descendant in $\mathbf{Z}$ through O_l due to $O_k \dashrightarrow O_l$. This contradicts the choice of O_k as being inactive along p' given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ which This concludes the proof. □

Given any path p m-connecting X and Y given $\mathbf{Z}$ in $\mathcal{M}$, for every collider Q on p, there is a directed path (possibly of length 0) from Q to a member of $\mathbf{Z}$. Define the distance-from-$\mathbf{Z}$ of Q to be the length of a shortest directed path (possibly of length 0) from Q to $\mathbf{Z}$, and define the distance-from-$\mathbf{Z}$ of p to be the sum of the distances from $\mathbf{Z}$ of the colliders on p.

lemma II: Let p be a shortest path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ such that no equally short m-connecting path has a shorter distance-from-$\mathbf{W} \cup \mathbf{Z}$ than p does. Let p^* denote the corresponding path constituted by the same sequence of nodes in $\mathcal{P}_{\overline{\mathbf{W}}}$. Then, p^* is a definite m-connecting path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$.

Proof. Since p is a shortest m-connecting path in $\mathcal{M}_{\overline{\mathbf{W}}}$, then by Lemma I, every non-endpoint vertex on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider. We note here that all the nodes along p including the endpoints are are not in $\mathbf{W}$ since (1) $\mathbf{X} \cap \mathbf{W} = \emptyset$ and (2) the path would not be connecting given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$. The rest of the proof follows exactly as [Zhang, 2006, Lemma 2; p. 213] while considering p in $\mathcal{M}_{\overline{\mathbf{W}}}$ instead of $\mathcal{M}$. □

Finally, the main result follows from Lemma II. □

Lemma 15. *Let $\mathcal{M}$ be a MAG over $\mathbf{V}$, and let $\mathcal{P}$ be the PAG that represents the Markov equivalence class of $\mathcal{M}$, i.e., $[\mathcal{M}]$. For $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, and disjoint sets $\mathbf{Z}, \mathbf{W} \subseteq \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y})$, if there is a proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$, then there is a proper path definitely m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$.*

Proof. Let $p = \langle O_0, \dots, O_m \rangle$ denote a proper active path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$. Then, p must be into $O_0 \in \mathbf{X}$ by definition of $\mathbf{X}$-lower manipulation of MAGs. The same path is defined and active in $\mathcal{M}_{\overline{\mathbf{W}}}$ with the only possible difference of having $O_0 \leftrightarrow O_1$ in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$ and $O_0 \rightarrow O_1$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ if $O_0 \rightarrow O_1$ is invisible in $\mathcal{M}$. It follows that the corresponding path of p in $\mathcal{M}_{\overline{\mathbf{W}}}$ is not out of $O_0 \in \mathbf{X}$ with a visible edge. By Lemma 16, there is a proper path definitely m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$ that is not out of $\mathbf{X}$ with a visible edge. Let $p^* = \langle A_0, \dots, A_p \rangle$ denote one such path in $\mathcal{P}_{\overline{\mathbf{W}}}$. Note that, except for A_0, none of the nodes along p^* and the corresponding directed paths from colliders to $\mathbf{Z}$ belong to $\mathbf{X}$ else the path is not proper. Since the edge between A_0 and A_1 is not a visible edge out of A_0 in $\mathcal{P}$ (and $\mathcal{P}_{\overline{\mathbf{W}}}$), then the corresponding path of p^* in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$ is defined and definitely m-connecting with the only possible difference of having $A_0 \leftrightarrow A_1$ in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$ and $A_0 \rightarrow A_1$ in $\mathcal{P}_{\overline{\mathbf{W}}}$ if $A_0 \rightarrow A_1$ is defined and invisible in $\mathcal{P}$. This concludes the proof. □

Lemma 16. *Let $\mathcal{M}$ be a MAG over $\mathbf{V}$, and let $\mathcal{P}$ be the PAG that represents the Markov equivalence class of $\mathcal{M}$, i.e., $[\mathcal{M}]$. For $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, and disjoint sets $\mathbf{Z}, \mathbf{W} \subseteq \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y})$, if there is a proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ such that the path is not out of $\mathbf{X}$ with a visible edge, then there is a proper path definitely m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$ that is not out of $\mathbf{X}$ with a visible edge.*

Proof. In order to prove the lemma, we start with the following weaker result.

lemma I: Let p be a shortest proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ whose corresponding path in $\mathcal{P}$, denoted by p^*, is not out of $\mathbf{X}$ with a visible edge. Then every non-endpoint node on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider.

Proof. If p satisfies the conditions of Lemma 12 in $\mathcal{M}$, then p^* is a definite status path in $\mathcal{P}$ by consequent of the same lemma. Otherwise, there exists a shorter m-connecting path in $\mathcal{M}$ over a subsequence of the nodes along p. Following the same argument as that of Lemma I in Lemma 14, we show that such a shorter path does not exist in $\mathcal{M}$. □

Given any path p m-connecting X and Y given $\mathbf{Z}$ in $\mathcal{M}$, for every collider Q on p, there is a directed path (possibly of length 0) from Q to a member of $\mathbf{Z}$. Define the distance-from-$\mathbf{Z}$ of Q to be the length of a shortest directed path (possibly of length 0) from Q to $\mathbf{Z}$, and define the distance-from-$\mathbf{Z}$ of p to be the sum of the distances from $\mathbf{Z}$ of the colliders on p.

lemma II: Let p be a shortest proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$ whose corresponding path in $\mathcal{P}$, denoted by p^*, is not out of $\mathbf{X}$ with a visible edge and such that no equally short proper m-connecting path has a shorter distance-from-$\mathbf{W} \cup \mathbf{Z}$ than p does in $\mathcal{M}_{\overline{\mathbf{W}}}$. Then, p^* is a definite m-connecting path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$.

Proof. Since p is a shortest proper m-connecting path in $\mathcal{M}_{\overline{\mathbf{W}}}$, then by Lemma I, every non-endpoint vertex on p^*, if any, is of a definite status in $\mathcal{P}$ and consequently in $\mathcal{P}_{\overline{\mathbf{W}}}$. We note here that all the nodes along p including the endpoints are not in $\mathbf{W}$ for two reasons: (1) sets $\mathbf{X}, \mathbf{Y}, \mathbf{W}$ are disjoint, and (2) if any non-endpoint node along p is in $\mathbf{W}$, p would not be active given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$. Also, for any collider Q along p, the shortest directed path from Q to $\mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$, denoted d, remains a shortest directed path from Q to $\mathbf{Z}$ in $\mathcal{M}$ over any subsequence of nodes along d since none of the nodes along it are in $\mathbf{W}$. It follows by the above two observations that every adjacency in $\mathcal{M}_{\overline{\mathbf{W}}}$ between nodes along p and/or the respective shortest directed paths from colliders along p to $\mathbf{Z}$ also exists in $\mathcal{M}$ and vice versa. Therefore, the rest of the proof follows exactly as in [Zhang, 2006, Lemma 2'; p. 218] while considering p in $\mathcal{M}_{\overline{\mathbf{W}}}$ instead of $\mathcal{M}$ and deriving a contradiction by establishing a shorter active path or an equally-long active path with a shorter distance to $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}}$. □

Finally, the main result follows from Lemma II. □

Lemma 17. *Let $\mathcal{M}$ be a MAG over $\mathbf{V}$, and let $\mathcal{P}$ be the PAG that represents the Markov equivalence class of $\mathcal{M}$, i.e., $[\mathcal{M}]$. For $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, and disjoint sets $\mathbf{Z}, \mathbf{W} \subseteq \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y})$, if there is a path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$, then there is a path definitely m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$.*

Proof. First, for any $X \in \mathbf{X}$, if $X \notin \mathbf{X}(\mathbf{Z})$ given $\mathcal{M}$ according to the definition in rule 3 of Thm. 8, then $X \in \texttt{An}(\mathbf{Z})_{\mathcal{M}_{\mathbf{V} \setminus \mathbf{W}}}$. It follows, by Lemma 8, that $X \in \texttt{PossAn}(\mathbf{Z})_{\mathcal{P}_{\mathbf{V} \setminus \mathbf{W}}}$, and consequently $X \notin \mathbf{X}(\mathbf{Z})$ given $\mathcal{P}$ according to the definition in rule 3 of Thm. 1. Hence, every path in $\mathcal{M}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$ is also defined in $\mathcal{P}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$ since the adjacencies in latter are a subset of the ones in the former. Next, we derive the following weaker claim.

lemma I: Let p be a shortest proper path m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$. Let p^* denote the corresponding path constituted by the same sequence of variables in $\mathcal{P}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$. Then every non-endpoint vertex on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider.

Proof. First, if p is an active path in $\mathcal{M}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$, then it is also active in $\mathcal{M}$. If p satisfies the if condition of Lemma 11, then p^* is a definite status path in $\mathcal{P}$ and consequently in $\mathcal{P}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$. Otherwise, there is a subsequence of the nodes along p that constitutes an active path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}$ but not in $\mathcal{M}_{\overline{\mathbf{W}}, \overline{\mathbf{X}(\mathbf{Z})}}$. In what follows, we argue that such a path does not exist which concludes the proof.

Assume for the sake of contradiction that there is a subsequence of the nodes along p that constitutes an active path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}$ but not in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, and let $p' = \langle X = O_0, \ldots, O_m = Y\rangle$ denote any such path, where $X \in \mathbf{X}$, $Y \in \mathbf{Y}$. We have two cases to consider.

In the first case, p' is not defined in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ which means that at least one of the edges along p' is into a node, denoted O_k, and $O_k \in \mathbf{W} \cup \mathbf{X}(\mathbf{Z})$. If $O_k \neq X$, then O_k is a non-collider along p with both edges out of O_k, otherwise the edges into O_k are removed in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ and p is not defined in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ which is a contradiction. However, if $O_k \in \mathbf{W}$ is a non-collider along p, then p is blocked given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. This contradicts the choice of p as a proper active path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Alternatively, we have $O_k = O_0 = X$ and $X \in \mathbf{X}(\mathbf{Z})$. Let O_p denote the node adjacent to O_k along p', and note that O_p is along path p as well. Then, the first edge along p starting from O_k, i.e. X, is out of O_k, else the edge would be removed in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ and path p is not defined because $O_k \in \mathbf{X}(\mathbf{Z})$. Given the presence of $O_k \leftarrow\!\!* O_p$ in $\mathcal{M}$, the subpath of p between O_k and O_p cannot be a directed path out of O_k since this creates a directed or almost directed cycle and violates the ancestral property of MAGs. Hence, there is a collider along the subpath of p between O_k and O_p and this collider is active in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Let C denote the first collider along p starting from O_k. Since the first edge along p is out of X, then C is a descendant of X in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ and $X \notin \mathbf{X}(\mathbf{Z})$ which is a contradiction.

In the second case, p' is defined in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, but it is not active given $\mathbf{W} \cup \mathbf{Z}$. Let O_k denote any inactive node along p' in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, and recall that O_k is active along p given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ since p is active in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ by definition. Hence, O_k must have different status (collider or non-collider) along p and p'. If O_k is a non-collider along p' and it is active in $\mathcal{M}$ given $\mathbf{W} \cup \mathbf{Z}$ by definition, then it must be active in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ as well, which contradicts the choice of O_k. Hence, O_k is a collider along p' and a non-collider along p. Let O_p denote the node that is adjacent to O_k along p' and closest to X, and let p'' denote the subpath of p between O_k and O_p. Since O_k is a non-collider along p, one of the edges incident on O_k must be out of it. Assume without loss of generality that p'' is out of O_k. We have two cases to consider here both of which lead to contradictions. If p'' is a directed path from O_k to O_p then we have a cycle or almost directed cycle due to $O_k \dashrightarrow O_p$ and one of $O_p \rightarrow O_k$ or $O_p \leftrightarrow O_k$. This violates the ancestral property in MAGs and hence it is not possible. Otherwise, let O_l denote the first collider along p'' starting from O_k. Recall O_l is active along p given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ by definition, then O_l has a descendant in $\mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Therefore, in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, O_k has a descendant in $\mathbf{Z}$ due to $O_k \dashrightarrow O_l$. This contradicts the choice of O_k as being inactive along p' given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ which concludes the proof. $\square$

Given any proper path p m-connecting $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{Z}$ in $\mathcal{M}$, for every collider Q on p, there is a directed path (possibly of length 0) from Q to a member of $\mathbf{Z}$. Define the distance-from-$\mathbf{Z}$ of Q to be the length of a shortest directed path (possibly of length 0) from Q to $\mathbf{Z}$, and define the distance-from-$\mathbf{Z}$ of p to be the sum of the distances from $\mathbf{Z}$ of the colliders on p.

lemma II: Let p be a shortest proper active path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ such that no equally short active path has a shorter distance-from-$\mathbf{W} \cup \mathbf{Z}$ than p does. Let p^* denote the corresponding path constituted by the same sequence of nodes in $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Then, p^* is a definite m-connecting path between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$.

Proof. This proof follows almost exactly as [Zhang, 2006, Lemma 2; p. 213] while accounting for a few special cases. We note first that all the nodes along p, except for the one in $\mathbf{X}$, are not in $\mathbf{W}$, or else the path would not be connecting in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Since p is a shortest active path in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, then by Lemma I, every non-endpoint vertex on p^*, if any, is of a definite status, i.e., either a definite collider or a definite non-collider. Since $\mathcal{P}$ is the PAG of $\mathcal{M}$, every definite non-collider on p^* corresponds to a non-collider on p, and hence is not in $\mathbf{W} \cup \mathbf{Z}$, for otherwise p would not be m-connecting given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$.

Similarly, for any definite collider Q on p^*, Q is also a collider on p. Hence, there is a directed path (possibly of length 0) from Q to a member of $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ and consequently in $\mathcal{M}$. Let d be a shortest such path from Q to, say, $C \in \mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Note that this directed path is also the shortest in $\mathcal{M}$ over any subsequence of d or else d would not exist in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. Let d^* denote the corresponding path in $\mathcal{P}$. Because $\mathcal{P}$ is a PAG of $\mathcal{M}$, d^* is a potentially directed path from Q to C in $\mathcal{P}$ and $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. We now show that no circle mark ($\circ$) appears on d^*, i.e., that d^* is (fully) directed. Suppose for contradiction that there is a circle on d^*. Then, the mark at Q on d^* must be a circle, for otherwise, an arrowhead would meet a circle on d^*, and by the property in Lemma 1, a proper subpath of d^* would constitute a potentially directed path from Q to C, which in turn implies that there is a shorter directed path from Q to C in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ than d is, a contradiction with our choice of d.

Let $Q \circ\!\!-\!\!* S$ be the first edge along d^*. Suppose S is not on p^* for the moment. Since Q is a definite collider on p^*, we have $Q_l *\!\!\rightarrow Q \leftarrow\!\!* Q_r$ in $\mathcal{P}$, Q_l, Q_r being the two nodes adjacent to Q on p^*. By Lemma 1, there is an edge between Q_l and S that is into S, and there is an edge between Q_r and S that is into S, i.e., $Q_l *\!\!\rightarrow S \leftarrow\!\!* Q_r$ in $\mathcal{P}$.

