Markov Equivalence for Ancestral Graphs*

R. Ayesha Ali

National University of Singapore, Singapore

staara@nus.edu.sg

Thomas Richardson

University of Washington, USA

tsr@stat.washington.edu

Peter Spirtes

Carnegie-Mellon University, USA

ps7z@andrew.cmu.edu

September 25, 2004

Abstract

Ancestral graph models can encode conditional independence relations that arise in directed acyclic graph (DAG) models with latent and selection variables. However, for any ancestral graph, there may be several other graphs to which it is Markov equivalent. We state and prove conditions under which two maximal ancestral graphs are Markov equivalent to each other, thereby extending analogous results for DAGs given by other authors.

*University of Washington Technical Report No. 466.

Contents

1 Introduction

A graphical Markov model is a set of distributions with independence structure described by a graph consisting of vertices and edges. The *independence model* associated with a graph is the set of conditional independence relations encoded by the graph through a *global Markov property.* In general, different graphs may encode the same independence model. In this paper we consider a particular class of graphs, called ancestral graphs, and characterize when two graphs encode the same sets of conditional independence relations.

The class of ancestral graphs is motivated in the following way: we suppose our observed data was generated by a process represented by a directed acyclic graph (DAG). The causal interpretation of such a DAG is described by Spirtes et al. (1993), and Pearl (2000). However, in general, we may only have observed a subset of these variables, in a specific subpopulation. Hence some variables in the underlying DAG are not observed ('latent'), while other variables, specifying the specific subpopulation from which our data is sampled, are conditioned upon ('selection variables').

Even though the underlying model is a DAG, the conditional independence structure holding amongst the observed variables, conditional on the selection variables cannot always be represented by a DAG containing only the observed variables. For this purpose the larger class of ancestral graphs is required. (See Figure 2 and Definition 2.2) The statistical models associated with ancestral graphs retain many of the desirable properties that are associated with DAG models. In particular, the statistical model associated with a ancestral graph consists of distributions (in a given family) that obey the global Markov property for the graph.

Like DAGs, two different ancestral graphs can represent the same set of conditional independence relations, and hence distributions. Such graphs are said to be *Markov equivalent*. A graphical characterization of the circumstances under which graphs are Markov equivalent is of importance for several reasons:

- Data cannot distinguish between Markov equivalent graphs, even asymptotically, because the set of distributions associated with the graphs are the same. Thus for the purposes of interpreting a model it is often important to characterize those features which are common to all the graphs in a given class (Spirtes et al. (1993), Meek (1995)).

- Different graphs correspond to different parameterizations. Though the set of distributions represented by the model are the same, some models may be simpler to fit than others. For example, in Figure 1(i) the graph on the left is an example of a seemingly unrelated regression (SUR) model (Zellner (1962)). In general there are no closed form expressions for the MLEs for SUR models, iterative methods are required (see Drton and Richardson (2004b)), and there may be multiple solutions to the likelihood equations (Drton and Richardson (2004c), Drton (2004)). However, the SUR model in Figure 1(i) is Markov equivalent to Figure 1(ii), which is a DAG. DAG models have closed form MLEs, and the likelihood is unimodal (see Lauritzen (1996)). Consequently none of the problems which may arise for general SUR models apply to the specific model in Figure 1(i) (see also Drton et al. (2003) and Drton and Richardson (2004a)).

- It is inefficient for a model search procedure to consider separately every graph in a given equivalence class if they are all assigned the same score (Chickering (2002a), Chickering (2002b); see also Gillispie and Perlman (2001)).

In this paper we provide necessary and sufficient graphical conditions un-

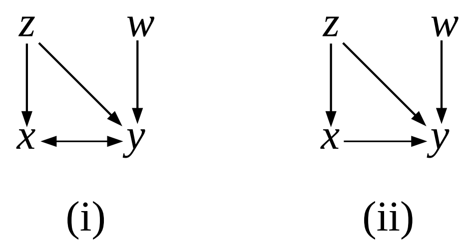

Figure 1: (i) A seemingly unrelated regression model and (ii) a Markov equivalent DAG model.

der which two ancestral graphs are Markov equivalent. Frydenberg (1990), Verma and Pearl (1991), Chickering (1995) and Andersson et al. (1997) solve the Markov equivalence problem for DAGs.

Section 2 defines the class of ancestral graphs, and outlines the motivation for the class. Section 3 contains the main result of the paper. Discussion and relation to prior work is in section 4.

2 Ancestral graphs

The basic motivation for ancestral graphs is to enable one to model the independence structure over the observed variables that results from the presence of latent or selection variables without explicitly including such variables in the model. To illustrate this, consider the DAG shown in Figure 2(i) in which Azt, Pcp, Ap and $CD4$ are observed variables, while H is unobserved. Azt and Ap may be thought of as treatments; Pcp and $CD4$ as responses correlated by underlying health status H. The DAG given in Figure 2(i) incorporates the assumption that Azt and Ap are both randomized, and further, that Azt does not affect $CD4$. The DAG implies the following conditional independence relations over the observed variables:

$$Azt \perp\!\!\!\perp Ap, CD4 \qquad Ap \perp\!\!\!\perp Azt, Pcp.$$

The ancestral graph which represents these conditional independence relations is shown in Figure 2(ii). (See §2.2 for the definition of an ancestral graph and the Markov property.) However, there is no DAG on the four observed variables which represents all and only these conditional independence relations.

As this example suggests, bi-directed ($\leftrightarrow$) edges are required in order to represent conditional independence relations resulting from unobserved parents. Likewise undirected edges (—) are required to represent children that have been conditioned on in the selected sub-population from which the sample is taken (see Cox and Wermuth (1996), Cooper (1995)). However, bi-directed and undirected edges may also arise in other contexts, where both marginalization and conditioning are present. Richardson and Spirtes (2003) provides a detailed discussion on the interpretation of edges in an ancestral graph.

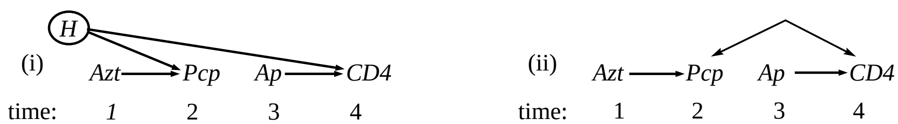

Figure 2: (i) A DAG with a latent variable H. (ii) The ancestral graph resulting from marginalizing over H adds a bi-directed edge between Pcp and $CD4$.

2.1 Basic graphical notation and terminology

We use the following terminology to describe relations between vertices in a graph $\mathcal{G}$:

$$\text{If}\left\{\begin{array}{c}\alpha - \beta \\ \alpha \leftrightarrow \beta \\ \alpha \rightarrow \beta \\ \alpha \leftarrow \beta\end{array}\right\} \text{in } \mathcal{G} \text{ then } \alpha \text{ is a} \left\{\begin{array}{c}\textit{neighbour} \\ \textit{spouse} \\ \textit{parent} \\ \textit{child}\end{array}\right\} \text{of } \beta, \text{ and} \left\{\begin{array}{c}\alpha \in \mathrm{ne}_{\mathcal{G}}(\beta) \\ \alpha \in \mathrm{sp}_{\mathcal{G}}(\beta) \\ \alpha \in \mathrm{pa}_{\mathcal{G}}(\beta) \\ \alpha \in \mathrm{ch}_{\mathcal{G}}(\beta)\end{array}\right\}.$$

A pair of vertices which are connected by some edge will be said to be *adjacent*. Note that the three edge types should be considered as distinct symbols, in particular,

$$\alpha - \beta \quad \neq \quad \alpha \underset{\leftarrow}{\overset{\rightarrow}{}} \beta \quad \neq \quad \alpha \leftrightarrow \beta.$$

The graphs we consider in this paper have at most one edge between each pair of vertices. If there is an edge $\alpha \rightarrow \beta$, or $\alpha \leftrightarrow \beta$ then there is said to be *an arrowhead at* β *on this edge*. Conversely, if there is an edge $\alpha \rightarrow \beta$, or $\alpha - \beta$ then there is said to be a *tail at* α. We do not allow a vertex to be adjacent to itself. A *path*, $\boldsymbol{\pi}$ is a sequence of distinct vertices that are adjacent. If a and b are vertices on a path $\boldsymbol{\pi}$ then the *subpath* of $\boldsymbol{\pi}$, with endpoints a and b, will be denoted by $\boldsymbol{\pi}(a, b)$.

2.2 Definition of ancestral graphs

DAGs are directed graphs in which directed cycles:

$$x \rightarrow \cdots \rightarrow x$$

are not permitted. Similarly, certain configurations of edges are not permitted in ancestral graphs:

Definition 2.1 *A graph, which may contain undirected* $(-)$*, directed* $(\rightarrow)$ *or bi-directed edges* $(\leftrightarrow)$ *is* ancestral *if:*

(a) *there are no directed cycles;*

(b) *whenever there is an edge $x \leftrightarrow y$, then there is no directed path from x to y, or from y to x;*

(c) *if there is an undirected edge $x - y$ then x and y have no spouses or parents.*

Conditions (a) and (b) may be summarized by saying that if x and y are joined by an edge and there is an arrowhead at x, then x is *not* an ancestor of y; this is the motivation for the term 'ancestral'.

A vertex α is said to be an *ancestor* of a vertex β if *either* there is a directed path $\alpha \to \cdots \to \beta$ from α to β, *or* $\alpha = \beta$. A vertex α is said to be *anterior* to a vertex β if there is a path $\boldsymbol{\mu}$ between α and β on which every edge is either of the form $\gamma - \delta$, or $\gamma \to \delta$ with δ between γ and β on $\boldsymbol{\mu}$. Such a path is said to be an *anterior path* from α to β. By (c) in Definition 2.2, the configurations $\to \gamma -$ and $\leftrightarrow \gamma -$ never occur in an ancestral graph, hence every anterior path takes the form:

$$\alpha - \cdots - \beta \to \cdots \to \gamma$$

where $\alpha = \beta$ and $\beta = \gamma$ are possible. We apply these definitions disjunctively to sets:

$$\begin{aligned} an(X) &= \{\alpha \mid \alpha \text{ is an ancestor of } \beta \text{ for some } \beta \in X\}; \\ ant(X) &= \{\alpha \mid \alpha \text{ is anterior to } \beta \text{ for some } \beta \in X\}. \end{aligned}$$

Note that by definition $X \subseteq \mathrm{an}(X) \subseteq \mathrm{ant}(X)$.