Now we show that there exists a node W (distinct from Q) on $p(X,Q)$ such that (i) there is an edge between W and S in $\mathcal{M}$ and $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ that is into S; and (ii) in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, the collider/non-collider status of W on p is the same as the collider/non-collider status of W on $p(X,W) \oplus \langle W,S\rangle$. To show this, it suffices to demonstrate that if no node between X and Q on p satisfies the two conditions, then X must satisfy them. Suppose no vertex between X and Q on p satisfies the two conditions. First, consider $Q_l = X$. X satisfies (i) and (ii) trivially in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ if the edge is out of X, i.e., we have $X \rightarrow S$, or $X \notin \mathbf{X}(\mathbf{Z})$. Suppose for the sake of contradiction that we have $X \leftrightarrow S$ and $X \in \mathbf{X}(\mathbf{Z})$. Then, we have $X \rightarrow Q$ or else the edge would be cut as part of the operation $\overline{\mathbf{X}(\mathbf{Z})}$ and p is not defined in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. However, in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, X would have a descendant in $\mathbf{Z}$ due to $X \rightarrow Q$ and Q being an active collider along d; a contradiction. Next, suppose $Q_l \neq X$. We argue by induction that every vertex between X and Q is a collider on p and is a parent of S in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. The rest of the proof follows almost exactly like [Zhang, 2006, Lemma 2; p. 213] in addition to the earlier argument for $X \notin \mathbf{X}(\mathbf{Z})$ to show that X satisfies conditions (i) and (ii) in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ if no node along $p(X,Q)$ satisfies them.

By symmetry, it follows that there exists a vertex V (distinct from Q) on $p(Q,Y)$ such that (i) there is an edge between V and S in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ that is into S; and (ii) the collider/non-collider status of V on p is the same as the collider/non-colllider status of V on $\langle S,V\rangle \oplus p(V,Y)$. Then, the path $p' = p(X,W) \oplus \langle W,S,V\rangle p(V,Y)$ (it could be that $X = W$ and/or $V = Y$) is obviously m-connecting given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. It is easy to check that either p' is shorter than p is, or p' is as long as p is (when $W = Q_l$ and $V = Q_r$) but has a shorter distance-from-$\mathbf{W} \cup \mathbf{Z}$ than p. Either case is a contradiction with our assumption about p.

Finally, if S is on p^*, the argument follows same as in [Zhang, 2006, Lemma 2; p. 213]. $\square$

Finally, the main result follows from Lemma II. $\square$

Proof of Theorem 2. **Rule 1.** If the condition fails, then there exists a proper definite m-connecting path p between some $X \in \mathbf{X}$ and some $Y \in \mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$. We use Lemma 6 to construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$. Since p is a definite m-connecting path in $\mathcal{P}_{\overline{\mathbf{W}}}$, then the corresponding path in $\mathcal{P}_{\overline{\mathbf{W}}}$, denoted p', is m-connecting as well. Next, we construct a causal diagram $\mathcal{G}$ in the equivalence class of $\mathcal{M}$ by keeping the directed edges and replacing bidirected edges with bidirected dashed arcs (special case of Lemma 10). It follows easily that the corresponding path to p' in $\mathcal{G}_{\overline{\mathbf{W}}}$ is active given $\mathbf{W} \cup \mathbf{Z}$. Hence, rule 1 of Theorem 9 is not applicable in $\mathcal{G}$.

Rule 2. If the condition fails, then there exists a proper definite m-connecting path between some $X \in \mathbf{X}$ and some $Y \in \mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$. Let p denote the path in $\mathcal{P}$ corresponding to the same sequence of nodes along the m-connecting path in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$. We use Lemma 6 to construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ such that there are no additional edges into X in $\mathcal{M}$. Let p' denote the path corresponding to p in $\mathcal{M}$. If the first edge along p incident on X is not into X in $\mathcal{P}$, then the corresponding edge in $\mathcal{M}$ is out of X by the construction in Lemma 6, denoted $X \rightarrow J$, and invisible by Lemma 7. Next, we construct a causal diagram $\mathcal{G}$ in the equivalence class of $\mathcal{M}$ using Lemma 10 such that we add $X \leftarrow\!\!\text{-}\,\text{-}\!\!\rightarrow J$ in $\mathcal{G}$ if $X \rightarrow J$ is along p' else we follow the special case as in **Rule 1**. It follows that there exists an active path between X and Y given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$. Hence, rule 2 of Theorem 9 is not applicable in $\mathcal{G}$.

Rule 3. If the condition fails, then there exists a proper definite m-connecting path p between some $X \in \mathbf{X}$ and some $Y \in \mathbf{Y}$ given $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$. We use Lemma 6 to construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ such that there are no additional edges into X in $\mathcal{M}$. If X is a possible ancestor of $\mathbf{Z}$ in $\mathcal{P}_{\mathbf{V} \setminus \mathbf{W}}$, then by Lemma 4, there exists in $\mathcal{P}$ an uncovered possibly directed path from X to a node in $\mathbf{Z}$, denoted p^d, such that no node along the path is in $\mathbf{W}$. It follows by the construction in Lemma 6 that the path corresponding to p^d in $\mathcal{M}$ is directed from X to the same node in $\mathbf{Z}$. Hence, the path corresponding to p in $\mathcal{M}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$, denoted p', is defined and is also m-connecting given $\mathbf{W} \cup \mathbf{Z}$. Finally, we construct a causal diagram $\mathcal{G}$ in the equivalence class of $\mathcal{M}$ by keeping the directed edges and replacing bidirected edges with bidirected dashed arcs (special case of Lemma 10). It follows easily that the corresponding path to p' in $\mathcal{G}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$ is active given $\mathbf{W} \cup \mathbf{Z}$. Hence, rule 3 of Theorem 9 is not applicable in $\mathcal{G}$. $\square$

D Proof of Theorem 3

Figure 7 shows the results involved in proving Theorem 3.

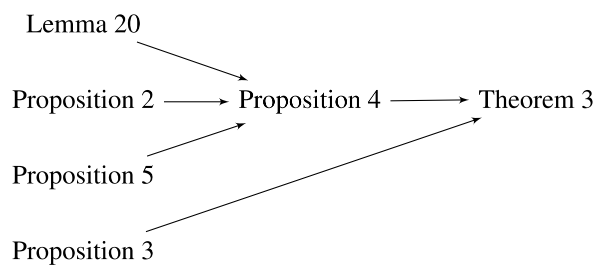

Figure 7: Roadmap to proving Theorem 3.

Proof of Theorem 3. Prop. 3 shows that **IDP**, as formulated in [Jaber et al., 2019a], fails whenever Alg. 1 fails. Following the completeness of the earlier [Jaber et al., 2019a, Corol. 2], we have Alg. 1 complete. This result along with Prop. 4 establishes the second part of the claim. $\square$

Proposition 3. *Given a PAG $\mathcal{P}$ and a target causal effect $P_{\mathbf{x}}(\mathbf{y})$, if Alg. 1 fails to identify the effect, then **IDP** in [Jaber et al., 2019a] fails to identify it as well.*

Proof. Suppose Alg. 1 fails to identify the target effect, and let $\textsc{Identify}(\mathbf{C}, \mathbf{T}, Q)$ be the recursive call that throws a fail. If any bucket/circle component is split between $\mathbf{C}$ and $\mathbf{T} \setminus \mathbf{C}$, then the original **IDP** algorithm would fail as well. Otherwise, each bucket in $\mathbf{T}$ is either a subset of $\mathbf{C}$ or $\mathbf{T} \setminus \mathbf{C}$. Since Alg. 1 fails the condition at line 6, then every bucket in $\mathbf{T} \setminus \mathbf{C}$ is in the same pc-component with a descendant. Also, since Alg. 1 fails the condition at line 9, then $\mathcal{R}_{\mathbf{B}} = \mathbf{C}$ for all buckets $\mathbf{B} \subseteq \mathbf{C}$. Let C^* be any node in any sink bucket in $\mathcal{P}_{\mathbf{C}}$, i.e., a bucket with only arrowheads incident on. Node C^* is in the same pc-component with at least one node from every bucket in $\mathcal{P}_{\mathbf{C}}$ due to [Zhang, 2006, Lem. 3.3.2], and the paths consistent with Def. 5 are into C^*.

To prove our claim, we start with the following observation. Let $\mathbf{A}$ denote any bucket in $\mathbf{T} \setminus \mathbf{C}$ and B denote a descendant node such that B is in the pc-component of $\mathbf{A}$. Suppose B is in the pc-component of $\mathbf{A}$ due to collider path that is not into B, and let D denote the second last node along the path before B. According to the PAG properties in [Zhang, 2008b, Lems. B.1 & B.2], there is a potentially directed from $\mathbf{A}$ to B that is into B ($E{*}{\rightarrow}\, B$). If we have $B{\circ}{\rightarrow}\, D$, then we have $E \rightarrow D$ or $E{\circ}{\rightarrow}\, D$ by [Zhang, 2008b, Lem. A.1]. Also, if we have $B \rightarrow D$, then we have $E \rightarrow D$ or $E{\circ}{\rightarrow}\, D$ by [Zhang, 2008b, FCI:R2] and since $B \rightarrow D$ is invisible. It follows that $\mathbf{A}$ has a descendant (D) that is in the same pc-component with it and the corresponding path consistent with Def. 5 is into that descendant. By the above observation, for every bucket $\mathbf{A}$ in $\mathcal{P}_{\mathbf{T}}$ we have a sequence of nodes $\langle T_1 \in \mathbf{A}, T_2, \ldots, T_p = C^* \rangle$ such that every consecutive pair $\langle T_i, T_{i+1} \rangle$ is connected with a path consistent with Def. 5 that is into T_{i+1}.

Finally, we show that every bucket in $\mathbf{T} \setminus \mathbf{C}$ is in the same pc-component with a possible child, and hence fails the condition of Prop. 2. This concludes the proof as the original **IDP** algorithm would fail as well. Let $\mathbf{A}$ denote any bucket in $\mathbf{T} \setminus \mathbf{C}$ and consider a sequence of nodes discussed earlier $\langle T_1 \in \mathbf{A}, T_2, \ldots, T_p = C^* \rangle$. We show by induction on the number of nodes in the latter sequence that $\mathbf{A}$, i.e., some node in $\mathbf{A}$, is in the same pc-component with C^* and the corresponding path is into C^*. The claim trivially holds for a sequence of two. Assume the claim holds for a sequence of i nodes, we prove it for $i+1$ nodes. Consider the first three nodes $\langle T_1, T_2, T_3 \rangle$. If T_2 and T_3 share a bi-directed path, then T_1 and T_3 share a collider path that is into T_3 and we obtain a sequence of i nodes excluding T_2 which concludes the proof. Otherwise, the path between T_2 and T_3 consistent with Def. 5 is not into T_2. Let $T_2 {*\!\!\rightarrow} B$ denote the first edge along such path. If any node along the path between T_1 and T_2 is adjacent to B with a bi-directed edge, then we are done as we can exclude T_2 and obtain a shorter sequence of nodes. Otherwise, by applying Lemma 18 recursively, every node along the path between T_1 and T_2 has an invisible edge into B, including T_1. This concludes the proof as we obtain a shorter sequence by excluding T_2 since T_1 and B share an invisible edge into B. It follows from the above that every bucket in $\mathbf{T} \setminus \mathbf{C}$ is in the same pc-component with C^* and the corresponding path is into C^*. Also, note that every bucket $\mathbf{A}$ in $\mathbf{T} \setminus \mathbf{C}$ has a possible child in another bucket $\mathbf{B}$ by the failing of the algorithm (line 6), and consequently every node in $\mathbf{B}$ is a possible child of $\mathbf{A}$ by [Zhang, 2008b, Lem. A.1]. Therefore, every bucket in $\mathbf{T} \setminus \mathbf{C}$ is in the same pc-component with a possible child which concludes the proof. □

Lemma 18. *In $\mathcal{P}_{\mathbf{A}}$, where $\mathcal{P}$ is a PAG over $\mathbf{V}$ and $\mathbf{A} \subseteq \mathbf{V}$, the following property holds:*

For any three vertices A, B, C, if $A {*\!\!\rightarrow} B ?\!\!\rightarrow C$ and both edges are invisible, then we have $A {*\!\!\rightarrow} C$ and the edge is invisible.[5]

Proof. We prove this property for $\mathbf{A} = \mathbf{V}$. The same trivially holds for $\mathbf{A} \subset \mathbf{V}$ since $\mathcal{P}_{\mathbf{A}}$ is an induced subgraph of $\mathcal{P}$. If $B {\circ\!\!\rightarrow} C$, then we have $A {*\!\!\rightarrow} C$ by [Zhang, 2008b, Lemma A.1]. Also, if $B \rightarrow C$ and the edge is invisible, then A and C are adjacent and the edge is into C by [Zhang, 2008b, FCI:R2].

We still need to show that the edge between A and C is invisible if it was oriented out of A. Suppose for the sake of contradiction that $A \rightarrow C$ is visible. We can construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ using [Perković et al., 2018, Lem. 43] with possible circles incident on A and B oriented out of both. Note that the construction is possible since A and B are in different circle components in $\mathcal{P}$. It follows by [Perković et al., 2018, Lem. 48] that, in $\mathcal{M}$, we have $A {*\!\!\rightarrow} B \rightarrow C$ and both edges are invisible while $A \rightarrow C$ is visible. This violates the MAG property in Lemma 19, and hence it is not possible. Therefore, a directed edge $A \rightarrow C$ would be invisible in $\mathcal{P}$. This concludes the proof. □

Lemma 19 ([Jaber et al., 2018b], Lemma 6)**.** *In $\mathcal{M}_{\mathbf{A}}$, where $\mathcal{M}$ is a MAG over $\mathbf{V}$ and $\mathbf{A} \subseteq \mathbf{V}$, the following property holds:*

For any three vertices A, B, C, if $A {*\!\!\rightarrow} B \rightarrow C$ and both edges are invisible, then we have $A {*\!\!\rightarrow} C$ and the edge is invisible.