Note that every DAG is an ancestral graph, since clauses (b) and (c) are trivially satisfied.

In the next Lemma and elsewhere we will make use of the shorthand notation $x ?\!\to y$ to indicate that either $x \to y$ or $x \leftrightarrow y$. Similarly, $x ?\!- y$ indicates that either $x \leftarrow y$ or $x - y$.

Lemma 2.1 *In an ancestral graph $\mathcal{G}$ if a and c are adjacent, there is an edge $b \to c$, and $a ?\!\!\to b$ then either $c \leftarrow a \to b$, or $c \leftrightarrow a \leftrightarrow b$, or $c \leftarrow a \leftrightarrow b$.*

In particular, note that if the edge ends at a on the (a, b) and (a, c) edges differ then we have $c \leftarrow a \leftrightarrow b$. We make use of this property in §3.4 and §3.8.

Proof: If there is no arrowhead at c on the (a, c) edge then $c \to a$ (the edge cannot be undirected since, by hypothesis there is an arrowhead at c on the (b, c) edge). But in this case the graph is not ancestral since we have $a ?\!\!\to b \to c \to a$. Hence $a ?\!\!\to c$. The conclusion then follows from noting that the configuration $a \leftarrow b \leftarrow c \leftrightarrow a$ is not ancestral. □

2.3 The global Markov property for ancestral graphs

In an ancestral graph a non-endpoint vertex v on a path is said to be a *collider* if two arrowheads meet at v, i.e. $\to v \leftarrow$, $\leftrightarrow v \leftrightarrow$, $\leftrightarrow v \leftarrow$ or $\to v \leftrightarrow$; all other non-endpoint vertices on a path are *non-colliders*, i.e. $- v -$, $- v \to$, $\to v \to$, $\leftarrow v \to$, $\leftrightarrow v \to$. A subpath of length 3 is called a *triple*. In an ancestral graph a triple is either a collider or a non-collider; we refer to this as the *type* of the triple. These definitions of collider and non-collider are direct extensions of the corresponding definitions for DAGs. A path along which every non-endpoint is a collider is called a *collider path*.

Verma and Pearl (1991) introduced *d-separation*, a set of graphical conditions by which conditional independence relations could be read off a DAG. Richardson and Spirtes (2002) apply a natural extension of Pearl's d-separation criterion, called m-separation, to ancestral graphs:

Definition 2.2 *In an ancestral graph $\mathcal{G}$ a path $\boldsymbol{\pi}$ between α and β is said to be* m-connecting *given Z ($\alpha, \beta \notin Z$) if the following hold:*

(i) no non-collider on $\boldsymbol{\pi}$ is in Z; and

(ii) every collider on $\boldsymbol{\pi}$ is an ancestor of a vertex in Z.

Two vertices α and β are said to be m-separated *given Z in $\mathcal{G}$ if there is no path m-connecting α and β given Z in $\mathcal{G}$. Likewise, sets A and B are* m-separated *given Z in $\mathcal{G}$ if for every pair $\alpha \in A$, and $\beta \in B$, α and β are m-separated given Z.*

For example in the ancestral graph in Figure 2(ii), Azt and Ap are m-separated given $CD4$. Definition 2.2 is an extension of the original definition of d-separation for DAGs in that the notions of 'collider' and 'non-collider' now allow for bi-directed and undirected edges; the definition of ancestor is unchanged. Furthermore, d-separation is equivalent to m-separation for DAGs. The following result is useful:

Lemma 2.2 *If $\boldsymbol{\pi}$ is a path m-connecting α and β given Z, γ is on $\boldsymbol{\pi}$ ($\alpha \neq \gamma \neq \beta$), and there is an arrowhead at γ on the subpath $\boldsymbol{\pi}(\alpha, \gamma)$ then either $\gamma \in \mathrm{an}(Z)$, or $\boldsymbol{\pi}(\gamma, \beta)$ is a directed path from γ to β.*

Proof: Suppose the result is false. Let γ be the vertex closest to β satisfying the antecedent of the Lemma, but not the conclusion. If γ is a collider, then by definition of m-connection, $\gamma \in \mathrm{an}(Z)$ which is a contradiction. If γ is a non-collider then let δ be the vertex after γ on $\boldsymbol{\pi}(\gamma, \beta)$. Now, $\gamma \to \delta$ since γ is a non-collider. If $\delta \in \mathrm{an}(Z)$, or $\boldsymbol{\pi}(\delta, \beta)$ forms a directed path from δ to β, then clearly γ satisfies the conclusion of the Lemma. But if $\delta \notin \mathrm{an}(Z)$ and $\boldsymbol{\pi}(\delta, \beta)$ is not a directed path to β then δ satisfies the conditions on γ, but is closer to β, again a contradiction. □

2.4 Independence models

An *independence model* $\mathfrak{I}$ over vertex set V is a set of relations, $\langle X, Y \mid Z \rangle$, where X, Y, Z are disjoint subsets of V, and Z may be empty. The set V also indexes a collection of random variables. If $\langle X, Y \mid Z \rangle \in \mathfrak{I}$ then we say that X is conditionally independent of Y given Z, which we also write as $X \perp\!\!\!\perp Y \mid Z$, using the notation of Dawid (1980).

The independence model associated with an ancestral graph is defined via m-separation. More precisely, we say that a distribution P obeys the *global Markov property* with respect to graph $\mathcal{G}$ if for all disjoint sets X, Y, Z, (Z may be empty):

$$X \text{ m-separated from } Y \text{ given } Z \text{ in } \mathcal{G} \quad \Longrightarrow \quad X \perp\!\!\!\perp Y \mid Z \;\; [P].$$

We will denote the set of independence relations implied to hold by an ancestral graph $\mathcal{G}$ via m-separation by $\mathfrak{I}_m(\mathcal{G})$.

2.5 Marginalizing and conditioning independence models ($\mathfrak{I}[^S_L$)

An independence model $\mathfrak{I}$ *after marginalizing out a subset* L is simply the subset of relations which do not involve any vertices in L. More formally we define:

$$\mathfrak{I}[_L \;\equiv\; \Big\{ \langle X, Y \mid Z \rangle \;\Big|\; \langle X, Y \mid Z \rangle \in \mathfrak{I};\;\; (X \cup Y \cup Z) \cap L = \emptyset \Big\}.$$

If $\mathfrak{I}$ contains the independence relations present in a distribution P, then $\mathfrak{I}[_L$ contains the subset of independence relations remaining after marginalizing out the 'Latent' variables in L; see Richardson and Spirtes (2002) Theorem 7.1.

An independence model $\mathfrak{I}$ *after conditioning on a subset* S is the set of independence relations defined as follows:

$$\mathfrak{I}[^S \;\equiv\; \Big\{ \langle X, Y \mid Z \rangle \;\Big|\; \langle X, Y \mid Z \cup S \rangle \in \mathfrak{I};\;\; (X \cup Y \cup Z) \cap S = \emptyset \Big\}.$$

Thus if $\mathfrak{I}$ contains the independence relations present in a distribution P then $\mathfrak{I}[^S$ constitutes the subset of independencies holding among the remaining variables after conditioning on S. The letter S is used because Selection effects represent one context in which conditioning may occur. See Cox and Wermuth (1996), p.44, for further discussion of conditioning in this context.

Combining these definitions we obtain, for disjoint sets S, L:

$$\mathfrak{I}[^S_L \;\equiv\; \Big\{\langle X, Y \mid Z\rangle \;\Big|\; \langle X, Y \mid Z \cup S\rangle \in \mathfrak{I};\;\; (X \cup Y \cup Z) \cap (S \cup L) = \emptyset\Big\}.$$

Note that, by definition, $\mathfrak{I}[^\emptyset_\emptyset = \mathfrak{I}$.

2.6 Graphical transforms corresponding to marginalizing and conditioning

Given an ancestral graph $\mathcal{G}$ with vertex set V, for arbitrary disjoint sets S, L (both possibly empty) we now define a graphical transformation

$$\mathcal{G} \mapsto \mathcal{G}[^S_L$$

in such a way that the independence model corresponding to the transformed graph will be the independence model obtained by marginalizing and conditioning the independence model of the original graph.

Though we define this transformation for any ancestral graph $\mathcal{G}$, our primary motivation is the case in which $\mathcal{G}$ is an underlying (causal) DAG that is partially observed.

Definition 2.3 *For disjoint sets S, L,* graph $\mathcal{G}[^S_L$ *has vertex set $V \setminus (S \cup L)$, and edges specified as follows:*

If α, β, are s.t. $\forall Z$, with $Z \subseteq V \setminus (S \cup L \cup \{\alpha, \beta\})$,

$$\langle\{\alpha\}, \{\beta\} | Z \cup S\rangle \notin \mathfrak{I}_m(\mathcal{G}),$$

$$\textit{and} \left\{\begin{array}{c} \alpha \in \mathrm{ant}_{\mathcal{G}}(\{\beta\} \cup S);\;\; \beta \in \mathrm{ant}_{\mathcal{G}}(\{\alpha\} \cup S) \\ \alpha \notin \mathrm{ant}_{\mathcal{G}}(\{\beta\} \cup S);\;\; \beta \in \mathrm{ant}_{\mathcal{G}}(\{\alpha\} \cup S) \\ \alpha \in \mathrm{ant}_{\mathcal{G}}(\{\beta\} \cup S);\;\; \beta \notin \mathrm{ant}_{\mathcal{G}}(\{\alpha\} \cup S) \\ \alpha \notin \mathrm{ant}_{\mathcal{G}}(\{\beta\} \cup S);\;\; \beta \notin \mathrm{ant}_{\mathcal{G}}(\{\alpha\} \cup S) \end{array}\right\} \textit{then} \left\{\begin{array}{c} \alpha - \beta \\ \alpha \leftarrow \beta \\ \alpha \rightarrow \beta \\ \alpha \leftrightarrow \beta \end{array}\right\} \textit{in } \mathcal{G}[^S_L.$$

In words, $\mathcal{G}[^S_L$ is a graph containing the vertices that are not in S or L. Two vertices α, β are adjacent in $\mathcal{G}[^S_L$ if α and β are m-connected in $\mathcal{G}$ given any subset that contains all vertices in S and no vertices in L. If α and β are adjacent in $\mathcal{G}[^S_L$ then there is an arrowhead at α if and only if α is not anterior either to β or any vertex in S, and a tail otherwise.