Proposition 4. *Alg. 1 can be mapped to a sequence of calculus rules from Thm. 1 along with standard probability manipulations.*

Proof. This follows from the derivation of Lem. 20, Prop. 2, and Prop. 5. □

Lemma 20. *Let $\mathbf{D} = \mathit{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V} \setminus \mathbf{X}}}$, then we have $P_{\mathbf{x}}(\mathbf{y}) = P_{\mathbf{v} \setminus \mathbf{d}}(\mathbf{y}) = \sum_{\mathbf{d} \setminus \mathbf{y}} Q[\mathbf{D}]$.*

Proof. We use rule 3 of Thm. 1 to prove that $P_{\mathbf{x}}(\mathbf{y}) = P_{\mathbf{v} \setminus \mathbf{d}}(\mathbf{y})$. Let $\mathbf{W} = \mathbf{V} \setminus (\mathbf{D} \cup \mathbf{X})$ and consider, for the sake of contradiction, p to be any proper definite m-connecting path between $W \in \mathbf{W}$ and $Y \in \mathbf{Y}$ given $\mathbf{X}$ in $\mathcal{P}_{\overline{\mathbf{X}}, \overline{\mathbf{W}}}$. If p is possibly directed from W to Y then no node in $\mathbf{X}$ is along p $W \in \mathbf{D}$ which is a contradiction. Else, let C be the collider along p that is closest to W. Since C is active then there is a directed path from C to $\mathbf{X}$ in $\mathcal{P}_{\overline{\mathbf{X}}, \overline{\mathbf{W}}}$. This is not possible since we cut the incoming to $\mathbf{X}$ and we have a contradiction. Therefore, rule 3 is applicable and we get $P_{\mathbf{x}}(\mathbf{y}) = P_{\mathbf{v} \setminus \mathbf{d}}(\mathbf{y})$. Finally, the second identity follows by standard probability manipulation. □

[5] $*$ is a wildcard for any of the possible marks $(-, \circ, >)$ and ? is a wildcard for tail or circle.

Proof of Proposition 2. Let $\mathbf{V}^- = \mathbf{T} \setminus \texttt{PossDe}(\mathbf{X})$, $\mathbf{V}^+ = \mathbf{T} \setminus (\mathbf{V}^- \cup \mathbf{X})$, and $\mathbf{W} = \mathbf{V} \setminus \mathbf{T}$. Let $\mathbf{B}_1 < \cdots < \mathbf{B}_m$ be a topological order of the circle components in $\mathcal{P}_\mathbf{T}$ relative to $\mathbf{X}$ such that all the circle components after $\mathbf{X}$ are possible descendants of $\mathbf{X}$, and let $\mathbf{B}^i$ denote $\bigcup_{\{k|k\leq i\}} \mathbf{B}_k$.

$$
\begin{aligned}
Q[\mathbf{T} \setminus \mathbf{X}] &= P_{\mathbf{w}\cup\mathbf{x}}(\mathbf{T} \setminus \mathbf{X}) && (2)\\
&= \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{V}^-\}} P_{\mathbf{w}\cup\mathbf{x}}(\mathbf{B}_i|\mathbf{B}^{(\mathbf{i-1})}) \times \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{V}^+\}} P_{\mathbf{w}\cup\mathbf{x}}(\mathbf{B}_i|\mathbf{B}^{(\mathbf{i-1})} \setminus \mathbf{X}) && (3)\\
&= \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{V}^-\}} P_{\mathbf{w}}(\mathbf{B}_i|\mathbf{B}^{(\mathbf{i-1})}) \times \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{V}^+\}} P_{\mathbf{w}\cup\mathbf{x}}(\mathbf{B}_i|\mathbf{B}^{(\mathbf{i-1})} \setminus \mathbf{X}) && (4)\\
&= \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{V}^-\}} P_{\mathbf{w}}(\mathbf{B}_i|\mathbf{B}^{(\mathbf{i-1})}) \times \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{V}^+\}} P_{\mathbf{w}}(\mathbf{B}_i|\mathbf{B}^{(\mathbf{i-1})}) && (5)\\
&= \frac{P_{\mathbf{v}\setminus\mathbf{t}}}{P_{\mathbf{v}\setminus\mathbf{t}}(\mathbf{X}|\mathbf{T} \setminus \texttt{PossDe}(\mathbf{X}))} && (6)
\end{aligned}
$$

Line 3 follows by standard probability manipulations.

Line 4 follows by rule 3 of Thm. 1 since $(\mathbf{B}_i \perp\!\!\!\perp \mathbf{X}|\mathbf{W} \cup \mathbf{B}^{(\mathbf{i-1})})_{\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}}}}$ where $\mathbf{B}^i \subseteq \mathbf{V}^-$. This is because any definite status path between $\mathbf{X}$ and $\mathbf{B}_i$ in $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}}}$ will not be into $\mathbf{X}$ and hence it includes a collider with no descendants in $\mathbf{W} \cup \mathbf{B}^{(\mathbf{i-1})}$.

Line 5 follows by applying rule 2 of Thm. 1, i.e., $(\mathbf{B}_i \perp\!\!\!\perp \mathbf{X}|\mathbf{W} \cup \mathbf{B}^{(\mathbf{i-1})} \setminus \mathbf{X})_{\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}}$, where $\mathbf{B}^i \subseteq \mathbf{V}^+$. First, the condition $C^\mathbf{X} \cap \texttt{PossDe}(\mathbf{X}) \subseteq \mathbf{X}$ implies that all the edges incident on $\mathbf{X}$ in $\mathcal{P}_\mathbf{T}$ are either into $\mathbf{X}$ or out of it and visible. Next, for the sake of contradiction, let p denote any proper definitely m-connecting path between $\mathbf{X}$ and $\mathbf{B}_i \subseteq \mathbf{V}^+$ in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$. It follows by the earlier observation that p is not a direct adjacency and it is into $\mathbf{X}$ since all visible edges incident on $\mathbf{X}$ are removed by definition of the lower manipulation in PAGs. If p contains any non-collider, then either it is in $\mathbf{B}^{(\mathbf{i-1})}$ which blocks the path or it lies after $\mathbf{B}_i$ in the partial order which leads to an inactive collider along p. Hence, every non-endpoint node along p must be a collider for it to be connecting. Finally, let D denote the last non-endpoint node along p and closest to $\mathbf{B}_i$. Either the edge between $\mathbf{B}_i$ and D is bidirected or D is a possible child of $\mathbf{B}_i$ and consequently a possible descendant of $\mathbf{X}$. In the earlier case, $\mathbf{B}_i$ violates the condition $C^\mathbf{X} \cap \texttt{PossDe}(\mathbf{X}) \subseteq \mathbf{X}$ while D violates it in the latter case. Therefore, path p does not exists which concludes the proof. ☐

Proposition 5. *Given a PAG $\mathcal{P}$ over $\mathbf{V}$ and set $\mathbf{C} \subseteq \mathbf{V}$, $Q[\mathbf{C}]$ can be decomposed as follows where $\mathbf{A} \subseteq \mathbf{C}$ and $\mathcal{R}_{(.)} = \mathcal{R}^\mathbf{C}_{(.)}$.*

$$
Q[\mathbf{C}] = \frac{Q[\mathcal{R}_\mathbf{A}].Q[\mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_\mathbf{A}}]}{Q[\mathcal{R}_\mathbf{A} \cap \mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_\mathbf{A}}]}
$$

Proof. Let $\mathbf{B}_1 < \cdots < \mathbf{B}_m$ be a topological order of the circle components in $\mathcal{P}_\mathbf{C}$ following the procedure in Lemma 21, and let $\mathbf{B}^i$ denote $\bigcup_{\{m|m\leq i\}} \mathbf{B}_m$. Also, let $\mathcal{R}^\mathbf{C}_{\overline{\mathcal{R}_\mathbf{A}}}$, $\mathcal{R}_{\overline{\mathcal{R}_\mathbf{A}}}$ for short, denote the set $\mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_{\mathbf{A}}}$. In what follows, we argue for the following derivation.

$$P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{C}) = \prod_{\{i|\mathbf{B}_i\subseteq\mathbf{C}\}} P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{B}_i|\mathbf{B^{(i-1)}}) \tag{7}$$

$$= \frac{\prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\}} P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{B}_i|\mathbf{B^{(i-1)}}) \times \prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\}} P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{B}_i|\mathbf{B^{(i-1)}})}{\prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\}} P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{B}_i|\mathbf{B^{(i-1)}})} \tag{8}$$

$$= \frac{\prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\}} P_{\mathbf{v}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)}(\mathbf{B}_i|\mathbf{B^{(i-1)}}) \times \prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\}} P_{\mathbf{v}\setminus(\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\cup\mathbf{B}^i)}(\mathbf{B}_i|\mathbf{B^{(i-1)}})}{\prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\}} P_{\mathbf{v}\setminus(\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\cup\mathbf{B}^i)}(\mathbf{B}_i|\mathbf{B^{(i-1)}})} \tag{9}$$

$$= \frac{\prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\}} P_{\mathbf{v}\setminus\mathcal{R}_{\mathbf{A}}}(\mathbf{B}_i|\mathbf{B^{(i-1)}}\cap\mathcal{R}_{\mathbf{A}}) \times \prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\}} P_{\mathbf{v}\setminus\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}}(\mathbf{B}_i|\mathbf{B^{(i-1)}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}})}{\prod_{\{i|\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\}} P_{\mathbf{v}\setminus(\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}})}(\mathbf{B}_i|\mathbf{B^{(i-1)}}\cap(\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}))} \tag{10}$$

$$= \frac{Q[\mathcal{R}_{\mathbf{A}}]\times Q[\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}]}{Q[\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}]} \tag{11}$$

Lines 7 & 8 follow by standard probability manipulations assuming the quantities in the denominator of Eq. 8 are not zero.

Line 9 holds by applying rule 3 of Thm. 1 to each conditional term as follows. Consider the term $P_{\mathbf{v}\setminus\mathbf{c}}(\mathbf{B}_i|\mathbf{B^{(i-1)}})$ for $\mathbf{B}_i \subseteq \mathcal{R}_{\mathbf{A}}$ and let $\mathbf{W} = \mathbf{C}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)$. We claim that $(\mathbf{W} \perp\!\!\!\perp \mathbf{B}_i|\mathbf{B}^{(i-1)}\cup(\mathbf{V}\setminus\mathbf{C}))_{\mathcal{P}_{\overline{\mathbf{V}\setminus\mathbf{C}},\overline{\mathbf{W}(\mathbf{B}^{(i-1)})}}}$, where $\mathbf{W}(\mathbf{B}^{(i-1)}) = \mathbf{W}\setminus\texttt{PossAn}(\mathbf{B}^{(i-1)})_{\mathcal{P}_{\mathbf{C}}} = \mathbf{W}$ by definition of $\mathbf{W}$. For the sake of contradiction, let p denote any proper definitely m-connecting path between $\mathbf{W}$ and $\mathbf{B}_i$ in $\mathcal{P}_{\overline{\mathbf{V}\setminus\mathbf{C}},\overline{\mathbf{W}}}$. Since $\mathbf{C}$ is an arbitrary subset of $\mathbf{V}$, then $\mathbf{W}$ can include strict subsets of circle component in $\mathcal{P}_{\overline{\mathbf{V}\setminus\mathbf{C}},\overline{\mathbf{W}}}$. If p starts with a circle edge $\circ\!\!-\!\!\circ$ incident on $\mathbf{W}$, then the other end of the circle edge is in $\mathbf{V}\setminus\mathbf{C}$ by construction of $\mathbf{W}$ as a subset of the buckets in $\mathcal{P}_{\mathbf{C}}$. It follows that p is blocked since we condition on $\mathbf{V}\setminus\mathbf{C}$. So p starts with a directed or partially directed edge out of $\mathbf{W}$. Since $\mathbf{W}$ comes after $\mathbf{B}^i$ in the partial order, p must include an inactive collider and it is blocked, a contradiction. Hence, the separation holds. The same argument applies to the conditional terms where $\mathbf{B}_i\subseteq\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}$ and $\mathbf{W} = \mathbf{C}\setminus(\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\cup\mathbf{B}^i)$, and where $\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}$ and $\mathbf{W} = \mathbf{C}\setminus(\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}\cup\mathbf{B}^i)$.

Line 10 holds by applying rule 2 of Thm. 1 to each conditional term as follows. Consider the term $P_{\mathbf{v}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)}(\mathbf{B}_i|\mathbf{B^{(i-1)}})$ for $\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}$ and let $\mathbf{W} = \mathbf{B}^i\setminus\mathcal{R}_{\mathbf{A}}$. We claim that $(\mathbf{W} \perp\!\!\!\perp \mathbf{B}_i|\mathbf{V}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)\cup\mathbf{B^{(i-1)}}\setminus\mathbf{W}))_{\mathcal{P}_{\overline{\mathbf{V}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)},\underline{\mathbf{W}}}}$. Suppose for the sake of contradiction there exists a definitely m-connecting path between $\mathbf{W}$ and $\mathbf{B}_i$ in $\mathcal{P}_{\overline{\mathbf{V}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)},\underline{\mathbf{W}}}$, and let p denote any such path. Since $\mathbf{W}$ comes before $\mathbf{B}_i$ in the partial order over $\mathcal{P}_{\mathbf{C}}$, then any direct adjacency between the sets in $\mathcal{P}_{\overline{\mathbf{V}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)},\underline{\mathbf{W}}}$ must be into $\mathbf{B}_i$. The edge cannot be visible since we cut the outgoing edges of $\mathbf{W}$. Also, the edge cannot be invisible by Lemma 13. Hence, p is not a direct edge and contains at least one non-endpoint node along it. Path p does not have a non-collider along it because (1) we condition on it if the non-collider comes before $\mathbf{B}_i$ and (2) p will have an inactive collider if the non-collider comes after $\mathbf{B}_i$. Hence, every non-endpoint node along p is a collider. Also, p is into $\mathbf{B}_i$ else it contains an inactive collider than comes after $\mathbf{B}^i$ in the partial order. Moreover, the first edge along p incident on $\mathbf{W}$ cannot be visible since we cut them in $\mathcal{P}_{\overline{\mathbf{V}\setminus(\mathcal{R}_{\mathbf{A}}\cup\mathbf{B}^i)},\underline{\mathbf{W}}}$. The resulting collider path into $\mathbf{B}_i$ violates the property in Lemma 13. Hence, p does not exist and the separation holds. Finally, we note that the property of regions in Lemma 13 applies to $\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}$ as well since a violation for the intersection implies the same in one of the two regions. Therefore, the same argument applies to the conditional terms where $\mathbf{B}_i\subseteq\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}$ and $\mathbf{W} = \mathbf{B}^i\setminus\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}$, and where $\mathbf{B}_i\subseteq\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}}$ and $\mathbf{W} = \mathbf{B}^i\setminus(\mathcal{R}_{\mathbf{A}}\cap\mathcal{R}_{\overline{\mathcal{R}_{\mathbf{A}}}})$. This concludes the proof. □

Lemma 21. *Given a PAG $\mathcal{P}$ over $\mathbf{V}$, set $\mathbf{C}\subseteq\mathbf{V}$, and a partial topological order $\mathcal{O} := \mathbf{B}_1 < \cdots < \mathbf{B}_m$ over the circle components in $\mathcal{P}$, then the following procedure constructs a sound partial order $\mathcal{O}'$ over the circle components in $\mathcal{P}_{\mathbf{C}}$.*

1. *Initialize $\mathcal{O}' = \mathcal{O}$.*

2. *For every bucket* $\mathbf{B}_i \in \mathcal{O}'$ *such that* $\mathbf{B}_i \cap \mathbf{C} = \emptyset$*, set* $\mathcal{O}' = \mathcal{O} \setminus \mathbf{B}_i$.
3. *For every bucket* $\mathbf{B}_i \in \mathcal{O}'$ *such that* $\mathbf{B}_i \setminus \mathbf{C} \neq \emptyset$*, replace* $\mathbf{B}_i$ *with the buckets in* $\mathcal{P}_{\mathbf{B}_i \setminus \mathbf{C}}$.