Richardson and Spirtes (2002) showed that given an ancestral graph $\mathcal{G}$, the graph $\mathcal{H} = \mathcal{G}[^S_L$, satisfies

$$\mathfrak{I}_m(\mathcal{H}) = (\mathfrak{I}_m(\mathcal{G}))[^S_L.$$

In particular, if $\mathcal{G}$ is a DAG with observed variables O, latent variables L, and selection variables S, then the graph $\mathcal{G}[^S_L$ formed through the graphical transformation represents those conditional independence relations implied to hold among the observed variables, conditional on any selection variables.

Given a maximal ancestral graph $\mathcal{G}$ (see 3.2) with vertex set V it is always possible to construct a DAG $\mathcal{D}(\mathcal{G})$ with vertex set $V \cup S_{\mathcal{D}(\mathcal{G})} \cup L_{\mathcal{D}(\mathcal{G})}$, such that:

$$\mathcal{G} = \mathcal{D}(\mathcal{G})[^{S_{\mathcal{D}(\mathcal{G})}}_{L_{\mathcal{D}(\mathcal{G})}}.$$

i.e. the graphical transformation of marginalizing over L and conditioning on S, when applied to $\mathcal{D}(\mathcal{G})$ results in the original graph $\mathcal{G}$.

It also follows from this result that the class of independence models associated with ancestral graphs is the smallest class that contains the DAG independence models and is closed under marginalizing and conditioning.

3 Markov equivalence

We introduce the following:

Definition 3.1 *Two graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ with the same vertex set are said to be Markov equivalent if for all disjoint sets A, B, Z (Z may be empty), A*

and B are m-separated given Z in $\mathcal{G}_1$ if and only if A and B are m-separated given Z in $\mathcal{G}_2$, i.e. $\mathfrak{I}_m(\mathcal{G}_1) = \mathfrak{I}_m(\mathcal{G}_2)$.

The graphs in Figure 3 are Markov equivalent, as are $\mathcal{G}_1$ and $\mathcal{G}_2$ in Figure 4. As outlined in the introduction, whether or not two graphs are Markov equivalent is of great significance since the associated models contain the same set of distributions. Consequently they are indistinguishable in the asymptotic limit. At the same time, some members of the equivalence class may lead more directly to closed form MLEs (see Drton and Richardson (2004a)). In addition, for the purposes of model search it is more efficient to perform a search across equivalence classes rather than scoring individual models (Chickering (2002a), Chickering (2002b)).

3.1 Markov equivalence for DAGs

Frydenberg (1990) and Verma and Pearl (1991) gave simple graphical conditions for determining whether two DAGs are Markov equivalent. A triple of vertices $\langle a, b, c\rangle$ is said to be *unshielded* if a and c are not adjacent, and *shielded* otherwise. (A triple is defined in §2.3.)

Theorem 3.1 *Two DAGs are Markov equivalent if and only if they have the same vertex set, the same adjacencies and the same unshielded colliders.*

3.2 Maximal ancestral graphs

As described in the last section, for two DAGs to be Markov equivalent it is necessary that they have the same adjacencies. This is a direct consequence of the fact that DAGs satisfy a pairwise Markov property:

Proposition 3.2 *In a DAG $\mathcal{D}$, if α and β are not adjacent, and $\alpha \notin \mathrm{de}(\beta)$ then α is m-separated from β by $V \setminus (\mathrm{de}(\beta) \cup \{\alpha\})$.*

Note that by acyclicity for any pair α, β, either $\alpha \notin \mathrm{de}(\beta)$ or $\beta \notin \mathrm{de}(\alpha)$. Consequently in a DAG every missing edge implies a conditional independence between the non-adjacent vertices.

In general, no such property holds for ancestral graphs. For example, there is no set that m-separates γ and δ in the graph in Figure 3(a). This motivates the following definition:

Definition 3.2 *An ancestral graph $\mathcal{G}$ is said to be* maximal *if, for every pair of non-adjacent vertices (α, β) there exists a set Z, $(\alpha, \beta \notin Z)$ such that α and β are m-separated conditional on Z.*

These graphs are maximal in the sense that no additional edge may be added to the graph without changing the associated independence model. It has been shown in Theorem 5.1 of Richardson and Spirtes (2002) that if an ancestral graph $\mathcal{G}$ is not maximal, then there exists a unique maximal ancestral graph $\mathcal{G}^*$ which is a supergraph of $\mathcal{G}$, and is Markov equivalent to $\mathcal{G}$. The graph $\mathcal{G}^*$ may be formed by adding bi-directed edges to $\mathcal{G}$ whenever two non-adjacent vertices α and β in $\mathcal{G}$ are joined by an *inducing path*, defined as:

Definition 3.3 *An inducing path $\boldsymbol{\pi}$ between α and β in an ancestral graph $\mathcal{G}$ is a path on which every non-endpoint vertex is both a collider and an ancestor of at least one of the endpoints, α, β.*

Note that strictly speaking an inducing 'path' $\boldsymbol{\pi} = \langle \alpha, \nu_1, \ldots \nu_p, \beta \rangle$ is a collection of paths: the collider path $\boldsymbol{\pi}$, together with directed paths from each vertex ν_i to one of the endpoints. (The definition given here is called a 'primitive' inducing path in Richardson and Spirtes (2002); the concept was introduced by Verma and Pearl (1990).)

Figure 3(a) shows an example of a non-maximal ancestral graph. The path $\langle \gamma, \beta, \alpha, \delta \rangle$ forms an inducing path between γ and δ. By adding the

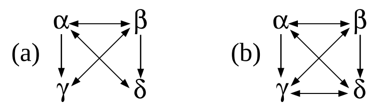

Figure 3: (a) The path $\langle\gamma, \beta, \alpha, \delta\rangle$ is an example of an inducing path in an ancestral graph. (b) A maximal ancestral graph Markov equivalent to (a).

bi-directed edge $\gamma \leftrightarrow \delta$ the graph is made maximal without changing the associated independence model (which is empty), as shown in Figure 3(b).

Since for every non-maximal ancestral graph there is a unique maximal ancestral graph of which it is a subgraph, the problem of characterizing Markov equivalence for ancestral graphs naturally reduces to that of characterizing equivalence in the case where both graphs are maximal. In the remainder of this paper we will restrict attention to maximal ancestral graphs.

Proposition 3.3 *If $\mathcal{G}_1$ and $\mathcal{G}_2$ are Markov equivalent maximal ancestral graphs then $\mathcal{G}_1$ and $\mathcal{G}_2$ have the same adjacencies and unshielded colliders.*

Proof: Since $\mathcal{G}_1$ is maximal for each pair of non-adjacent vertices (x, y) in $\mathcal{G}_1$ there is some set Z such that x and y are m-separated given Z in $\mathcal{G}_1$. If x and y are adjacent in $\mathcal{G}_2$ then they are not m-separated by Z, contradicting Markov equivalence of $\mathcal{G}_1$ and $\mathcal{G}_2$. So adjacencies in $\mathcal{G}_1$ are a subset of those in $\mathcal{G}_2$. By a symmetric argument, the adjacencies in $\mathcal{G}_2$ are a subset of those in $\mathcal{G}_1$.

Suppose for a contradiction that $\langle\alpha, \beta, \gamma\rangle$ is an unshielded collider in $\mathcal{G}_1$, but not in $\mathcal{G}_2$. Since $\mathcal{G}_1$ is maximal, for some set Z, α and γ are m-separated by Z, and $\beta \notin Z$. If $\langle\alpha, \beta, \gamma\rangle$ is a non-collider in $\mathcal{G}_2$ then α and γ are m-connected given Z, which is a contradiction. Hence every unshielded collider in $\mathcal{G}_1$ is present in $\mathcal{G}_2$. The conclusion follows by symmetry. $\square$

An important consequence of this proposition is that if $\mathcal{G}_1$ and $\mathcal{G}_2$ are maximal and Markov equivalent, then a sequence of vertices forming a path in $\mathcal{G}_1$, also forms a path in $\mathcal{G}_2$ and vice-versa, though the edge-types on these paths may differ. Consequently we will often refer to the path $\boldsymbol{\pi}^*$ in $\mathcal{G}_2$ *corresponding to* a given path $\boldsymbol{\pi}$ in $\mathcal{G}_1$.

3.3 Discriminating paths in ancestral graphs

A key difference between DAGs and ancestral graphs is that the having the same adjacencies and unshielded colliders, though necessary, are not sufficient for Markov equivalence of ancestral graphs.

Consider the graphs shown in Figure 4: $\mathcal{G}_1$ and $\mathcal{G}_3$ contain the same adjacencies and the same unshielded colliders, but these two graphs are not Markov equivalent to each other: In $\mathcal{G}_1$, x is m-separated from y given q; but according to $\mathcal{G}_3$, x is m-connected to y given q. In fact in any graph Markov equivalence to $\mathcal{G}_1$, $\langle q, \beta, y\rangle$ forms a shielded collider. (There is only one such graph, $\mathcal{G}_2$, so $\{\mathcal{G}_1, \mathcal{G}_2\}$ forms a Markov equivalence class.) However, in general, it is clearly not necessary that two graphs have all of the same *shielded* colliders in order for them to be Markov equivalent.

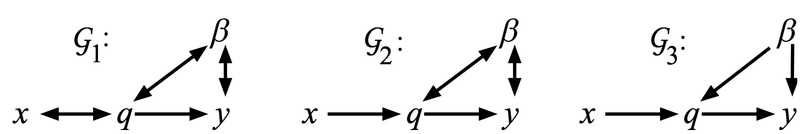

Figure 4: $\mathcal{G}_1$, $\mathcal{G}_2$, $\mathcal{G}_3$ have the same adjacencies and the same unshielded colliders, but $\mathcal{G}_1$ and $\mathcal{G}_3$ are not Markov equivalent. $\boldsymbol{\pi} = \langle x, q, \beta, y\rangle$ forms a discriminating path for β in every graphs.

Discriminating paths are special paths that, if present in two Markov equivalent graphs, imply that certain shielded colliders (or non-colliders) will be present in both graphs:

Definition 3.4 *A path* $\boldsymbol{\pi} = \langle x, q_1, q_2, \ldots, q_p, \beta, y\rangle$ *is a* discriminating path for $\langle q_p, \beta, y\rangle$ *in an ancestral graph* $\mathcal{G}$ *if and only if:*

(i) x *is not adjacent to* y*, and*

(ii) For every vertex $q_i, 1 \leq i \leq p$ *on* $\boldsymbol{\pi}$*, (i.e. excluding* x, y*, and* β*),* q_i *is a collider on* $\boldsymbol{\pi}$ *and* q_i *is a parent of* y *in* $\mathcal{G}$*.*

The paths $\langle x, q, \beta, y\rangle$ in $\mathcal{G}_1$, $\mathcal{G}_2$ and $\mathcal{G}_3$ from Figure 4 are examples of discriminating paths for β.

Like inducing paths, a discriminating 'path' $\boldsymbol{\pi} = \langle x, q_1, \ldots, q_p, \beta, y\rangle$ is in fact a collection of paths:

$$x \,?\!\rightarrow q_1 \leftrightarrow \cdots \leftrightarrow q_p \leftarrow\!?\, \beta ?\!\rightarrow y$$

$$x \,?\!\rightarrow q_1 \leftrightarrow \cdots \leftrightarrow q_j \rightarrow y \qquad\qquad (1 \leq j \leq p)$$

with the requirement that the endpoints x and y are not adjacent. We will often refer to some path $\boldsymbol{\pi}$ between x and y as *a discriminating path for* β, thereby implicitly specifying the triple for which the path is discriminating to be the triple containing y and β and the vertex prior to β on the path; by convention we order the endpoints of the discriminating path so that it is the second endpoint (in this case, y) which is in the discriminated triple.

Consider a discriminating path $\boldsymbol{\pi} = \langle x, q_1, \ldots, q_p, \beta, y\rangle$ in an ancestral graph $\mathcal{G}$. If a given set Z does not contain all vertices $q_i, 1 \leq i \leq p$, then for some j, $q_j \notin Z$ and $q_k \in Z$ for all $k < j$, thus the path $\langle x, q_1, \ldots, q_j, y\rangle$ m-connects x and y given Z (because $q_1, \ldots, q_{j-1}$ are colliders and q_j is a non-collider). See Figure 6. Hence, if Z m-separates x and y then $\{q_1, \ldots, q_p\} \subseteq Z$. Consequently, if β is a collider on the path $\boldsymbol{\pi}$ in the graph $\mathcal{G}$ and Z m-separates x and y, then $\beta \notin Z$: otherwise the path $\boldsymbol{\pi}$ would m-connect x

and y, since every non-endpoint vertex on $\boldsymbol{\pi}$ would be a collider and in Z. Conversely, if β is a non-collider on the path $\boldsymbol{\pi}$ then β is a member of any set Z that m-separates x and y.

Thus, whenever $\langle x, q_1, \ldots, q_p, \beta, y \rangle$ forms a discriminating path in $\mathcal{G}$, then β is a collider (non-collider) if and only if every set Z m-separating x and y is such that $\beta \notin Z$ ($\beta \in Z$). It follows that in any graph $\mathcal{G}^*$ Markov equivalent to $\mathcal{G}$ in which the path $\boldsymbol{\pi}^*$ corresponding to $\boldsymbol{\pi}$ forms a discriminating path for β, β is a collider on $\boldsymbol{\pi}^*$ in $\mathcal{G}^*$ if and only if β is a collider on $\boldsymbol{\pi}$ in $\mathcal{G}$. Thus we have proved the following:

Lemma 3.4 *Let $\boldsymbol{\pi} = \langle x, q_1, q_2, \ldots, q_p, \beta, y \rangle$ be a discriminating path for β in the maximal ancestral graph $\mathcal{G}$. If $\mathcal{G}^*$ is a maximal ancestral graph Markov equivalent to $\mathcal{G}$, and the corresponding path $\boldsymbol{\pi}^*$ forms a discriminating path for β in $\mathcal{G}^*$, then β is a collider on $\boldsymbol{\pi}$ in $\mathcal{G}$ if and only if β is a collider on $\boldsymbol{\pi}^*$ in $\mathcal{G}^*$.*

Thus even though q_p and y are adjacent, $\langle q_p, \beta, y \rangle$ is 'discriminated' by the path $\boldsymbol{\pi}$ to be of the same type (collider or non-collider) on the corresponding path in any graph $\mathcal{G}^*$ Markov equivalent to $\mathcal{G}$ in which *the corresponding path $\boldsymbol{\pi}^*$ also forms a discriminating path.*

The following Lemma shows that a necessary condition for $\boldsymbol{\pi}^*$ to form a discriminating path in $\mathcal{G}^*$ is that the colliders $\langle x, q_1, q_2 \rangle, \ldots, \langle q_{p-1}, q_p, \beta \rangle$ on $\boldsymbol{\pi}$ are also colliders on $\boldsymbol{\pi}^*$.

Lemma 3.5 *If $\boldsymbol{\pi} = \langle x, q_1, \ldots, q_p, \beta, y \rangle$ is a discriminating path in $\mathcal{G}$ then in any Markov equivalent graph $\mathcal{G}^*$ in which the q_i are colliders on the corresponding path $\boldsymbol{\pi}^*$, the edge between q_i and y is $q_i \to y$, ($1 \leq i \leq p$).*

Proof: The proof proceeds by induction on i. First consider the (q_1, y) edge in $\mathcal{G}^*$. If there is an arrowhead at q_1, then $\langle x, q_1, y \rangle$ forms an unshielded

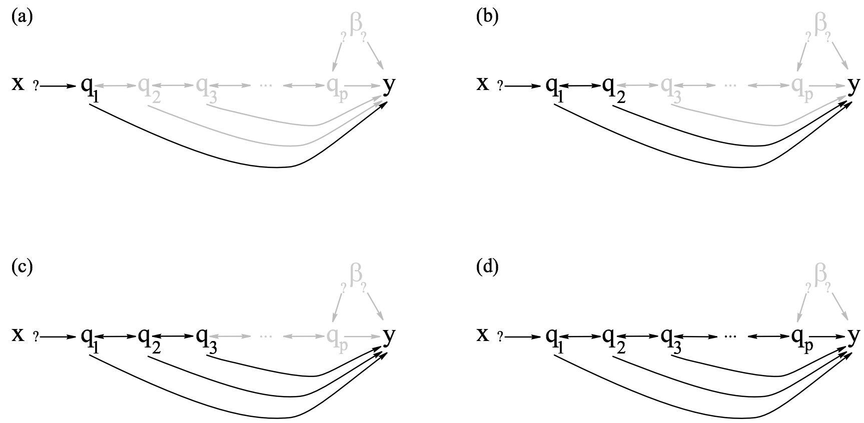

Figure 5: The sequence of discriminating paths for the non-colliders $\langle x, q_1, y\rangle$, $\langle q_{j-1}, q_j, q_{j+1}\rangle$ $(1 < j < p)$. See Lemma 3.5 and Lemma 3.7

collider in $\mathcal{G}^*$ but an unshielded non-collider in $\mathcal{G}$. But then by Lemma 3.3, $\mathcal{G}$ and $\mathcal{G}^*$ are not Markov equivalent, which is a contradiction. Since $x \,?\!\rightarrow q_1$, $q_1 - y$ is ruled out, hence $q_1 \rightarrow y$ in $\mathcal{G}^*$.

Suppose that $q_j \rightarrow y$ for $1 \leq j < i$ in $\mathcal{G}^*$. Then the path $\langle x, q_1, \ldots, q_i, y\rangle$ forms a discriminating path for q_i in both $\mathcal{G}^*$ and $\mathcal{G}$. If $q_i \leftarrow\!?\, y$ in $\mathcal{G}^*$ then $\langle q_{i-1}, q_i, y\rangle$ forms a collider in $\mathcal{G}^*$ but a non-collider in $\mathcal{G}$. But then by Lemma 3.4 $\mathcal{G}$ and $\mathcal{G}^*$ are not Markov equivalent, which is a contradiction. Since $q_{i-1} \leftrightarrow q_i$, $q_i - y$ is ruled out, so $q_i \rightarrow y$, as required. □

Though discriminating paths can exist in DAGs, they are not important for determining Markov equivalence because such paths always discriminate non-colliders (see $\mathcal{G}_3$ in Figure 4): if $\langle x, q, \beta\rangle$ forms a collider, then since there are no bi-directed edges in a DAG, it follows that β and x are parents of q.

One might hope that if $\mathcal{G}_1$ and $\mathcal{G}_2$ are Markov equivalent, then they would have the same discriminating paths. Unfortunately, this is not the

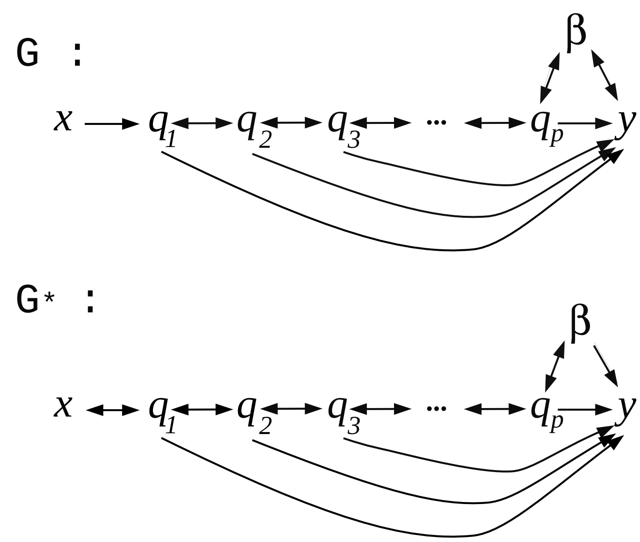

Figure 6: $\boldsymbol{\pi} = \langle x, q_1, \ldots, q_p, \beta, y\rangle$ forms a discriminating path for β in $\mathcal{G}$ and $\mathcal{G}^*$. β is a collider in $\mathcal{G}$ but a non-collider in $\mathcal{G}^*$ so the two graphs are not Markov equivalent. See Lemma 3.4 for further details.

case: it is possible for a path $\boldsymbol{\pi}$ to be discriminating in $\mathcal{G}$, and yet the corresponding path $\boldsymbol{\pi}^*$ not be discriminating in $\mathcal{G}^*$ even though $\mathcal{G}$ and $\mathcal{G}^*$ are Markov equivalent. Hence the antecedent in Lemma 3.4 will not hold for all pairs of Markov equivalent graphs. Thus, the fact that a non-collider is discriminated by a path in $\mathcal{G}$, does not mean that it will be present in every graph Markov equivalent to $\mathcal{G}$.