Proof. The partial order over the circle components after step 2 is sound by the soundness of the input partial topological order over $\mathbf{V}$ in $\mathcal{P}$. Given any bucket $\mathbf{B}_i \in \mathcal{O}'$ in step 3 such that $\mathbf{B}_i \setminus \mathbf{C} \neq \emptyset$, let $\mathbf{D}_1, \ldots, \mathbf{D}_p$ denote the resultant circle components in $\mathcal{P}_{\mathbf{B}_i \setminus \mathbf{C}}$.

To prove the soundness of step 3, it suffices to show that there are no possibly directed paths between any pair $\mathbf{D}_j, \mathbf{D}_k$ outside $\mathbf{B}_i$ in $\mathcal{P}$, and consequently in $\mathcal{P}_{\mathbf{C}}$. Suppose for the sake of contradiction there exists a possibly directed path from $\mathbf{D}_j$ to $\mathbf{D}_k$ in $\mathcal{P}$ that goes outside $\mathbf{B}_i$. Then such a path includes at least one directed or partially directed edge ($\rightarrow$ or $\circ\!\!\rightarrow$). Also, there is a possibly directed path from $\mathbf{D}_k$ to $\mathbf{D}_j$ composed of circle edges since both $\mathbf{D}_j, \mathbf{D}_k$ are subsets of the circle component $\mathbf{B}_i$. This violates the property in Lemma 3 and hence it is not possible. This concludes the proof. $\square$

E Proof of Theorem 4

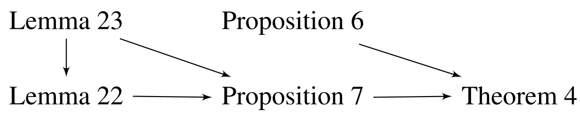

Figure 8: Roadmap to proving Theorem 4.

Proof of Theorem 4. The result follows by Propositions 6 and 7. $\square$

Proposition 6. *If Alg. 2 fails at line 7, we can construct a causal diagram in the Markov equivalence class of the PAG* $\mathcal{P}$ *such that the input query* $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ *is not identifiable by Thm. 5.*

Proof. If the algorithm fails at line 7, then we have $(\mathbf{X}' \not\perp\!\!\!\perp \mathbf{Y} | (\mathbf{X} \setminus \mathbf{X}') \cup \mathbf{Z})_{\mathcal{P}_{\overline{\mathbf{X} \setminus \mathbf{X}'}, \underline{\mathbf{X}'}}}$. Let p denote any proper definitely m-connecting path between $X \in \mathbf{X}'$ and $Y \in \mathbf{Y}$. Path p' in $\mathcal{P}$ corresponding to the same sequence of nodes is also active given $(\mathbf{X} \setminus \mathbf{X}') \cup \mathbf{Z}$. We have two cases to consider.

Case (1): p' is not into X and the edge incident on X is not visible.

Case (1a): If there is no collider along p', then p' is a possibly directed path from X to Y. By Lemma 9, we obtain a MAG in the equivalence class of $\mathcal{P}$ where there is a directed path from X to Y over a subsequence of the nodes in p', and where the path starts with an invisible edge out of X. Then, by Lemma 10, we construct a causal diagram $\mathcal{G}$ in the Markov equivalence class where the invisible edge is confounded. Let A denote the node adjacent to X along the directed path from X to Y. Then, $\mathbf{F} = \{X, A\}, \mathbf{F}' = \{A\}$ form a hedge for $P_x(y)$ (by Def. 8) and $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ is not identifiable in $\mathcal{G}$ by Thm. 5.

Case (1b): There is at least one collider along p' (and consequently p). Let C denote the first collider along p' starting from X. Since the subpath of p' from X to C, denoted p'_{XC}, is a possibly directed path that starts with an invisible edge, then by Lemma 9, we can construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ with a directed path from X to C over a subsequence of the nodes along p'_{XC} and such that the directed path starts with an invisible edge in $\mathcal{M}$. We denote the directed path from X to C in $\mathcal{M}$ as p^d_{XC}. Next, we use Lemma 10 to construct a causal diagram $\mathcal{G}$ in the Markov equivalence class where the first edge along p^d_{XC} is confounded. Let A denote the node adjacent to X along the directed path from X to C in $\mathcal{M}$ and let Z^* denote the descendant of C such that C is active along p' in $\mathcal{P}$. Then, $\mathbf{F} = \{X, A\}, \mathbf{F}' = \{A\}$ form a hedge for $P_x(z^*)$ (by Def. 8). It is left to show that Z^* is not in the unique maximal set $\mathbf{W}$ according to Thm. 5. Since p' is a definitely m-connecting path in $\mathcal{P}$, then the path corresponding to the same sequence of nodes in $\mathcal{M}$ is also active, and consequently in $\mathcal{G}$. Finally, in $\mathcal{G}$, let p^* denote the concatenated path composed of the directed path from C to Z^* and the subpath of p' between C and Y. It is easy to see that p^* is active given $\mathbf{X} \cup (\mathbf{Z} \setminus \{Z^*\})$ in $\mathcal{G}_{\overline{\mathbf{X}}, \underline{Z^*}}$ since p' is active given $\mathbf{X} \setminus \{X'\} \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{X} \setminus \{X'\}}, \underline{\mathbf{Z}}}$. Therefore, Z^* does not belong to the unique maximal set $\mathbf{W}$ in Thm. 5. This concludes **Case (1b)**.

Case (2): p' is into X.

Case (2a) If there exists a node $X^* \in \mathbf{B}_i \cap \mathbf{X}$ (where $\mathbf{B}_i$ is the bucket of X) such that $X^* \in \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus(\mathbf{X}\cup\mathbf{Z})}}$, then we proceed as in **Case (1a)**.

Case(2b): Let $X^* \in \mathbf{B}_i \cap \mathbf{X}$ such that $X^* \circ\!\!-\!\!\circ V_j$ and $V_j \in \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$, where $V_j \in \mathbf{V}$. Such an intervention node exists by the condition in Line 2. Hence, consider a possibly directed path $p = \langle X^*, V_j, \ldots, Z^* \rangle$, where $Z^* \in \mathbf{Z}$, let A denote the node closest to Z^* along p such that A is in the same bucket as X^* in $\mathcal{P}$, and let D denote the node adjacent to A along p that is closest to Z^*. Note that D is not in the same bucket as A and the edge between them is into D.

Next, we construct a MAG $\mathcal{M}$ following the procedure in Lemma 6 such that we have no new edges into X^* and D. The step of not having additional arrowheads incident on both X^* and D in $\mathcal{M}$ is possible because the two nodes are in different buckets in $\mathcal{P}$. In $\mathcal{M}$, we have a directed path from X^* to A that does not start with a visible edge. This follows by Lemma 9 as the claim is based on applying the procedure in Lemma 6 and not orienting further arrowheads into X^*. Also, in $\mathcal{M}$, we have a directed path from D to Z^* due to the following. There is a possibly directed path from D to Z^* in $\mathcal{P}$, then by Lemma 4, there is an uncovered possibly directed path from D to Z^* in $\mathcal{P}$. Also, since we orient the PAG according to Lemma 6 such that there are no additional arrowheads into D, then the uncovered possibly directed path from D to Z^* is oriented as a directed path in the MAG $\mathcal{M}$. It follows that there is a directed path $\langle X^*, \ldots, A, D, \ldots, Z^* \rangle$ in $\mathcal{M}$ that starts with an invisible edge. Finally, we use Lemma 10 to construct a causal diagram $\mathcal{G}$ such that the first edge along the directed path from X^* to Z^*, i.e. $X^* \to C$ is confounded. This leads to the hedge $\mathbf{F} = \{X^*, C\}$,$\mathbf{F}' = \{C\}$ for $P^*_x(z^*)$. It is left to show that Z^* does not belong to the maximal set $\mathbf{W}$ such that $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) = P_{\mathbf{x},\mathbf{w}}(\mathbf{y}|\mathbf{z} \setminus \mathbf{w})$ according to Thm. 5.

Back to p', the definitely m-connecting path between X and Y given $(\mathbf{X} \setminus \mathbf{X}') \cup \mathbf{Z}$ in $\mathcal{P}$. Let B denote the node closest to X along p'. Since p' is into X, then B is adjacent to every node in the bucket of X including A and the edge is into A by Lem. 2. Also, by Lem. 2, the edge between B and A is bi-directed if and only if the edge between B and X is bi-directed in $\mathcal{P}$. It follows that B and A are adjacent in $\mathcal{G}$ and the edge is bi-directed if and only if the same is bi-directed in $\mathcal{P}$. Let p^* denote the concatenated path in $\mathcal{G}$ composed of p'_{BY} (the subpath of p'), $B{*\!\!\to} A$, and the directed path from A to Z^*. First, p'_{BY} is active given $(\mathbf{X} \setminus \mathbf{X}') \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{X}\setminus\mathbf{X}'},\underline{\mathbf{X}'}}$ since the same is active in $\mathcal{P}_{\overline{\mathbf{X}\setminus\mathbf{X}'},\underline{\mathbf{X}'}}$. Then, p'_{BY} is active given $\mathbf{X} \cup \mathbf{Z}$ in $\mathcal{G}_{\overline{\mathbf{X}}}$ since no node in $\mathbf{X}'$ is along p'_{BY}. If p'_{BY} is not active in $\mathcal{G}_{\overline{\mathbf{X}},\underline{Z^*}}$ given $\mathbf{X} \cup (\mathbf{Z} \setminus \{Z^*\})$, then $(\{Z^*\} \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup (\mathbf{Z} \setminus \{Z^*\}))_{\mathcal{G}_{\overline{\mathbf{X}},\underline{Z^*}}}$ and $\{Z^*\}$ does not belong to the maximal set $\mathbf{W}$ in Thm. 5 which concludes the proof. Otherwise, p'_{BY} is active in $\mathcal{G}_{\overline{\mathbf{X}},\underline{Z^*}}$ given $\mathbf{X} \cup (\mathbf{Z} \setminus \{Z^*\})$. Second, B is a collider along p^* if and only if it is a collider along p' due to the PAG property in Lem. 2 and the construction of the MAG from the PAG by Lem. 6 and the causal diagram from the MAG by Lem. 10. Similar to the argument in the first point, B is active along p^* given $\mathbf{X} \cup (\mathbf{Z} \setminus \{Z^*\})$ in $\mathcal{G}_{\overline{\mathbf{X}},\underline{Z^*}}$ else Z^* does not belong to the maximal set $\mathbf{W}$ in Thm. 5 which concludes the proof. Third, the subpath p^*_{BZ} has no colliders along it and all the non-endpoint nodes are not in $\mathbf{X} \cup \mathbf{Z}$. Therefore, p^* is active in $\mathcal{G}_{\overline{\mathbf{X}},\underline{Z^*}}$ given $\mathbf{X} \cup (\mathbf{Z} \setminus \{Z^*\})$ and Z^* does not belong to the maximal set $\mathbf{W}$ in Thm. 5. This concludes the proof. □

Proposition 7. *If the call to* **IDP**$(\cdot)$ *in Alg. 2 fails, we can construct a causal diagram in the Markov equivalence class of the PAG $\mathcal{P}$ such that the input query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ is not identifiable by Thm. 5.*

Proof. Since **IDP**$(\cdot)$ is complete for marginal effect identification (Thm. 7), one of the two graphical conditions in Thm. 6 is true whenever the call in Line 10 of Alg. 2 fails.

Case 1: There exists a proper possibly directed path from $A \in \mathbf{X}$ to $B \in \mathbf{Y} \cup \mathbf{Z}$ in $\mathcal{P}$ that starts with an invisible edge. We have three sub-cases to consider.

Case 1a: $B \in \mathbf{Y}$. Hence, we have a possibly directed path from A to $B \in \mathbf{Y}$, denoted p, that starts with an invisible edge. By Lemma 9, there is a MAG $\mathcal{M}$ in the Markov equivalence class of $\mathcal{P}$ such that the path in $\mathcal{M}$, consisting of the same sequence of nodes as p in $\mathcal{P}$, contains a subsequence p^* that is a directed path from A to B starting with an invisible edge in $\mathcal{M}$. By Lemma 10, we construct a causal diagram in the Markov equivalence class of $\mathcal{M}$, and consequently $\mathcal{P}$, where p^* is directed from A to B and the first edge is confounded. Let C denote the first node along p^* after A. It follows easily that $\mathbf{F} = \{A, C\}, \mathbf{F}' = \{C\}$ form a hedge for $P_{\mathbf{a}}(b)$ in $\mathcal{G}$ while $A \in \mathbf{X}'$ and $B \in \mathbf{Y}'$ according to Thm. 5. Hence, the input query is not identifiable in $\mathcal{G}$, and consequently in $\mathcal{P}$.

Alternatively, $B \in \mathbf{Z}$ and we have a possibly directed path from A to $B \in \mathbf{Z}$, denoted p, that starts with an invisible edge. Since B remains in $\mathbf{Z}$ after Lines 8- 9 in Alg. 2, there exists a definitely m-connecting path between B and some $Y \in \mathbf{Y}$ given $\mathbf{X} \cup \mathbf{Z} \setminus B$ in $\mathcal{P}_{\overline{\mathbf{X}},\underline{B}}$. Let $p^\dagger$ denote the path in $\mathcal{P}$ corresponding to the same sequence of nodes as the definitely m-connecting path between B and Y in $\mathcal{P}_{\overline{\mathbf{X}},\underline{B}}$. We have two cases to consider.

Case 1b: $p^\dagger$ is into B in $\mathcal{P}$. By Lemma 9, we construct a MAG $\mathcal{M}$ in the Markov equivalence class of $\mathcal{P}$ such that the path in $\mathcal{M}$, consisting of the same sequence of nodes as p in $\mathcal{P}$, contains a subsequence p^* that is a directed path from A to B starting with an invisible edge in $\mathcal{M}$. By Lemma 10, we construct a causal diagram $\mathcal{G}$ in the Markov equivalence class such that p^* remains in $\mathcal{G}$ and the first edge out of B is confounded. Since $p^\dagger$ is a definitely m-connecting path between B and Y given $\mathbf{X} \cup \mathbf{Z} \setminus \{B\}$ in $\mathcal{P}$ that is into B, then the same corresponding path is also active and into B in $\mathbf{M}$, and consequently in $\mathcal{G}$. It follows that B is not in the maximal set $\mathbf{W}$ in Thm. 5. Let C denote the node adjacent to A along p^* in $\mathcal{G}$. Then, the pair $\mathbf{F} = \{A, C\}, \mathbf{F}' = \{C\}$ form a hedge for $P_a(b)$ while $A \in \mathbf{X}'$ and $B \in \mathbf{Y}'$ according to Thm. 5. Hence, the input query is not identifiable in $\mathcal{G}$, and consequently in $\mathcal{P}$.