Consider the example given by the two graphs in Figure 7(i). Note that q is a collider on the path $\langle x, q, \beta, y\rangle$ in $\mathcal{G}_1$, but not in $\mathcal{G}_2$; $\langle x, q, \beta, y\rangle$ forms a discriminating path in $\mathcal{G}_1$, but not in $\mathcal{G}_2$, though they are Markov equivalent. Hence, although $\langle q, \beta, y\rangle$ is a non-collider in any graph Markov equivalent to $\mathcal{G}_1$ in which $\langle x, q, \beta, y\rangle$ forms a discriminating path, $\langle q, \beta, y\rangle$ need not be a non-collider in graphs such as $\mathcal{G}_2$, where the corresponding path is not discriminating for β.

We believe that a similar situation cannot occur for colliders. In particular, we conjecture that if a collider is discriminated by some path in $\mathcal{G}$, then this collider will be present in every graph $\mathcal{G}^*$ Markov equivalent to $\mathcal{G}$, even though there may not be a discriminating path for this collider in $\mathcal{G}^*$.

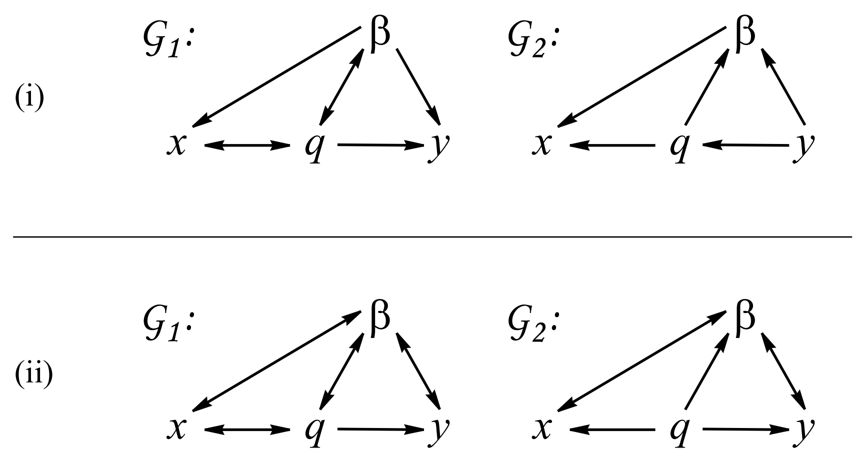

Figure 7: Two examples of maximal ancestral graphs that are Markov equivalent and $\langle x, q, \beta, y \rangle$ forms a discriminating path in $\mathcal{G}_1$, but not in $\mathcal{G}_2$. See text for further explanation.

The results in this section present a dilemma: it is clear that discriminating paths, when present in both graphs, lead directly to necessary conditions for Markov equivalence. However, a discriminating path for a given triple may not be present in all graphs within a Markov equivalence class. We avoid this problem by identifying, via a recursive definition, a sub-class of discriminating paths (those 'with order') which are always present, and show that these paths provide sufficient conditions for determining Markov equivalence (in conjunction with the earlier requirements that the graphs have the same adjacencies and unshielded colliders).

We now define a hierarchy of triples as follows:

Definition 3.5 *Let* $\mathfrak{O}_i$ *(*$i \geq 0$*) be the set of* triples of order i*, defined recursively as follows:*

Order 0: *A triple* $\langle \alpha, \beta, \gamma \rangle \in \mathfrak{O}_0$ *if* α *and* γ *are not adjacent in* $\mathcal{G}$*.*

Order $i+1$: *A triple* $\langle \alpha, \beta, \gamma \rangle \in \mathfrak{O}_{i+1}$ *if*

(1) $\langle \alpha, \beta, \gamma \rangle \notin \mathfrak{O}_j$*, for some* $j < i+1$ *and*

(2) there exists a discriminating path $\boldsymbol{\pi} = \langle x, q_1, \ldots, q_p = \alpha, \beta, \gamma \rangle$ *for* β *in* $\mathcal{G}$*, and each of the colliders on the path:*

$$\langle x, q_1, q_2 \rangle, \ldots \langle q_{p-1}, q_p, \beta \rangle \in \bigcup_{j \leq i} \mathfrak{O}_j .$$

If $\langle \alpha, \beta, \gamma \rangle \in \mathfrak{O}_i$ *then the triple is said to have order* i*. A discriminating path is said to have order* i *if every triple on the path has order at most* i*, and at least one triple has order* i*. If a triple has order* i *for some* i*, then we will say that the triple* has order*, likewise for discriminating paths.*

For example, in each graph in Figure 4, the triple $\langle x, q, \beta \rangle$ has order 0, while $\langle q, \beta, y \rangle$ has order 1. It is important to note that not every triple in a graph will have an order. For example, the triple $\langle q, \beta, y \rangle$ in Figure 7(i) does not have order. Further, it is possible for a triple without order to be of the same type (collider or non-collider) in every graph in the Markov equivalence class.

We now show that a necessary condition for two graphs to be Markov equivalent is that they have the same colliders with order.

Proposition 3.6 *If* $\langle \alpha, \beta, \gamma \rangle$ *has order in* $\mathcal{G}$ *then* $\langle \alpha, \beta, \gamma \rangle$ *has order in any graph* $\mathcal{G}^*$ *Markov equivalent to* $\mathcal{G}$*, and further* $\langle \alpha, \beta, \gamma \rangle$ *is a collider in* $\mathcal{G}$ *if and only if it is a collider in* $\mathcal{G}^*$*.*

Proof: The proof is by induction on i, the order of the triple $\langle \alpha, \beta, \gamma \rangle$. For $i = 0$, the result follows from Proposition 3.3. For $i > 0$, it follows by the definition of order that there exists some discriminating path $\boldsymbol{\pi} = \langle q_0, q_1, \ldots, q_p = \alpha, \beta, \gamma \rangle$. By definition, with the exception of $\langle \alpha, \beta, \gamma \rangle$, every other triple on $\boldsymbol{\pi}$ is a collider, and has order less than i. Consequently, by the induction hypothesis, in $\mathcal{G}^*$ these triples have order and also form colliders. By Lemma 3.5 since the q_i's $(i > 0)$ are colliders on the corresponding path $\boldsymbol{\pi}^*$ in $\mathcal{G}^*$, $q_i \rightarrow y$ $(1 \leq i \leq p)$ in $\mathcal{G}^*$. The result then follows from Lemma 3.4. □

We next establish that having the same non-colliders with order is implied by having the same adjacencies and colliders with order.

Proposition 3.7 *If $\mathcal{G}_1$ and $\mathcal{G}_2$ are maximal ancestral graphs with the same adjacencies and the same colliders with order then $\mathcal{G}_1$ and $\mathcal{G}_2$ also have the same non-colliders with order.*

Though Proposition 3.7 appears similar to Lemma 3.5, the antecedent in the latter assumes the two graphs are Markov equivalent, while in the former it does not.

Proof: Suppose $\langle q_p, \beta, y \rangle$ is a non-collider with order in $\mathcal{G}_1$. Then there is a discriminating path $\boldsymbol{\mu} = \langle q_0, q_1, \ldots, q_p, \beta, y \rangle$ for β, in $\mathcal{G}_1$ and each collider q_i $(1 \leq i \leq p)$ on $\boldsymbol{\mu}$ has order. By hypothesis each collider q_i is also a collider on the corresponding path $\boldsymbol{\mu}^*$ in $\mathcal{G}_2$ and has order.

We claim that $\boldsymbol{\mu}^*$ also forms a discriminating path in $\mathcal{G}_2$. It is sufficient to show by induction that $q_j \rightarrow y$ $(1 \leq j \leq p)$ in $\mathcal{G}_2$ (see Figure 5). If $\langle q_0, q_1, y \rangle$ forms a collider in $\mathcal{G}_2$ then it has order because q_0 and y are not adjacent. This is a contradiction because $\langle q_0, q_1, y \rangle$ forms a non-collider with order in $\mathcal{G}_1$. Hence $\langle q_0, q_1, y \rangle$ is a non-collider in $\mathcal{G}_2$, and further, $q_1 \rightarrow y$, because $q_0 \,?\!\!\rightarrow q_1$. Now suppose that $q_i \rightarrow y$ $(1 \leq i < j)$ in $\mathcal{G}_2$. Now $\langle q_0, q_1, \ldots, q_j, y \rangle$

forms a discriminating path with order for $\langle q_{j-1}, q_j, y\rangle$. Consequently if $\langle q_{j-1}, q_j, y\rangle$ formed a collider in $\mathcal{G}_2$ then $\mathcal{G}_1$ and $\mathcal{G}_2$ would have different colliders with order, which is a contradiction.

Hence $\boldsymbol{\mu}$ forms a discriminating path in $\mathcal{G}_2$ which has order because the colliders q_j $(1 \leq j \leq p)$ have order in $\mathcal{G}_1$, hence also in $\mathcal{G}_2$. Finally, if $\langle q_p, \beta, y\rangle$ were a collider in $\mathcal{G}_2$ then again the graphs would not have the same colliders with order. □

3.4 Discriminating sub-paths of a path

It follows from Propositions 3.6 and 3.7 that having the same colliders with order and non-colliders with order is a necessary condition for Markov equivalence. As a step towards showing that this condition (together with the same adjacencies) is sufficient, we will show that every triple on a 'minimal' m-connecting path has order (see §3.5). We first consider the relationships between discriminating paths which are sub-paths of a given path.

Given two non-endpoint vertices a, b on a path $\boldsymbol{\pi}$, we define a relation $a \prec_\pi b$ if there is a sub-path $\boldsymbol{\pi}(x, y)$ which forms a discriminating path for b, $a \neq b$, and a is a collider on $\boldsymbol{\pi}(x, y)$ (note that by our convention, b is adjacent to y; see page 18). In this section we prove that, as the symbol $\prec_\pi$ suggests, this relation is acyclic, so that for $p > 1$,

$$b_1 \prec_\pi \cdots \prec_\pi b_p \qquad \text{implies} \qquad b_p \not\prec_\pi b_1. \tag{1}$$

Note however that the relation $\prec_\pi$ is not transitive in general. The acyclic property (1) is central to establishing that every triple on a 'minimal' m-connecting path has order; see Lemma 3.14.

Lemma 3.8 *In the maximal ancestral graph $\mathcal{G}$, if $\boldsymbol{\pi}$ is a path then there is no pair of distinct vertices $\{b, j\}$ such that $b \prec_\pi j$ and $j \prec_\pi b$. Consequently, the relation is anti-symmetric: $b \prec_\pi j$ implies $j \not\prec_\pi b$.*

This Lemma establishes (1) in the case where $p = 2$, hence if q is a collider on a discriminating path for β then β is not a collider on a discriminating path for q.

Proof: For a contradiction suppose that $b \prec_\pi j$ and $j \prec_\pi b$. In this case there exists a subpath $\boldsymbol{\pi}(i,k)$ of $\boldsymbol{\pi}$ forming a discriminating path for j, and b is a collider on this path, and there exists a subpath $\boldsymbol{\pi}(a,c)$ forming a discriminating path for b and j is a collider on this path. By definition, $b \neq j$. Furthermore, $c \neq k$: otherwise either b would not lie on $\boldsymbol{\pi}(i,k)$ and j would not lie on $\boldsymbol{\pi}(a,c)$, or $b = j$, both of which are contradictions (see Figure 8).

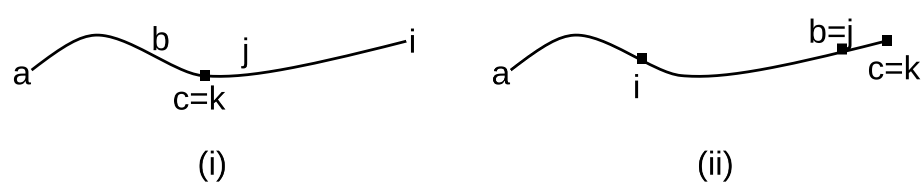

Figure 8: Diagram 1 for proof of Lemma 3.8.

Now, $j \neq a$ and $j \neq c$ because j is a collider on $\boldsymbol{\pi}(a,c)$, and similarly, $i \neq b \neq k$. Since j is a collider on $\boldsymbol{\pi}(a,c)$, by definition of a discriminating path, j is a parent of c. Consequently, $c \neq i$ because otherwise $\boldsymbol{\pi}(i,k)$ forms an inducing path and by maximality, i is adjacent to k which is a contradiction (see Figure 9).

Similarly, b is a parent of k, and $k \neq a$. Further, $a \neq i$ because then either j is not on $\boldsymbol{\pi}(a,c)$ or b is not on $\boldsymbol{\pi}(i,k)$ or both. Therefore, (i) k lies on $\boldsymbol{\pi}(a,c)$ and (ii) c lies on $\boldsymbol{\pi}(i,k)$, which is a contradiction: by (i) k is a parent of c, but by (ii) c is a parent of k (see Figure 10). $\square$

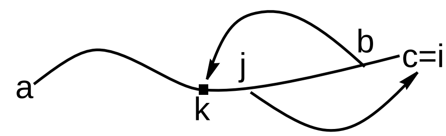

Figure 9: Diagram 2 for proof of Lemma 3.8.

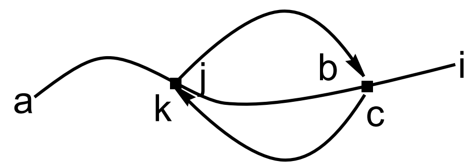

Figure 10: Diagram 3 for proof of Lemma 3.8.

Corollary 3.9 *On a path* $\boldsymbol{\pi}$ *in a maximal ancestral graph* $\mathcal{G}$ *there is no set of distinct vertices* $\{x, y, z\}$ *on* $\boldsymbol{\pi}$ *such that* $x \prec_\pi y \prec_\pi z$ *and* z *is on the subpath* $\boldsymbol{\pi}(x, y)$*.*

Proof: Suppose for a contradiction that such distinct vertices $\{x, y, z\}$ exist in $\mathcal{G}$ on $\boldsymbol{\pi}$ as shown in Figure 11. Since z is between x and y, and $x \prec_\pi y$, z is a collider on the discriminating path for y, hence $z \prec_\pi y$. But $y \prec_\pi z$ by hypothesis, which contradicts Lemma 3.8. $\square$

Corollary 3.10 *On a path* $\boldsymbol{\pi}$ *in a maximal ancestral graph* $\mathcal{G}$*, there is no quadruple of distinct vertices* $\langle q, b^*, b, q^* \rangle$*, in that order, on* $\boldsymbol{\pi}$ *such that* $q \prec_\pi b$ *and* $q^* \prec_\pi b^*$*.*

Proof: For a contradiction, suppose that such distinct vertices $\langle q, b^*, b, q^* \rangle$ exist on $\boldsymbol{\pi}$ in $\mathcal{G}$. Let $\boldsymbol{\pi}(x, y)$ be the discriminating path for b and $\boldsymbol{\pi}(x^*, y^*)$

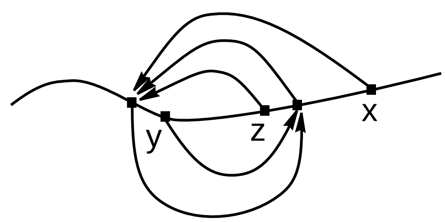

Figure 11: Diagram for proof of Corollary 3.9. See text for further explanation.

be the discriminating path for b^* (see Figure 12). Note that possibly $x = b^*$ or $x^* = b$; likewise $x = y^*$ or $x^* = y$ are also possible.

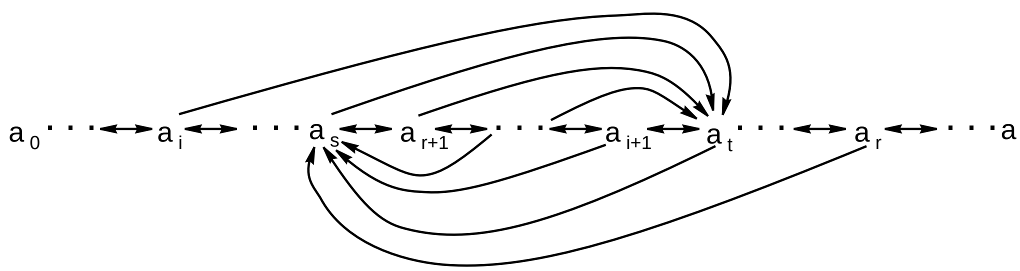

Figure 12: Diagram for proof of Corollary 3.10.

Now, b is a collider on $\boldsymbol{\pi}(x^*, y^*)$ because b is between q^* and b^*, so $b \prec_\pi b^*$. Likewise, b^* is a collider on $\boldsymbol{\pi}(x, y)$, so $b^* \prec_\pi b$. But this contradicts Lemma 3.8. □

Corollary 3.11 *On a path* $\boldsymbol{\pi}$ *in a maximal ancestral graph* $\mathcal{G}$ *there is no sequence of distinct vertices* $\langle b_1, b_2, \ldots, b_n \rangle, n > 1$ *such that* $b_i \prec_\pi b_{i+1}$, $1 \leq i < n$, *and* $b_n \prec_\pi b_1$.

Proof: Suppose for a contradiction that such a sequence of vertices $\langle b_1, b_2, \ldots, b_n \rangle$ exists. (Note that the order of the indices need not correspond to relative positions on $\boldsymbol{\pi}$.) By Lemma 3.8 $n \neq 2$ so $n > 2$. Let $\boldsymbol{\pi}$ have endpoints x and y. Without loss of generality, suppose b_1 is between b_n and y on $\boldsymbol{\pi}$. Let j be the highest index such that b_j is between b_1 and y, if such a vertex exists; otherwise let $j = 1$. Since, by hypothesis, $b_{n-1} \prec_\pi b_n \prec_\pi b_1$, by Corollary 3.9 b_1 is not on $\boldsymbol{\pi}(b_n, b_{n-1})$. Hence $j < n-1$. (Note that $b_{n-1} \neq b_1$ since $n > 2$.)

We will now show that b_{j+1} is between x and b_n.

> If $j = 1$, then by the definition of j, b_2 is between x and b_1. Further, since $b_n \prec_\pi b_1 \prec_\pi b_2$, by Corollary 3.9, b_2 is not on $\boldsymbol{\pi}(b_n, b_1)$ so b_2 is between x and b_n.
>
> If $1 < j < n-1$ then by the definition of j, b_{j+1} does not lie between b_1 and y. Since $j \neq n-1$, $b_{j+1} \neq b_n$. As $b_j \prec_\pi b_{j+1}$ and $b_n \prec_\pi b_1$, by Corollary 3.10, b_{j+1} does not lie on $\boldsymbol{\pi}(b_n, b_1)$. Hence b_{j+1} is between x and b_n as required.

We have shown that $j < n-1$. If $j = n-2$ then b_n is between $b_{j+1} = b_{n-1}$ and $b_j = b_{n-2}$ on $\boldsymbol{\pi}$. However, this contradicts Corollary 3.9 since $b_{n-2} \prec_\pi b_{n-1} \prec_\pi b_n$. Hence $j < n-2$.

We will now show that for each m $(j+2 \leq m < n)$, the vertex b_m is between x and b_{m-1} on $\boldsymbol{\pi}$, which will lead to a contradiction. Note that by the definition of b_j, no vertex b_m $(j < m < n)$ is between b_j and y, hence it is sufficient to show that b_m does not lie on $\boldsymbol{\pi}(b_{m-1}, b_j)$.

Base Case: $(m = j+2 < n)$ Since $b_j \prec_\pi b_{j+1} \prec_\pi b_{j+2}$, it follows by Corollary 3.9 that b_{j+2} does not lie on $\boldsymbol{\pi}(b_{j+1}, b_j)$ and b_{j+1} is between x and b_n. Hence $b_m = b_{j+2}$ is between x and $b_{m-1} = b_{j+1}$ as required.

Inductive Case: $(m = j+k)$ Assume for each s $(1 < s < k)$ b_{j+s} is between x and b_{j+s-1}. Since $b_{m-2} \prec_\pi b_{m-1} \prec_\pi b_m$, by Corollary 3.9 b_m does not lie on $\boldsymbol{\pi}(b_{m-1}, b_{m-2})$. Further, b_m does not lie on $\boldsymbol{\pi}(b_{j+s}, b_{j+s-1})$

for any s $(0 < s < m - j - 1)$: by the induction hypothesis, we would have $b_{m-1} \prec_\pi b_m$ and $b_{j+s} \prec_\pi b_{j+s-1}$, but the vertices occuring in the order b_{m-1}, b_{j+s}, b_m, b_{j+s-1} on $\boldsymbol{\pi}$ thereby violating Corollary 3.10. Hence b_m does not lie on $\boldsymbol{\pi}(b_{m-1}, b_j)$ as required.