Case 1c: $p^\dagger$ is not into B in $\mathcal{P}$. Then $p^\dagger$ starts from B with $\circ\!\!-\!\!\circ$, $\circ\!\!\rightarrow$, or invisible $\rightarrow$. By Lemma 23, there exists an uncovered possibly directed path from A to B in $\mathcal{P}$, denoted p^*, that starts with an invisible edge. Let C denote the node along the uncovered possibly directed path that is adjacent to A. An important note here is that A and B are not in the same circle component in $\mathcal{P}$ following Lemma 22. So, we construct a MAG $\mathcal{M}$ following the procedure in Lemma 6 and such that no new edges are into A and B. This is possible since A and B are in different circle components and each circle component can be oriented independently. First, p', the path corresponding to P^* in $\mathcal{M}$, is a directed path out of A since the edge between A and C is out of A in $\mathcal{M}$ and every non-endpoint node along p^* is a definite non-collider in $\mathcal{P}$. Second, $p^\ddagger$, the path corresponding to $p^\dagger$ in $\mathcal{M}$, is out of B. Third, by Lemma 7, we have invisible $A \rightarrow C$ in $\mathcal{M}$ and the first edge out of B along p' is also invisible. Let D denote the node along p' in $\mathcal{M}$ that is adjacent to B and let E denote the node along $p^\ddagger$ that is adjacent to B. In $\mathcal{M}$, we have invisible $B \rightarrow E$ and $D \rightarrow B \rightarrow E$, then D and E are adjacent in $\mathcal{M}$ by [Zhang, 2008a, Def. 8]. Also, we have $D \rightarrow E$ in $\mathcal{M}$ else we have $D \rightarrow B \rightarrow E$ and $E{*}\!\!\rightarrow D$ and we violate the ancestral property in $\mathcal{M}$. Finally, we construct a causal diagram $\mathcal{G}$ in the equivalence class of $\mathcal{M}$ following Lemma 10 where $A \rightarrow C$ is confounded. Since $p^\dagger$ is a definitely m-connecting path between B and Y given $\mathbf{X} \cup \mathbf{Z} \setminus \{B\}$ in $\mathcal{P}$, then the same corresponding path is also active in $\mathbf{M}$, and consequently in $\mathcal{G}$. Recall $B \rightarrow E$ and $D \rightarrow E$ are in $\mathcal{M}$, and consequently $\mathcal{G}$. Also, $D \notin \mathbf{Z}$ since D is along a proper causal path from $A \in \mathbf{X}$ to $B \in \mathbf{Z}$. Then, the concatenated path composed of $B \leftarrow D \rightarrow E$ and $p^\ddagger(E, Y)$, the subpath of $p^\ddagger$ between E and Y, is active given $\mathbf{X} \cup \mathbf{Z} \setminus \{B\}$ in $\mathcal{G}_{\overline{\mathbf{X}},\underline{B}}$ and B does not belong to the unique maximal set $\mathbf{W}$ in Thm. 5. Therefore, the pair $\mathbf{F} = \{A, C\}, \mathbf{F}' = \{C\}$ form a hedge for $P_a(b)$ while $A \in \mathbf{X}'$ and $B \in \mathbf{Y}'$ according to Thm. 5. Thus, the input query is not identifiable in $\mathcal{G}$, and consequently in $\mathcal{P}$.

Case 2: There exist dc-forests $\mathbf{F}, \mathbf{F}'$ in $\mathcal{P}$ forming a $\mathcal{P}$-hedge for $P_{\mathbf{x}}(\mathbf{y} \cup \mathbf{z})$ such that $\mathbf{X}$ and $\mathbf{Z}$ are the sets obtained right before executing Line 10 in Alg. 2. Recall by Def. 11 that dc-forests $\mathbf{F}, \mathbf{F}'$ have a subset $\mathbf{R} \subseteq \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V} \setminus \mathbf{X}}}$ as a root set in $\mathcal{P}$. Let $\mathbf{R}' \subseteq \mathbf{R}$ be such that no pair of nodes in $\mathbf{R}'$ belong to the same bucket in $\mathcal{P}$. Also, let $\mathcal{F}, \mathcal{F}'$ denote the corresponding induced subgraphs of $\mathcal{P}$ over the nodes in $\mathbf{F}, \mathbf{F}'$, respectively, excluding $\mathbf{R} \setminus \mathbf{R}'$. Finally, let I be any node in $\mathbf{F}$ (or $\mathbf{F}'$), J be any node in $\mathbf{R} \setminus \mathbf{R}'$, and $J^* \in \mathbf{R}'$ such that J and J^* are in the same bucket in $\mathcal{P}$. By the property in Lemma 2, if $I \rightarrow J$ or $I \circ\!\!\rightarrow J$ is in $\mathcal{P}$ then $I \rightarrow J^*$ or $I \circ\!\!\rightarrow J^*$ is in $\mathcal{P}$ as well. Moreover, we have $I \leftrightarrow J$ in $\mathcal{P}$ iff $I \leftrightarrow J^*$ is in $\mathcal{P}$ as well. It follows that $\mathcal{F}, \mathcal{F}'$ preserve the properties of a $\mathcal{P}$-Hedge for $P_{\mathbf{x}}(\mathbf{y} \cup \mathbf{z})$ except for the one where each node has at most a single possible child via $\rightarrow$ or $\circ\!\!\rightarrow$. This can be easily obtained by dropping unnecessary edges from the subgraphs $\mathcal{F}, \mathcal{F}'$. Next, we have a couple of cases to consider.

Case 2a: $\mathbf{R}' \cap \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Z})}} \neq \emptyset$. Using Lemma 6, we construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ with no additional arrowheads into all the nodes in $\mathcal{R}'$. This is possible since each node in $\mathcal{R}'$ is in a different circle component in $\mathcal{P}$. By Lemma 4 and the above construction of the MAG, there exists a directed path from each node in $\mathcal{R}'$ to some node in $\mathbf{Z} \cup \mathbf{Y}$ such that the directed paths do not go through $\mathbf{X}$. Also, at least one node in $\mathbf{R}'$ has a directed path to $\mathbf{Y}$ that is not intercepted by $\mathbf{Z} \cup \mathbf{X}$. Next, a causal diagram $\mathcal{G}$ can be trivially constructed from $\mathcal{M}$ by keeping every directed edge and replacing bi-directed edges with dashed arcs (a special case of Lemma 10). Let $\mathbf{A} = \mathcal{F} \cap \mathbf{X}$ and let $\mathbf{B} \subseteq \mathbf{Y} \cup \mathbf{Z}$ be a minimal set such that $\mathbf{R}' \subseteq \texttt{An}(\mathbf{B})_{\mathcal{G}_{\overline{\mathbf{X}},\underline{\mathbf{Y}},\underline{\mathbf{Z}}}}$ and such that $\mathbf{B} \cap \mathbf{Y} \neq \emptyset$. The subgraphs corresponding to $\mathcal{F}, \mathcal{F}'$ in $\mathcal{G}$ form a hedge for $P_{\mathbf{a}}(\mathbf{b})$. Next, let Z' denote any node in $\mathbf{B} \cap \mathbf{Z}$ and consider the path $Z' \dashleftarrow V_i \dashleftarrow\!\!\dashrightarrow V_j \dashrightarrow Y$ where $Y \in \mathbf{B} \cap \mathbf{Y}$, $V_i, V_j \in \mathbf{R}'$, and every node along the bi-directed path between V_i and V_j, inclusive, is in $\mathcal{F}'$. The directed paths $Z' \dashleftarrow V_i$ and $V_j \dashrightarrow Y$ are not intercepted by $\mathbf{X} \cup \mathbf{Z}$ by definition of $\mathbf{B}$, and all the nodes along the bi-directed path between V_i and V_j are not in $\mathbf{X}$ by definition of $\mathcal{F}'$ and have descendants in $\mathbf{Z}$. Hence, this path is active in $\mathcal{G}_{\overline{\mathbf{X}},\underline{Z'}}$ given $\mathbf{X} \cup \mathbf{Z} \setminus \{Z'\}$. Following Thm. 5, $\mathcal{F}, \mathcal{F}'$ forms a hedge for $P_{\mathbf{a}}(\mathbf{b})$ and no node in $\mathbf{B}$ is in the maximal set $\mathbf{W}$. Hence, the input query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ is not identifiable.

Case 2b: $\mathbf{R}' \cap \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus(\mathbf{X}\cup\mathbf{Z})}} = \emptyset$, and there exists a node $J \in \mathbf{Z}$ such that $(\{J\} \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{J\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{J}}}$ and J is in the same bucket as some node in $\mathbf{R}'$, denoted Q. We can replace Q with J in $\mathbf{R}'$ while still ensuring that the pair $\mathcal{F}, \mathcal{F}'$ forms a $\mathcal{P}$-Hedge for $P_{\mathbf{x}}(\mathbf{y} \cup \mathbf{z})$ (as argued earlier via Lemma 2). Similar to **case 2a**, we use Lemma 6 to construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$. Next, we construct a causal diagram $\mathcal{G}$ as follows. Recall that $J \in \mathbf{R}'$ and $(\{J\} \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{J\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{J}}}$. Let p denote a corresponding active path between J and $\mathbf{Y}$ given $\mathbf{X} \cup \mathbf{Z} \setminus \{J\}$ in $\mathcal{P}_{\overline{\mathbf{X}},\underline{J}}$ and let p' denote the corresponding path in $\mathcal{P}$. If p' starts with an edge that is not into J in $\mathcal{P}$, then the corresponding path in $\mathcal{M}$ starts with an invisible edge out of J. If so, then we add a latent confounder to the invisible edge out of J in the construction of $\mathcal{G}$ using Lemma 10. Otherwise, p' is into J in $\mathcal{P}$ and consequently $\mathcal{M}$, and we construct the causal diagram $\mathcal{G}$ similar to **case 2a**. The rest follows like in **case 2a** to show the presence of a hedge and that every node in $\mathbf{Z}$ involved in the hedge has a backdoor active path to $\mathbf{Y}$ through J.

Case 2c: The last case is that $\mathbf{R}' \cap \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus(\mathbf{X}\cup\mathbf{Z})}} = \emptyset$, and the buckets in $\mathcal{P}$ corresponding to the nodes in $\mathbf{R}'$ do not intersect with $\mathbf{Z}$. The latter condition holds due to Lines 8-9 of Alg. 2 and the absence of a node J satisfying the conditions in **Case 2b**. Fix $\mathbf{B} \subseteq \mathbf{Z}$ as a minimal set such that $\mathbf{R}' \subseteq \texttt{PossAn}(\mathbf{B})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{\mathbf{Y}},\underline{\mathbf{Z}}}}$, and let I be an arbitrary node in $\mathbf{B}$. If $(\{I\} \perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{I\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{I}}}$, then Lines 8-9 of Alg. 2 dictate the existence of a node $J \in \mathbf{Z}$ in the same bucket as I where $(\{J\} \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{J\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{J}}}$. Then, by the property in Lemma 2, we can select J instead of I to be in $\mathbf{B}$ and still maintain $\mathbf{R}' \subseteq \texttt{PossAn}(\mathbf{B})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{\mathbf{Y}},\underline{\mathbf{Z}}}}$. Using Lemma 6, we construct a MAG $\mathcal{M}$ in the equivalence class of $\mathcal{P}$ with no additional arrowheads into all the nodes in $\mathcal{R}'$ and J. This is possible since each of those nodes is in a different bucket in $\mathcal{P}$. By Lemma 4 and the above construction of the MAG, there exists a directed path from each node in $\mathcal{R}'$ to some node in $\mathbf{B}$ such that the directed paths do not go through $\mathbf{X}$. Next, we construct a causal diagram $\mathcal{G}$ via Lemma 10 similar to **case 2b** while ensuring that the active path between J and $\mathbf{Y}$ in $\mathcal{G}$ is either into J or out of J and the edge is confounded. The rest follows as in previous cases to establish a hedge and the presence a backdoor active path from each node in $\mathbf{Z}$ involved in the hedge to $\mathbf{Y}$. This concludes the proof. $\square$

Lemma 22. *Right before executing Line 10 of Alg. 2, let* $\mathbf{D} = \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$. *Then, every bucket* $\mathbf{B}$ *in* $\mathcal{P}$ *is such that either* $\mathbf{B} \subseteq \mathbf{D}$ *or* $\mathbf{B} \cap \mathbf{D} = \emptyset$.

Proof. After the while loop at Lines 2-7, every bucket $\mathbf{B}$ in $\mathcal{P}$ satisfies the condition in the lemma. Then any violation of the condition must be due to Lines 8-9. In what follows, we show that the routine at Lines 8-9 does not violate the condition.

Suppose for the sake of contradiction that the condition of the lemma is violated right before executing Line 10. Then, we have two nodes A, B in the same bucket in $\mathcal{P}$ such that $A \in \mathbf{X}$, $B \in \texttt{PossAn}(\mathbf{Y} \cup \mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$, and A, B are adjacent with a circle edge $A \circ\!\!-\!\!\circ B$. Now A does not belong to the initial input set $\mathbf{X}$ for the following reason. We note that Lines 8-9 only move nodes from the conditioning to the intervention set. So, at Line 10, if there exists a possibly directed path from A to $\mathbf{Y} \cup \mathbf{Z}$ through B such that A is the only node along this path that is in $\mathbf{X}$, then the same path exists in $\mathcal{P}$ before executing the routine at Lines 8-9. This is not possible since the condition in the lemma holds after the while loop at Lines 2-7.

Alternatively, assume for the sake of contradiction that A is a node in the conditioning set that was moved to the intervention set while executing the loop at Line 8. We exhaust the cases and show that all lead to contradictions which asserts the lemma.

Case 1: $B \in \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus(\mathbf{X}\cup\mathbf{Z})}}$. Hence, we have a possibly directed path from A to $Y^* \in \mathbf{Y}$, denoted p, that starts with an invisible edge. It follows by Lemma 23 that a subsequence of p constitutes an uncovered possibly directed path from A to Y^* that does not start with a visible edge out of A. The same path is a definitely m-connecting path between A and Y^* given $\mathbf{X} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{X}},\underline{\mathbf{Z}}}$

and we have $(A \not\perp\!\!\!\perp Y^* | \mathbf{X} \cup \mathbf{Z} \setminus \{A\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{A}}}$. This contradicts the assumption that A was moved from conditioning to intervention because it satisfies the condition at Line 8.