We have shown that $j < n - 2$ and the vertices b_k, $j \leq k \leq n$ occur on $\boldsymbol{\pi}$ in the following order: $b_{n-1}, b_{n-2}, \ldots, b_{j+1}, b_n, b_j$. Since $b_{n-1} \prec_\pi b_n$ and $b_j \prec_\pi b_{j+1}$ this contradicts Corollary 3.10. □

We have now established (1), the acyclicity of $\prec_\pi$.

3.5 Minimal m-connecting paths

As we have seen, a triple may be a collider on an m-connecting path $\boldsymbol{\pi}$ in a graph $\mathcal{G}$, and yet be a non-collider on the corresponding path $\boldsymbol{\pi}^*$ in some graph $\mathcal{G}^*$, Markov equivalent to $\mathcal{G}$. In this section we prove that every triple on a 'minimal' m-connecting path in $\mathcal{G}$ has order, and thus, by Proposition 3.6, is of the same type in any graph Markov equivalent to $\mathcal{G}$.

Definition 3.6 *A path $\boldsymbol{\mu}$, m-connecting x and y given Z, will be said to be* minimal *if no proper subsequence of the vertices on $\boldsymbol{\mu}$ form an m-connecting path between x and y given Z.*

It is simple to see that if there is some path m-connecting x and y given Z then there is a minimal path which m-connects x and y given Z. If $\boldsymbol{\mu} = \langle \nu_1, \ldots, \nu_p \rangle$ is a path, then we will refer to any pair of vertices (ν_i, ν_j) for which $|i - j| > 1$ as *non-consecutive vertices on $\boldsymbol{\mu}$*. As the next Lemma shows, on a minimal m-connecting path only certain non-consecutive vertices may be adjacent.

Lemma 3.12 *Let $\boldsymbol{\pi}$ be a minimal m-connecting path between α and β given Z in the maximal ancestral graph $\mathcal{G}$; let δ and γ be two non-consecutive vertices on $\boldsymbol{\pi}$ that are adjacent in $\mathcal{G}$. Then at least one of δ and γ is a* *collider on $\boldsymbol{\pi}$, this collider is in Z and the (δ, γ) edge is directed out of the said collider.*

Note that the existence of non-consecutive vertices implies that there are at least four vertices on the path. Lemma 3.12 is illustrated in Figure 13.

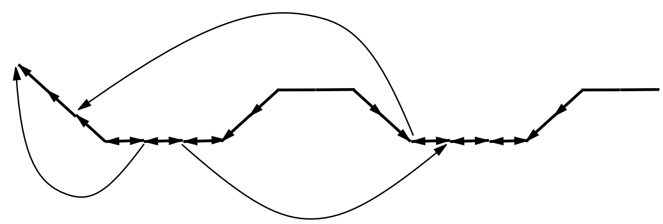

Figure 13: Example of a minimal m-connecting path given Z in an ancestral graph.

Proof: There are three cases to consider:

(i) *At least one of δ or γ is an endpoint.*
Suppose without loss that δ is an endpoint, and $\delta = \beta$ (the other cases are symmetric). By minimality of $\boldsymbol{\pi}$, γ is not an endpoint (since then the path consisting of the (γ, δ) edge would be m-connecting given Z). Let $\boldsymbol{\eta}$ be the path formed by concatenating $\boldsymbol{\pi}(\alpha, \gamma)$ and the (γ, δ) edge. We will now show that if δ is an endpoint then γ cannot be a non-collider on $\boldsymbol{\pi}$, and so is a collider on $\boldsymbol{\pi}$.
Suppose for a contradiction that γ is a non-collider on $\boldsymbol{\pi}$. Then $\gamma \notin Z$ since $\boldsymbol{\pi}$ is m-connecting given Z. If there is a tail at γ on the path $\boldsymbol{\pi}(\alpha, \gamma)$, then $\boldsymbol{\eta}$ is m-connecting given Z since γ is also a non-collider on $\boldsymbol{\eta}$. If there is an arrowhead at γ on $\boldsymbol{\pi}(\alpha, \gamma)$, then by Lemma 2.2, either $\gamma \in \text{an}(Z)$, or $\boldsymbol{\pi}(\gamma, \delta)$ forms a directed path. If $\gamma \in \text{an}(Z)$, then again $\boldsymbol{\eta}$ is m-connecting given Z, regardless of whether γ is a collider

or a non-collider on $\boldsymbol{\eta}$. If $\boldsymbol{\pi}(\gamma, \delta)$ forms a directed path, then $\gamma \to \delta$, since $\mathcal{G}$ is ancestral. Again $\boldsymbol{\eta}$ m-connects given Z. But the vertices on $\boldsymbol{\eta}$ form a proper subsequence of the vertices on $\boldsymbol{\pi}$, contradicting minimality. Hence γ is a collider on $\boldsymbol{\pi}$. Further note that there is no arrowhead at γ on the (γ, δ) edge because otherwise $\boldsymbol{\eta}$ would be m-connecting given Z, contradicting minimality. Since $\mathcal{G}$ is ancestral, we have $\gamma \to \delta$ as required.

(ii) *At least one of δ and γ is a non-collider on $\boldsymbol{\pi}$; neither is an endpoint.* We first show that at most one of δ and γ is a non-collider on $\boldsymbol{\pi}$. We will then show that if one of δ and γ is a collider on $\boldsymbol{\pi}$, then the said collider meets the conditions required by the Lemma.

For a contradiction, suppose both δ and γ are non-colliders on $\boldsymbol{\pi}$. Note that neither γ nor δ is in Z because $\boldsymbol{\pi}$ is m-connecting given Z. Let $\boldsymbol{\eta}$ be the path formed by concatenating the subpath $\boldsymbol{\pi}(\alpha, \gamma)$, the (γ, δ) edge, and the subpath $\boldsymbol{\pi}(\delta, \beta)$. If both δ and γ are non-colliders on $\boldsymbol{\eta}$ then $\boldsymbol{\eta}$ is m-connecting given Z, which violates the minimality of $\boldsymbol{\pi}$. Hence, at least one of δ and γ is a collider on $\boldsymbol{\eta}$. Suppose without loss of generality that γ is a collider on $\boldsymbol{\eta}$. Since there is an arrowhead at γ on $\boldsymbol{\pi}(\alpha, \gamma)$, by Lemma 2.2 either: (a) $\gamma \in \text{an}(Z)$, or (b) $\boldsymbol{\pi}(\gamma, \beta)$ is a directed path from γ to β. We consider these separately:

(a) Recall that both δ and γ are non-colliders on $\boldsymbol{\pi}$, $\delta, \gamma \notin Z$, and γ is a collider on $\boldsymbol{\eta}$. It follows that δ is also a collider on $\boldsymbol{\eta}$, and $\delta \notin \text{an}(Z)$: otherwise $\boldsymbol{\eta}$ is m-connecting given Z, regardless of whether γ is a collider or non-collider, but this violates the minimality of $\boldsymbol{\pi}$. Since there is an arrowhead at δ on $\boldsymbol{\pi}(\delta, \beta)$, but $\delta \notin \text{an}(Z)$, by Lemma 2.2, $\boldsymbol{\pi}(\delta, \alpha)$ forms a directed path to α, which contradicts the arrowhead at γ on $\boldsymbol{\pi}(\alpha, \gamma)$.

(b) Since δ is between γ and β along $\boldsymbol{\pi}$, $\boldsymbol{\pi}(\gamma, \delta)$ forms a directed path from γ to δ. Further $\gamma \to \delta$, since the graph is ancestral. Consequently δ is a non-collider on $\boldsymbol{\eta}$ and $\boldsymbol{\pi}$, so $\boldsymbol{\eta}$ is m-connecting given Z, which violates the minimality of $\boldsymbol{\pi}$.

Thus at most one of δ and γ is a non-collider on $\boldsymbol{\pi}$: without loss of generality suppose it is γ. Then δ is a collider on $\boldsymbol{\pi}$, and since $\boldsymbol{\pi}$ is m-connecting given Z, $\delta \in \mathrm{an}(Z)$ and $\gamma \notin Z$. We will now show that $\delta \to \gamma$. Suppose for a contradiction that δ is a collider on $\boldsymbol{\eta}$. Since $\boldsymbol{\eta}$ is *not* m-connecting given Z, γ is also a collider on $\boldsymbol{\eta}$, and $\gamma \notin \mathrm{an}(Z)$. But here we reach a contradiction: if γ is a collider on $\boldsymbol{\eta}$, then there is an arrowhead at γ on $\boldsymbol{\pi}(\alpha, \gamma)$, in which case by Lemma 2.2, $\gamma \in \mathrm{an}(Z)$ (since $\boldsymbol{\pi}(\gamma, \beta)$ cannot be a directed path because δ is a collider on $\boldsymbol{\pi}$). Hence, δ is not a collider on $\boldsymbol{\eta}$ and $\delta \to \gamma$ ($\delta - \gamma$ is ruled out because δ is a collider on $\boldsymbol{\pi}$).

We have now established: on $\boldsymbol{\pi}$ γ is a non-collider, while δ is a collider; $\delta \in \mathrm{an}(Z)$; $\gamma \notin Z$; $\delta \to \gamma$, so δ is a non-collider on $\boldsymbol{\eta}$. Suppose for a contradiction that $\delta \notin Z$. If γ is a non-collider on $\boldsymbol{\eta}$ then $\boldsymbol{\eta}$ is m-connecting given Z, which violates the minimality of $\boldsymbol{\pi}$. If γ is a collider on $\boldsymbol{\eta}$, then, as before, by Lemma 2.2, $\gamma \in \mathrm{an}(Z)$ and $\boldsymbol{\eta}$ is m-connecting given Z, again violating the minimality of $\boldsymbol{\pi}$. Hence we have shown that $\delta \in Z$.