Case 2: $B \in \texttt{PossAn}(\mathbf{Z})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$. Then there is a possibly directed path from A to $Z \in \mathbf{Z}$, denoted p, that starts with an invisible edge and such that Z is the only node along the path that is in $\mathbf{Z}$. By Lemma 23, a subsequence of p, denoted p', constitutes an uncovered possibly directed path from A to Z that does not start with a visible edge out of A. Since $Z \in \mathbf{Z}$ and has not been moved to the intervention set, then (1) A and Z are in different buckets in $\mathcal{P}$ and (2) there exists a node Z^* in the same bucket of Z such that $(Z^* \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{Z^*\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{Z^*}}}$. Point (1) implies that p' is into Z by Lemma 5. Let $\langle A = V_0, \dots, V_{k-2}, V_{k-1}, V_k = Z\rangle$ denote the sequence of nodes along p'. Since $V_{k-1} {*\!\!\rightarrow} Z$ is in $\mathcal{P}$ and Z, Z^* are in the same bucket, then $V_{k-1} {*\!\!\rightarrow} Z^*$ is in $\mathcal{P}$. Also, $V_{k-1} {*\!\!\rightarrow} Z^*$ is not into V_{k-1} else we have a possibly directed path $\langle V_{k-1}, Z, \dots, Z^*\rangle$ and $V_{k-1} {*\!\!\rightarrow} Z^*$ which contradicts Lemma 3. Then, the subpath $p'(A, V_{k-1})$ concatenated with $V_{k-1} {*\!\!\rightarrow} Z^*$ forms a possibly directed path, denoted p^*, that starts with an invisible edge. Since $p^*(A, V_{k-1})$ is uncovered, then p^* is uncovered if V_{k-2} and Z^* are not adjacent. Suppose for the sake of contradiction that V_{k-2} and Z^* are adjacent. Then, $V_{k-2} {*\!\!\rightarrow} Z^*$ is in $\mathcal{P}$ else we have a possibly directed path $\langle Z^*, V_{k-2}, V_{k-1}\rangle$ and $V_{k-1} {*\!\!\rightarrow} Z^*$ which contradicts Lemma 3. By Lemma 1, V_{k-2} and Z are adjacent as well which contradicts p' being uncovered. Therefore, p^* is an uncovered possibly directed path from A to Z^* that starts with an invisible edge. Recall we have $(Z^* \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{Z^*\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{Z^*}}}$, so let $p^\dagger$ denote the path in $\mathcal{P}$ corresponding to a definitely m-connecting path between Z^* and $Y \in \mathbf{Y}$ in $\mathcal{P}_{\overline{\mathbf{X}},\underline{Z^*}}$. We have two sub-cases to consider:

Case 2a: $p^\dagger$ is into Z^* in $\mathcal{P}$. Consider the concatenation of p^* and $p^\dagger$ given $\mathbf{X} \cup \mathbf{Z} \setminus \{A\}$. Path p^* is uncovered possibly directed from A to Z^*, starts with an invisible edge, and is proper. Also, Z^* is a collider along the concatenated path. Hence, the concatenated path is active given $\mathbf{X} \cup \mathbf{Z} \setminus \{A\}$ in $\mathcal{P}_{\overline{\mathbf{X}},\underline{A}}$ and we have $(A \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{A\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{A}}}$. This contradicts that assumption that A satisfies the condition in Line 8 of Alg. 2 and has been moved from the conditioning to the intervention set concluding this case.

Case 2b: $p^\dagger$ is not into Z^* in $\mathcal{P}$. Let L be the node along p^* that is closest to Z^* and let $J {*\!\!\rightarrow} K$ denote such an edge along $p^\dagger$ that is closest to Z^*. Since $p^\dagger$ is of definite status and not into Z^*, then $J {*\!\!\rightarrow} K$, if defined, is either $J \rightarrow K$ or $J {\circ\!\!\rightarrow} K$ (,and not $J \leftrightarrow K$). We have three sub-cases to consider:

(i) If $J {*\!\!\rightarrow} K$ is not defined, then $p^\dagger$ is an uncovered circle path between Z^* and Y. Since $L {*\!\!\rightarrow} Z^*$ is in $\mathcal{P}$, L is adjacent to and into every node along $p^\dagger$ including Y by Lemma 1. Also, $L {*\!\!\rightarrow} Y$ is not bi-directed since the path $\langle L, Z^*, \dots, Y\rangle$ is possibly directed and having $L \leftrightarrow Y$ in $\mathcal{P}$ contradicts the property in Lemma 3. Let $\mathcal{M}$ be the MAG constructed following Lemma 6 such that no additional edges are into A. The corresponding subpath of $p^*(A, L)$ in $\mathcal{M}$ is directed out of A and starts with an invisible edge due to Lemma 7 and the fact that every non-endpoint node along p^* is a non-collider. Also, the edge between L and Y is directed out of L in $\mathcal{M}$ following the construction in Lemma 6. It follows, in $\mathcal{M}$, we have a directed path from A to Y that starts with an invisible edge and does not include, except for A, any node in $\mathbf{X} \cup \mathbf{Z}$. Hence, the same path is active given $\mathbf{X} \cup \mathbf{Z} \setminus \{A\}$ in $\mathcal{M}_{\overline{\mathbf{X}},\underline{A}}$.Then, by Lemma 15, we have $(A \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{A\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{A}}}$, a contradiction.

(ii) $J {*\!\!\rightarrow} K$ is defined and $J \neq Z^*$. We note that every edge before $J {*\!\!\rightarrow} K$ along $p^\dagger$ is a circle edge. Then L is adjacent to and into every node along the subpath $p^\dagger(Z^*, J)$ by Lemma 1. Let $\mathcal{M}$ be the MAG constructed following Lemma 6 such that no additional edges are into A. Similar to the argument in **(i)**, we have a directed path from A to J that starts with an invisible edge and does not include, except for A, any node in $\mathbf{X} \cup \mathbf{Z}$. Also, the edge between J and K is directed out of J in $\mathcal{M}$ due to the construction in Lemma 6 and the earlier conclusion that $J \leftrightarrow K$ is not in $\mathcal{P}$. Next, consider the concatenation of the directed path from A to J, $J \rightarrow K$, and the corresponding subpath $p^\dagger(K, Y)$, possibly $K = Y$. By definition of $J {*\!\!\rightarrow} K$ in $\mathcal{P}$, J is a non-collider along $p^\dagger$ and $J \notin \mathbf{Z}$ else $p^\dagger$ is not active. It follows that the concatenated path is active given $\mathbf{X} \cup \mathbf{Z} \setminus \{A\}$ in $\mathcal{M}_{\overline{\mathbf{X}},\underline{A}}$. Then, by Lemma 15, we have $(A \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{A\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{A}}}$, a contradiction.

(iii) $J {*\!\!\rightarrow} K$ is defined and $J = Z^*$. Recall that the first edge along $p^\dagger$, i.e., $J {*\!\!\rightarrow} K$, is not a directed visible edge. Since $L {*\!\!\rightarrow} Z^* (= J)$ is in $\mathcal{P}$ and $J {*\!\!\rightarrow} K$ is $J {\circ\!\!\rightarrow} K$ or invisible $J \rightarrow K$, then L and K are adjacent in $\mathcal{P}$ by Lemma 1 or [Zhang, 2008a, Def. 8], respectively. Also, the edge between L and K is into K else we have a possibly directed path $\langle K, L, J\rangle$ and $J {*\!\!\rightarrow} K$ which contradicts

Lemma 3. Let $\mathcal{M}$ be the MAG constructed following Lemma 6 such that no additional edges are into A. In $\mathcal{M}$, similar to **(ii)**, we argue that the concatenation of the directed path from A to K and the corresponding subpath $p^\dagger(K, Y)$, possibly $K = Y$, is active given $\mathbf{X} \cup \mathbf{Z} \setminus \{A\}$ in $\mathcal{M}_{\overline{\mathbf{X}},\underline{A}}$. Then, by Lemma 15, we have $(A \not\perp\!\!\!\perp \mathbf{Y} | \mathbf{X} \cup \mathbf{Z} \setminus \{A\})_{\mathcal{P}_{\overline{\mathbf{X}},\underline{A}}}$, a contradiction. □

Lemma 23. *Let X and Y be distinct nodes in a PAG $\mathcal{P}$ such that there is a possibly directed path p from X to Y that does not start with a visible edge out of X. Then, a subsequence of p constitutes an uncovered possibly directed path from X to Y that does not start with a visible edge out of X.*

Proof. Let p^* be a shortest subsequence of p such that p^* is also a possibly directed path from X to Y in $\mathcal{P}$ that does not start with a visible edge out of X. We denote p^* by the sequence $\langle X = V_0, V_1, \dots, V_r = Y \rangle$, $r \geq 1$. If p^* is a definite status path in $\mathcal{P}$, then we are done.

Else, p^* is not of definite status in $\mathcal{P}$ and $r \geq 2$. Note that by the choice of p^*, $p^*(V_1, Y)$ is a shortest possibly directed path from V_1 to Y in $\mathcal{P}$. Hence, it is uncovered by Lemma 4. If X and V_2 are not adjacent or $V_1 \rightarrow V_2$ is in $\mathcal{P}$, then V_1 is of definite status and the lemma holds. In what follows, we show that the alternative options lead to a contradiction which concludes the proof.

Suppose for the sake of contradiction that X and V_2 are adjacent and the edge between V_1 and V_2 is not out of V_1, i.e., $V_1 \circ\!\!-\!\!\circ V_2$ or $V_1 \circ\!\!\rightarrow V_2$. First, the edge between X and V_2 is a visible edge out of X for the following reason. Since $\langle X, V_1, V_2 \rangle$ is possibly directed in $\mathcal{P}$, then $X \leftarrow\!\!* V_2$ cannot exist as it violates the property in Lemma 3. Also, if the edge between X and V_2 is invisible, then we contradict the choice of p^* as the shortest subsequence of p that is possibly directed from X to Y and does not start with a visible edge out of X. Since $X \rightarrow V_2$ is visible, there is a node D that is not adjacent to V_2 and such that (1) $D *\!\!\rightarrow X$ is in $\mathcal{P}$, or (2) there is a collider path $\langle D_0, D_1, \dots, D_k, X \rangle$, $k \geq 1$, that is into X in $\mathcal{P}$ such that every D_i, $1 \leq i \leq k$ is a parent of V_2. We consider these cases separately and show that we arrive at a contradiction.

(1) Since $D *\!\!\rightarrow X$ and $X \circ\!\!-\!\!\circ V_1$, $X \circ\!\!\rightarrow V_1$, or $X \rightarrow V1$ is invisible in $\mathcal{P}$, by Lemma 1 and the graphical condition of visibility [Zhang, 2008a, Def. 8], an edge between D and V_1 is in $\mathcal{P}$. This edge is into V_1, otherwise both a possibly directed path $\langle X, V_1, D \rangle$ and $D *\!\!\rightarrow X$ are in $\mathcal{P}$ (contrary to Lemma 3). Then, $D *\!\!\rightarrow V_1 \circ\!\!-\!\!* V_2$ is in $\mathcal{P}$ and Lemma 1 implies that D and V_2 are adjacent, a contradiction.

(2) First, $D_i \leftrightarrow V_1$, $1 \leq i \leq k$, does not exist in $\mathcal{P}$ for the following reason. Suppose for the sake of contradiction that such an edge does exist in $\mathcal{P}$ for some D_j, $1 \leq j \leq k$. Then, $\langle D_0, \dots, D_j, V_1, V_2 \rangle$ forms a discriminating path for V_1 and we have $V_1 \rightarrow V_2$ in $\mathcal{P}$ by [Zhang, 2008b, FCI:R4], a contradiction. Note that we rule out the case of $V_1 \leftrightarrow V_2$ in $\mathcal{P}$ by [Zhang, 2008b, FCI:R4] since the edge is along p^*, a possibly directed path from X to Y. Next, we argue by induction that $D_i *\!\!\rightarrow V_1$ is in $\mathcal{P}$ for all $0 \leq i \leq k$. In the base case, D_k is adjacent to V_1 else we have $X \rightarrow V_1$ by [Zhang, 2008a, FCI:R1] and $X \rightarrow V_1$ is visible by [Zhang, 2008a, Def. 8], a contradiction. Also, if the edge between D_k and V_1 is not into V_1, then we have a possibly directed path $\langle X, V_1, D_k \rangle$ and $D_k \leftrightarrow X$ which contradicts Lemma 3. In the induction step, we assume the property holds for $j + 1 \leq i \leq k$, and prove it for D_j. If D_j and V_1 are not adjacent, $\langle D_j, D_{j+1}, V_1 \rangle$ is an uncovered triple and $D_{j+1} \rightarrow V_1$ is in $\mathcal{P}$ by [Zhang, 2008b, FCI:R1] (the edge cannot be bi-directed as argued earlier). It follows that $\langle D_j, D_{j+1}, D_{j+2}, V_1 \rangle$ forms a discriminating path for D_{j+2} and we have $D_{j+2} \rightarrow V_1$ in $\mathcal{P}$ by [Zhang, 2008b, FCI:R4]. The last argument applies recursively for each D_i, $j + 2 \leq i \leq k$, where $\langle D_j, \dots, D_i, V_1 \rangle$ forms a discriminating path for D_i and we have $D_i \rightarrow V_1$ in $\mathcal{P}$. Therefore, $\langle D_j, \dots, D_k, X, V_1 \rangle$ forms a discriminating path for X and we have $X \rightarrow V_1$ in $\mathcal{P}$. We rule out $X \leftrightarrow V_1$ again since the edge is along p^*, a possibly directed path from X to Y. Also, the collider path $\langle D_j, \dots, D_k, X \rangle$ implies that $X \rightarrow V_1$ is visible by [Zhang, 2008a, Def. 8]. This contradicts the choice of p^* where the edge between X and V_1 is invisible, thus D_j is adjacent to V_1. If the edge between D_j and V_1 is not into V_1, then we have a possibly directed path $\langle D_{j+1}, V_1, D_j \rangle$ and $D_j *\!\!\rightarrow D_{j+1}$ which contradicts Lemma 3. Finally, consider the uncovered triple $\langle D_0, V_1, V_2 \rangle$ where $D_0 *\!\!\rightarrow V_1$ is in $\mathcal{P}$ by the induction argument. Then, we have $V_1 \rightarrow V_2$ in $\mathcal{P}$ by [Zhang, 2008b, FCI:R1] and we reach a contradiction since the edge between V_1 and V_2 is assumed not to be out of V_1. This concludes the proof. □

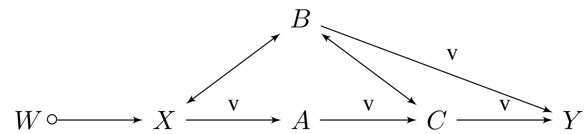

Figure 9: A query $P_x(\mathbf{v} \setminus \{x\})$ where the criterion in Prop. 8 is applicable and that in Prop. 2 is not.

F Comparison between the IDP Versions

Jaber et al. [2018a] derived the identification criterion shown in Prop. 8 where the intervention is on a bucket rather than a single node and the input distribution is possibly interventional. Note that the criterion in Prop. 8 is equivalent to testing the condition $C^{\mathbf{X}} \cap \texttt{PossCh}(\mathbf{X}) \subseteq \mathbf{X}$, by Lemma 24. Obviously, the condition in Prop. 2, i.e, $C^{\mathbf{X}} \cap \texttt{PossDe}(\mathbf{X}) \subseteq \mathbf{X}$, implies $C^{\mathbf{X}} \cap \texttt{PossCh}(\mathbf{X}) \subseteq \mathbf{X}$; however, the opposite is not true. In words, if a bucket's pc-component does not intersect with its possible descendants except for itself, i.e., the bucket, then the bucket's pc-component does not intersect with its possible children except for itself. Therefore, the criterion in Prop. 8 is superior to that in Prop. 2 in that the earlier allows us to identify more causal effects than the latter. On the other hand, whenever the causal effect is identifiable by Prop. 8, the resulting expression is convoluted and usually large as we see in Eq. 12. The following example illustrates the above points.

Proposition 8. *Let $\mathcal{P}$ denote a PAG over $\mathbf{V}$, $\mathbf{T}$ be a union of a subset of the buckets in $\mathcal{P}$, and $\mathbf{X} \subset \mathbf{T}$ be a bucket. Given $P_{\mathbf{v} \setminus \mathbf{t}}$ (i.e., $Q[\mathbf{T}]$), and a partial topological order $\mathbf{B}_1 < \cdots < \mathbf{B}_m$ with respect to $\mathcal{P}_{\mathbf{T}}$, $Q[\mathbf{T} \setminus \mathbf{X}]$ is identifiable if and only if, in $\mathcal{P}_{\mathbf{T}}$, there does not exist $Z \in \mathbf{X}$ such that Z has a possible child $C \notin \mathbf{X}$ that is in the pc-component of Z. If identifiable, then the expression is given by*

$$Q[\mathbf{T} \setminus \mathbf{X}] = \frac{P_{\mathbf{v} \setminus \mathbf{t}}}{\prod_{\{i | \mathbf{B_i} \subseteq S^{\mathbf{X}}\}} P_{\mathbf{v} \setminus \mathbf{t}}(\mathbf{B_i} | \mathbf{B^{(i-1)}})} \times \sum_{\mathbf{x}} \prod_{\{i | \mathbf{B_i} \subseteq S^{\mathbf{X}}\}} P_{\mathbf{v} \setminus \mathbf{t}}(\mathbf{B_i} | \mathbf{B^{(i-1)}}), \tag{12}$$

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

Consider PAG $\mathcal{P}$ in Fig. 9 and the causal query $P_x(\mathbf{v} \setminus \{x\})$. In $\mathcal{P}$, we have $C^X = \{W, X, B, C\}$ as the set of nodes in the pc-component of X, $\texttt{PossCh}(X) = \{X, A\}$ as the possible children of X, and $\texttt{PossDe}(X) = \{X, A, C, Y\}$ as the possible descendants of X. First, we try to use the criterion in Prop. 2 to compute $P_x(\mathbf{v} \setminus \{x\}) = Q[\mathbf{V} \setminus \{X\}]$ from $P(\mathbf{V}) = Q[\mathbf{V}]$; however, $C^X \cap \texttt{PossDe}(X) = \{X, C\} \nsubseteq \{X\}$ and the criterion is not applicable. However, $C^X \cap \texttt{PossCh}(X) = \{X\}$ and the criterion in Prop. 8 is applicable to compute $Q[\mathbf{V} \setminus \{X\}]$ from $P(\mathbf{V})$. The dc-component of X in $\mathcal{P}$ is $S^X = \{X, B, C\}$ and we assume the following partial topological order $W < X < B < A < C < Y$. Accordingly, we have the following expression.

$$P_x(\mathbf{v} \setminus \{x\}) = \frac{P(\mathbf{v})}{P(x|w)P(b|w,x)P(c|w,x,b,a)} \times \sum_x P(x|w)P(b|w,x)P(c|w,x,b,a)$$

Alternatively, if we use the partial order $W < B < X < A < C < Y$, we get the following expression which can be simplified to Eq. 14 since the term $P(b|w)$ is independent of X and can be cancelled out with the corresponding term in the denominator. This shows that the complexity of the expression is partially dependent on the choice of the partial order, among other factors.

$$P_x(\mathbf{v} \setminus \{x\}) = \frac{P(\mathbf{v})}{P(b|w)P(x|w,b)P(c|w,b,x,a)} \times \sum_x P(b|w)P(x|w,b)P(c|w,b,x,a) \tag{13}$$

$$= \frac{P(\mathbf{v})}{P(x|w,b)P(c|w,b,x,a)} \times \sum_x P(x|w,b)P(c|w,b,x,a) \tag{14}$$

The above example illustrates that the criterion in Prop. 8 is more powerful than that in Prop. 2. However, the expression is somewhat convoluted and possibly intractable. This earlier criterion is utilized in [Jaber et al., 2019a] to formulate an algorithm for marginal effect identification as shown in Alg. 3. This algorithm is almost identical to the one proposed in Alg. 1 except for line 6 where it tests for the criterion in Prop. 8 as opposed to Prop. 2 in Alg. 1. Interestingly, the proposed version in Alg. 1 remains as expressive as Alg. 3 and complete, by Thm. 3, despite using the weaker criterion in

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

Input: PAG $\mathcal{P}$ and two disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$
Output: Expression for $P_{\mathbf{x}}(\mathbf{y})$ or `FAIL`

1: Let $\mathbf{D} = \texttt{PossAn}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V}\setminus\mathbf{X}}}$
2: **return** $\sum_{\mathbf{d}\setminus\mathbf{y}}$ IDENTIFY($\mathbf{D}, \mathbf{V}, P$)
3: **function** IDENTIFY($\mathbf{C}$, $\mathbf{T}$, $Q = Q[\mathbf{T}]$)
4:     **if** $\mathbf{C} = \emptyset$ **then return** 1
5:     **if** $\mathbf{C} = \mathbf{T}$ **then return** Q
/* *In* $\mathcal{P}_{\mathbf{T}}$, *let* $\mathbf{B}$ *denote a bucket, and let* $C^{\mathbf{B}}$ *denote the pc-component of* $\mathbf{B}$ */
6:     **if** $\exists \mathbf{B} \subset \mathbf{T} \setminus \mathbf{C}$ such that $C^{\mathbf{B}} \cap \texttt{PossCh}(\mathbf{B})_{\mathcal{P}_{\mathbf{T}}} \subseteq \mathbf{B}$ **then**
7:         Compute $Q[\mathbf{T} \setminus \mathbf{B}]$ from Q; ▷ via Prop. 8
8:         **return** IDENTIFY($\mathbf{C}$,$\mathbf{T} \setminus \mathbf{B}$,$Q[\mathbf{T} \setminus \mathbf{B}]$)
9:     **else if** $\exists \mathbf{B} \subset \mathbf{C}$ such that $\mathcal{R}_{\mathbf{B}} \neq \mathbf{C}$ **then** ▷ $\mathcal{R}_{\mathbf{B}}$ is equivalent to $\mathcal{R}^{\mathbf{C}}_{\mathbf{B}}$
10:         **return** $\dfrac{\text{IDENTIFY}(\mathcal{R}_{\mathbf{B}},\mathbf{T},Q) \times \text{IDENTIFY}(\mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_{\mathbf{B}}},\mathbf{T},Q)}{\text{IDENTIFY}(\mathcal{R}_{\mathbf{B}}\cap\mathcal{R}_{\mathbf{C}\setminus\mathcal{R}_{\mathbf{B}}},\mathbf{T},Q)}$
11:     **else** **throw** `FAIL`

Prop. 2. The advantage of using Prop. 2 is dealing with a simpler expression as shown in Eq. 1. This is illustrated in the following example.

Consider again PAG $\mathcal{P}$ in Fig. 9 and causal query $P_x(y)$. We follows the steps in Alg. 1 to try and identify the target effect. We have $\mathbf{D} = \{Y, A, B, C\}$ and $P_x(y) = \sum_{a,b,c} Q[\mathbf{V} \setminus \{W, X\}]$. Then, we call IDENTIFY$(\cdot)$ to compute $Q[\mathbf{V} \setminus \{W, X\}]$ from $P(\mathbf{V})$. In IDENTIFY$(\cdot)$, we first check in line 6 if there exists a bucket, a singleton node in this case, in $\{W, X\}$ that satisfies the criterion in Prop. 2. Node X does not satisfy the criterion as discussed in the earlier example and node W is in the same pc-component with its possible child/descendant X. Next we try to decompose the query in line 9 using Prop. 5. Assuming $\mathbf{B}$ in line 9 is equal to $\{A\}$, we have $\mathcal{R}^{\mathbf{D}}_{\mathbf{B}} = \{A\} \neq \mathbf{D}$ and we get the decomposition $Q[\mathbf{V} \setminus \{W, X\}] = Q[A] \cdot Q[\{B, C, Y\}]$ since $\mathcal{R}_{\mathbf{B}} \cap \mathcal{R}_{\mathbf{D}\setminus\mathcal{R}_{\mathbf{B}}} = \{A\} \cap \{B, C, Y\} = \emptyset$. Subsequently, **IDP** calls IDENTIFY$(\cdot)$ at line 10 to compute $Q[A]$ and $Q[\{B, C, Y\}]$ from $P(\mathbf{V})$. Starting with $Q[A]$, node Y trivially satisfies the criterion of Prop. 2 as it has no descendants and we compute $Q[\mathbf{V} \setminus \{Y\}] = \frac{P(\mathbf{v})}{P(y|\mathbf{v}\setminus\{y\})} = P(\mathbf{v} \setminus \{y\})$. Similarly in the subsequent calls to IDENTIFY$(\cdot)$, we intervene on C and B to obtain $Q[W, X, A] = P(w, x, a)$. Next, X satisfies the criterion of Prop. 2 in the induced subgraph $\mathcal{P}_{W,X,A}$ and $Q[W, A] = \frac{P(w,x,a)}{P(x|w)} = P(w) \times P(a|w, x)$. Finally W satisfies the criterion and we get $Q[A] = \frac{P(w)\times P(a|w,x)}{\frac{P(w)\times P(a|w,x)}{\sum_w P(w)\times P(a|w,x)}} = \sum_w P(w) \times P(a|w, x)$. Similarly, **IDP** computes $Q[B, C, Y]$ by intervening on A, X, then W. Hence, we get the expression $Q[B, C, Y] = \sum_{w,x} P(y, c|w, x, a, b) \times P(w, x, b)$. The final expression for $P_x(y)$ is shown below.

$$P_x(y) = \sum_{a,b,c} Q[A] \cdot Q[B, C, Y] = \sum_{a,b,c} \big(\sum_{w} P(w) \cdot P(a|w, x)\big)\big(\sum_{w,x} P(y, c|w, x, a, b) \cdot P(w, x, b)\big)$$

Lemma 24. *Let $\mathcal{P}$ denote a PAG over $\mathbf{V}$, let $\mathbf{T} \subseteq \mathbf{V}$ be a subset of the buckets in $\mathcal{P}$, and let $\mathbf{X} \subset \mathbf{T}$ be a bucket. Then, the following two conditions are equivalent:*

1. *$C^{\mathbf{X}} \cap \texttt{PossCh}(\mathbf{X}) \subseteq \mathbf{X}$.*
2. *there does not exist $Z \in \mathbf{X}$ such that Z has a possible child $C \notin \mathbf{X}$ that is in the pc-component of Z.*

Proof. $1 \implies 2$. This direction follows trivially.

$2 \implies 1$. We prove this by contrapositive. Suppose $C^{\mathbf{X}} \cap \texttt{PossCh}(\mathbf{X}) \not\subseteq \mathbf{X}$. If there exists a node $Z \in \mathbf{X}$ that has a possible child $C \notin \mathbf{X}$ such that $Z \circ\!\!-\!\!\circ C$, $Z \circ\!\!\rightarrow C$, or an invisible $Z \rightarrow C$, then C is in the pc-component of Z which concludes the proof.

Alternatively, assume that every possible child of $\mathbf{X}$ is due to a visible edge out of $\mathbf{X}$. Following the initial assumption ($C^{\mathbf{X}} \cap \texttt{PossCh}(\mathbf{X}) \not\subseteq \mathbf{X}$), there exists at least one child of $\mathbf{X}$, denoted C,

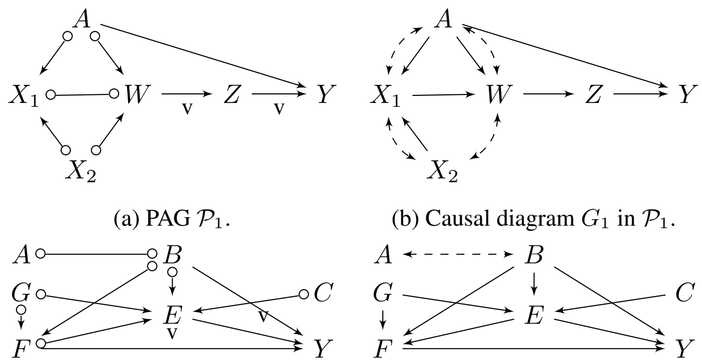

(a) PAG $\mathcal{P}_1$. (b) Causal diagram G_1 in $\mathcal{P}_1$. (c) PAG $\mathcal{P}_2$. (d) Causal diagram G_2 in $\mathcal{P}_2$.

Figure 10: (a) $\mathcal{P}_1$, a PAG (same as $\mathcal{P}$ in Fig. 3a); (b) G_1, a causal diagram in the Markov equivalence class of $\mathcal{P}_1$; (c) $\mathcal{P}_2$, a PAG (same as $\mathcal{P}$ in Fig. 5a); (d) G_2, a causal diagram in the Markov equivalence class of $\mathcal{P}_2$.

($X_i \rightarrow C$) such that the edge is visible and C is also in the pc-component of $\mathbf{X}$. Next, we prove that C is in the pc-component of X_i which concludes the proof. Let $X_j \neq X_i$ denote the node in $\mathbf{X}$ such that X_j and C are in the same pc-component. First, X_j and C can not be in the same pc-component due to $X_j \circ\!\!-\!\!\circ C$, $C \circ\!\!\rightarrow X_j$, or invisible $C \rightarrow X_j$ as it violates a PAG property [Maathuis and Colombo, 2015, Lemma 7.5]. Also, X_j and C can not be in the same pc-component due to an invisible $X_j \rightarrow C$ since we assume that every possible child of $\mathbf{X}$ is due to a visible edge out of $\mathbf{X}$. Hence, X_j and C are in the same pc-component due to a collider path p consistent with [Jaber et al., 2019a, Def. 4]. The first edge along p starting from X_j has to be into X_j ($X_j \leftrightarrow A$), else X_j has a possible child (A) which is also in the pc-component of $\mathbf{X}$ violating our assumption that every possible child of $\mathbf{X}$ is due to a visible edge out of $\mathbf{X}$. Hence, it follows by [Zhang, 2006, Lemma 3.3.2] that there exists a bi-directed edge between A and every node in $\mathbf{X}$, including X_i. Therefore, C is a possible child of X_i and C is in the same pc-component of X_i. This concludes the proof. $\square$

G Experiments

In this section, we empirically evaluate the soundness and performance of **CIDP** (Algorithm 2). **CIDP** and the auxiliary functions were implemented in R v4.1.2 [R Core Team, 2021], using standard graph and causal inference packages, including *dagitty* v0.3.1 [Textor et al., 2016], *igraph* v1.2.8.9014 [Csardi and Nepusz, 2006], *pcalg* v2.7.3 [Hauser and Bühlmann, 2012], and *causaleffect* v1.3.13 [Tikka and Karvanen, 2017]. The experiments were performed on an Intel® Core™ i9 processor at 2.3 GHz with 16 GB RAM. The R package will be made freely available.

G.1 Empirical Evaluation of the CIDP Algorithm

We first empirically evaluate **CIDP** (Alg. 2) in its ability to soundly identify the interventional distribution $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ from a PAG $\mathcal{P}$. Given a causal diagram $\mathcal{G}$ in the equivalence class of $\mathcal{P}$, we randomly generated 30 datasets of binary variables according to $\mathcal{G}$ consisting of N observations, with $N = \{5000, 10000, 50000, 100000, 500000\}$. This was performed using the function `simulateLogistic` of the dagitty R package. Then we generated the PAG using the FCI algorithm [Zhang, 2008b] with the true oracle and estimated $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z})$ for all possible values of $\mathbf{x}, \mathbf{y}$, and $\mathbf{z}$ using both the identification formula given $\mathcal{G}$, provided by the conditional ID algorithm [Tian, 2004, Shpitser and Pearl, 2006] and the identification formula given $\mathcal{P}$, provided by **CIDP**. Fig. 11 shows the average absolute differences and standard errors in the estimates of the conditional query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) \coloneqq P_{x_1,x_2,w}(y|z)$ given the causal diagram $\mathcal{G}_1$ in Fig. 10b and given the corresponding PAG $\mathcal{P}_1$ in Fig. 10a (same as PAG $\mathcal{P}$ in Fig. 3a). Similarly, Fig. 12 shows the average absolute differences and standard errors in the estimates of the conditional query $P_{\mathbf{x}}(\mathbf{y}|\mathbf{z}) \coloneqq P_{a,f}(y|b,e)$ given the causal diagram $\mathcal{G}_2$ in Fig. 10d and given the corresponding PAG $\mathcal{P}_2$ in Fig. 10c (same as PAG $\mathcal{P}$ in Fig. 5a).

As expected, results show a negligible average difference between the estimates derived given the true causal diagram and given the corresponding PAG. Additionally, the standard errors greatly decreases as the number of samples in the dataset increases. This illustrates that **CIDP** accurately identifies target causal effects given a PAG learned from observational data whenever the causal knowledge in the PAG is sufficient for identification.

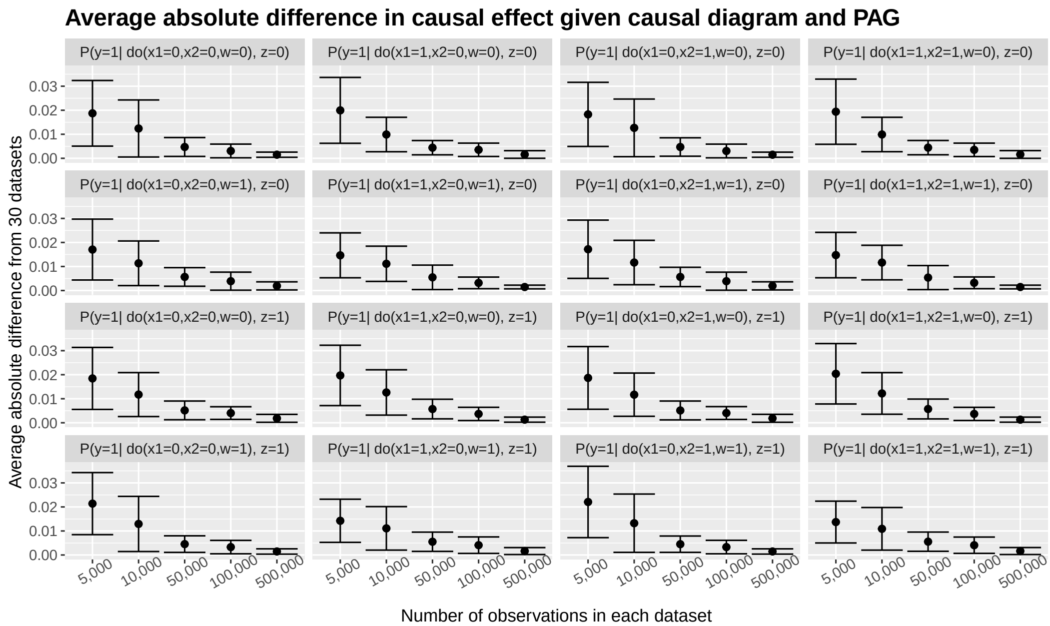

Figure 11: Average and standard deviation of the absolute difference between the estimates of $P_{x_1,x_2,w}(y|z)$ computed using the identification formula provided by **CIDP** given the PAG $\mathcal{P}$ in Fig. 3a, and using the identification formula provided by **IDC** given the true underlying causal diagram $\mathcal{G}$, a member of $\mathcal{P}$. The average is over 30 randomly generated data sets of binary variables consisting of N observations, with $N = \{5000, 10000, 50000, 100000, 500000\}$.

G.2 Comparison to Hyttinen et al. (2015)'s Algorithm

Hyttinen et al. [2015] propose an alternative approach, henceforth denoted by **HEJ**, to identify conditional causal effects $P(y|do(x), z)$ from the Markov equivalence class. **HEJ** determines effect identifiability by attempting to identify the effect in each causal diagram in the equivalence class. First, it translates the set of d-separation constraints of the equivalence class into a logical representation using an answer set programming (ASP) solver. Then, it repeatedly queries the constraint solver for a causal diagram $\mathcal{G}$ in the equivalence class and calls the **IDC** algorithm [Shpitser and Pearl, 2006] on $\mathcal{G}$ to identify the desired effect. The effect is determined as identifiable from the Markov equivalence class if all returned formulae are the same.

Even though some constraints are applied in the sampling procedure to avoid an explicit enumeration of all members of the equivalence class, **HEJ** is still very time-consuming. To witness, we next show a comparison of the running times to demonstrate the advantage in performance of **CIDP** over **HEJ**. We use the implementation of **HEJ** as provided from the first author's website at `https://www.cs.helsinki.fi/u/ajhyttin/` on April 18th, 2022, and we executed it using the ASP solver clingo v5.5.2 [Gebser et al., 2017].

In Fig. 13, we show boxplots of the running times (in seconds and in log scale) for identifying the causal effect $P(y|do(x), z)$ in 30 randomly generated PAGs with number of variables $n = \{5, 6, \ldots, 12\}$. Each PAG was generated using the true oracle corresponding to a random causal diagram with a fixed edge probability of 0.4. Running times for both algorithms do not include the time spent generating the PAG. We set a timeout of 90 minutes (5400 seconds) for each instance.

Clearly, the average running time of **HEJ** increases exponentially with the number of observed variables. Note that it ran out of time for models with more than 10 variables when the causal effect is identifiable and for models with more than 5 variables when the causal effect is not identifiable. Consistently, **CIDP** determines the identifiability of the causal effect much faster than **HEJ**. For all instances, a result is obtained in less than a second. To ensure efficiency, we implemented a linear-time algorithm for testing definite m-separability in PAGs based on the Bayes-Ball algorithm [Shachter, 1998]. A similar algorithm has been proposed by [Perković et al., 2018] for deciding

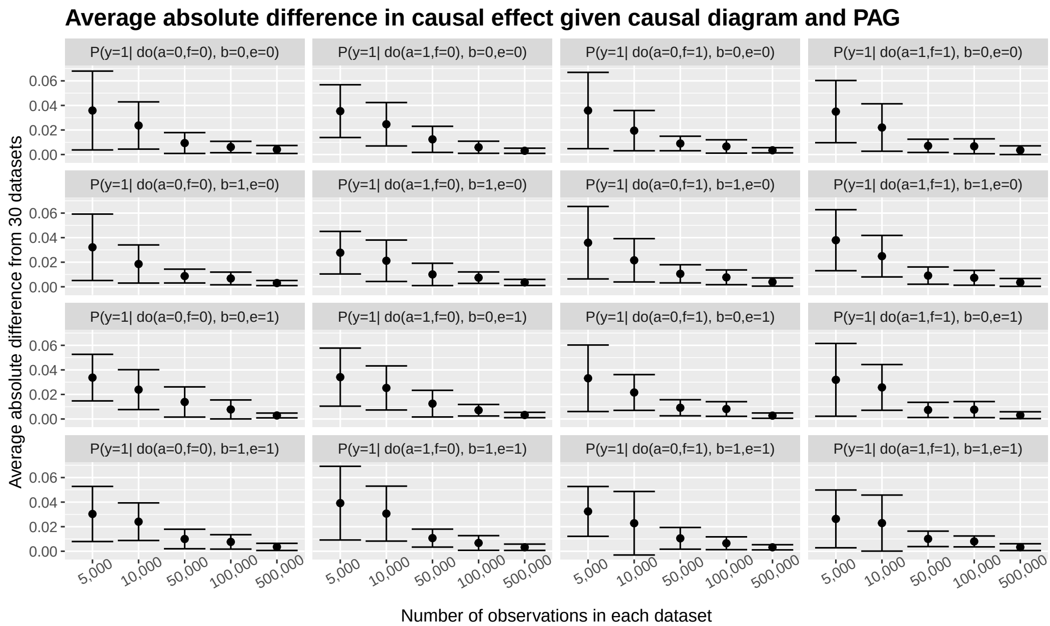

Figure 12: Average and standard deviation of the absolute difference between the estimates of $P_{a,f}(y|b,e)$ computed using the identification formula provided by **CIDP** given the PAG $\mathcal{P}$ in Fig. 5a, and using the identification formula provided by **IDC** given the true underlying causal diagram $\mathcal{G}$, a member of $\mathcal{P}$. The average is over 30 randomly generated data sets of binary variables consisting of N observations, with $N = \{5000, 10000, 50000, 100000, 500000\}$.

m-separability when testing whether a set is admissible for adjustment in PAGs and by [van der Zander et al., 2019] for deciding m-separability in MAGs.

References

E. Bareinboim and J. Pearl. Causal inference and the data-fusion problem. In R. M. Shiffrin, editor, *Proceedings of the National Academy of Sciences*, volume 113, pages 7345–7352. National Academy of Sciences, 2016.

Gabor Csardi and Tamas Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006.

R.A. Fisher. *The Design of Experiments*. Oliver and Boyd, Edinburgh, 6th edition, 1951.

Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Multi-shot ASP solving with clingo. *CoRR*, abs/1705.09811, 2017.

Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464, 2012.

Yimin Huang and Marco Valtorta. Pearl's calculus of intervention is complete. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, 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. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018a.

Amin Jaber, Jiji Zhang, and Elias Bareinboim. Causal identification under Markov equivalence: Completeness results. Technical report, R-42, CausalAI Lab, Columbia University, 2018b.

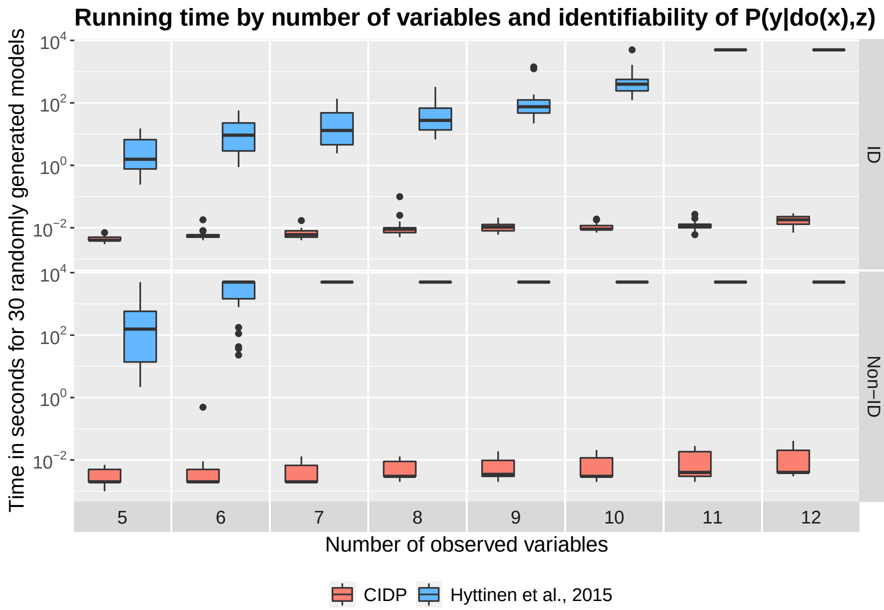

Figure 13: Boxplots of the running times (in seconds) using our algorithm, **CIDP** (red), and **HEJ** (blue), computed for 30 randomly generated models with increasing number of observed variables $n = \{5, 6, \ldots, 12\}$ and a fixed edge probability of 0.4. Results are shown separately for models where the conditional effect $P(y|do(x), z)$ is identifiable (top) and non-identifiable (bottom). Each instance has a timeout of 90 minutes (5400 seconds).

Amin Jaber, Jiji Zhang, and Elias Bareinboim. Causal identification under markov equivalence: Completeness results. In *International Conference on Machine Learning*, pages 2981–2989. PMLR, 2019a.

Amin Jaber, Jiji Zhang, and Elias Bareinboim. Identification of conditional causal effects under markov equivalence. In *33rd Annual Conference on Neural Information Processing Systems, NeurIPS 2019*. Neural Information Processing Systems Foundation, 2019b.

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

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

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

Judea Pearl. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60, 2019.

Judea Pearl and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, Inc., New York, NY, USA, 1st edition, 2018. ISBN 046509760X, 9780465097609.

Emilija Perković, Johannes Textor, and Markus Kalisch. Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs. *Journal of Machine Learning Research*, 18:1–62, 2018.

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. *Journal of Machine Learning Research*, 18(220):1–62, 2018. URL `http://jmlr.org/papers/v18/16-319.html`.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021.

Thomas Richardson and Peter Spirtes. Ancestral graph markov models. *The Annals of Statistics*, 30 (4):962–1030, 2002.

Ross D. Shachter. Bayes-ball: Rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams). In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, UAI'98, page 480–487, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc. ISBN 155860555X.

Ilya Shpitser and Judea Pearl. Identification of conditional interventional distributions. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, pages 437–444, 2006.

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

Johannes Textor, Benito van der Zander, Mark S Gilthorpe, Maciej Liśkiewicz, and George TH Ellison. Robust causal inference using directed acyclic graphs: the r package 'dagitty'. *International Journal of Epidemiology*, 45(6):1887–1894, 2016.

Jin Tian. Identifying conditional causal effects. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 561–568, 2004.

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

Santtu Tikka and Juha Karvanen. Identifying causal effects with the R package causaleffect. *Journal of Statistical Software*, 76(12):1–30, 2017.

Benito van der Zander, Maciej Liśkiewicz, and Johannes Textor. Separators and adjustment sets in causal graphs: Complete criteria and an algorithmic framework. *Artificial Intelligence*, 270:1–40, 2019.

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

Jiji Zhang. *Causal inference and reasoning in causally insufficient systems*. PhD thesis, Citeseer, 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: 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(16-17):1873–1896, 2008b.