(iii) *Both δ and γ are colliders.*

Since $\mathcal{G}$ is ancestral, the (γ, δ) edge is not undirected. The (γ, δ) edge cannot be bi-directed because then the subsequence $\langle \alpha, \ldots, \gamma, \delta, \ldots, \beta \rangle$ would form an m-connecting path given Z, contradicting minimality of $\boldsymbol{\pi}$. Without loss of generality, assume the (γ, δ) edge is directed out of γ. By the minimality of $\boldsymbol{\pi}$ the shorter path $\langle \alpha, \ldots, \gamma, \delta, \ldots, \beta \rangle$ is not m-connecting given Z, hence γ is in Z. $\square$

3.6 Discriminating paths on minimal m-connecting paths

The next Lemma shows that if a triple $\langle \alpha, \beta, \gamma \rangle$ on a minimal m-connecting path $\boldsymbol{\pi}$ is shielded then a subsequence of the path forms a discriminating path for β. Thus in the notation of §3.4 on a minimal m-connecting path the following holds:

$$\langle \alpha, \beta, \gamma \rangle \text{ a shielded triple on } \boldsymbol{\pi} \implies \text{there exists a non-endpoint vertex } \delta \text{ on } \boldsymbol{\pi} \text{ such that } \delta \prec_{\pi} \beta.$$

Lemma 3.13 *Let $\boldsymbol{\pi}$ be a minimal m-connecting path between a and b given Z in the maximal ancestral graph $\mathcal{G}$. If $\langle \delta, \beta, y \rangle$ is a sub-path of $\boldsymbol{\pi}$ and δ and y are adjacent then $\boldsymbol{\pi}$ contains a unique sub-path that forms a discriminating path for β.*

Proof: By Lemma 3.12, since δ and y are adjacent, non-consecutive, non-endpoint vertices, at least one is both a collider on $\boldsymbol{\pi}$ and a parent of the other. Suppose without loss of generality that $\delta \to y$, and that δ is a collider on the subpath $\boldsymbol{\pi}(a, y)$.

Claim: (See Figure 14.)

If $\boldsymbol{\pi}$ contains a vertex q_i such that every non-endpoint vertex on the subpath $\boldsymbol{\pi}(q_i, \ldots, q_0 = \delta)$ is a collider on $\boldsymbol{\pi}$ and a parent of y in $\mathcal{G}$ then $q_i \neq a$, and the vertex before q_i on $\boldsymbol{\pi}(a, q_i)$, say q_{i+1}, is such that either

(i) $\boldsymbol{\pi}(q_{i+1}, y)$ forms a discriminating path for β in $\mathcal{G}$, or

(ii) $q_{i+1} \to y$ in $\mathcal{G}$ and q_{i+1} is a collider on $\boldsymbol{\pi}$.

Assume that the antecedent for the claim holds. If $q_i = a$, then by Lemma 2.2 the subsequence $\langle a, y, \ldots, b \rangle$ forms an m-connecting path given Z that violates the minimality of $\boldsymbol{\pi}$. Thus, $q_i \neq a$, and there exists a vertex q_{i+1} before q_i on $\boldsymbol{\pi}(a, q_i)$.

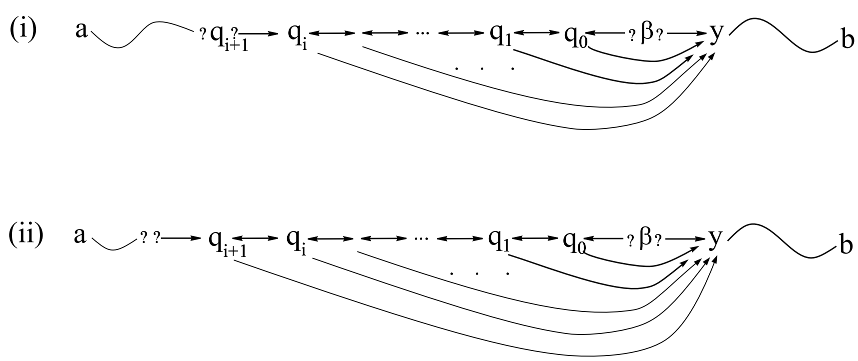

Figure 14: Unique sub-path of π forms a discriminating path for β in $\mathcal{G}$. See Lemma 3.13 for further explanation.

If q_{i+1} is not adjacent to y, then by the conditions in the antecedent of the claim, $\boldsymbol{\pi}(q_{i+1}, y)$ forms a discriminating path for β. If q_{i+1} is adjacent to y, then by the conditions on q_i we have $q_{i+1}? \!\rightarrow q_i \rightarrow y$, hence by Lemma 2.1, $q_{i+1}? \!\rightarrow y$. Further, $q_{i+1} \neq a$ since otherwise the subsequence $\langle a, y, \ldots, b\rangle$ forms an m-connecting path given Z, contradicting the minimality of $\boldsymbol{\pi}$. By Lemma 3.12, $q_{i+1} \rightarrow y$ in $\mathcal{G}$ and q_{i+1} is a collider on $\boldsymbol{\pi}$ as required.

The Lemma now follows directly. The antecedent of the claim holds trivially for $q_0 = \delta$ since δ is a collider on $\boldsymbol{\pi}$ and $\delta \rightarrow y$ in $\mathcal{G}$. By repeated applications of the claim, either there is a discriminating path $\boldsymbol{\pi}(q_{k+1}, \ldots, q_0, \beta, y\rangle$ or q_{k+1} is a non-endpoint vertex satisfying the antecedent of the claim. Since there are only finitely many vertices on the path $\boldsymbol{\pi}(a, q_0)$, for some $k \geq 0$, $\boldsymbol{\pi}(q_{k+1}, \ldots, q_0, \beta, y\rangle$ forms a discriminating path for β. Hence the conclusion holds. Uniqueness follows directly from the fact that q_{k+1} is not adjacent to y, but every vertex q_i for $i \leq k$ is a parent of y, in particular $q_0 \rightarrow y$. It follows that no other subpath of $\boldsymbol{\pi}$ forms a discriminating path for β. $\square$

3.7 Triples on minimal m-connecting paths

In this section we prove that every triple on a minimal m-connecting path has an order, and hence will be of the same type in every member of the equivalence class.

Lemma 3.14 *If $\langle \alpha, \beta, \gamma \rangle$ is a triple on a minimal m-connecting path $\boldsymbol{\pi}$ between x and y given Z in the maximal ancestral graph $\mathcal{G}$ then $\langle \alpha, \beta, \gamma \rangle$ has order.*

Proof: Suppose for a contradiction that $\langle \alpha, \beta, \gamma \rangle$ does not have order. α and γ are adjacent, since otherwise $\langle \alpha, \beta, \gamma \rangle$ is unshielded, and hence of order 0. It follows from Lemma 3.13 that there is a unique subpath of $\boldsymbol{\pi}$ which forms a discriminating path for $\langle \alpha, \beta, \gamma \rangle$. If every triple on this discriminating path has order, then by definition, $\langle \alpha, \beta, \gamma \rangle$ has order. Hence there is at least one triple which does not have order, call this $\langle \alpha_1, \beta_1, \gamma_1 \rangle$. As before it follows that α_1 and γ_1 are adjacent, and hence there is a unique subpath of $\boldsymbol{\pi}$ which forms a discriminating path for $\langle \alpha_1, \beta_1, \gamma_1 \rangle$. Arguing in this way, we can construct an infinite sequence of triples on $\boldsymbol{\pi}$ $\langle \alpha_i, \beta_i, \gamma_i \rangle$, none of which have order and such that:

$$\cdots \prec_\pi \beta_i \prec_\pi \cdots \prec_\pi \beta_1 \prec_\pi \beta$$

However, by Corollary 3.11 all of the β_i's are distinct, which is a contradiction since $\boldsymbol{\pi}$ is finite. Thus every triple on $\boldsymbol{\pi}$ has an order. □

Note that this argument shows that every triple on a minimal path $\boldsymbol{\pi}$ has *some* order; to determine *which* order, it would be necessary to consider all discriminating paths for the given triple not merely those which are subpaths of $\boldsymbol{\pi}$.

Corollary 3.15 *Suppose that $\mathcal{G}_1$ and $\mathcal{G}_2$ are maximal ancestral graphs with the same adjacencies, and the same colliders with order. If $\boldsymbol{\pi}$ is a minimal* *m-connecting path between x and y given Z in $\mathcal{G}_1$, then $\langle\alpha,\beta,\gamma\rangle$ is a collider (non-collider) on $\boldsymbol{\pi}$ in $\mathcal{G}_1$ if and only if $\langle\alpha,\beta,\gamma\rangle$ is a collider (non-collider) on the corresponding path $\boldsymbol{\pi}^*$ in $\mathcal{G}_2$.*

Proof: This follows directly from Proposition 3.7 and Lemma 3.14. $\square$

3.8 Directed paths from colliders to vertices in Z

To show that the corresponding path $\boldsymbol{\pi}^*$ is m-connecting given Z in $\mathcal{G}^*$, it is also necessary to show that every collider along the path is an ancestor of Z in $\mathcal{G}^*$. The next Lemmas establish that if there is an m-connecting path between x and y given Z in $\mathcal{G}_1$ then we can always find a path $\boldsymbol{\pi}$ m-connecting x and y given Z in $\mathcal{G}_1$ such that if a collider on $\boldsymbol{\pi}$ is an ancestor of a vertex z in $\mathcal{G}_1$ then the collider is also an ancestor of z in any graph $\mathcal{G}_2$ which contains the same adjacencies and triples with order as $\mathcal{G}_1$. For this purpose we require the following:

Definition 3.7 *An m-connecting path $\boldsymbol{\pi}$ between x and y given Z is said to be a closest m-connecting path to Z if there is no collider $\langle\alpha,\beta,\gamma\rangle$ on $\boldsymbol{\pi}$ such that:*

(i) The first vertex after β on a directed path from β to a vertex in Z, say β^*, is adjacent to a vertex α^* on $\boldsymbol{\pi}(x,\alpha)$, and β^* is adjacent to a vertex γ^* on $\boldsymbol{\pi}(\beta,y)$, and

(ii) $\langle\alpha^*,\beta^*,\gamma^*\rangle$ forms a collider, and

(iii) The path formed by concatenating $\boldsymbol{\pi}(x,\alpha^*)$, $\langle\alpha^*,\beta^*,\gamma^*\rangle$ and $\boldsymbol{\pi}(\gamma^*,y)$ forms a path m-connecting given Z.

In other words, if $\boldsymbol{\pi}$ is a closest path to Z then it is not possible to form a new m-connecting path with a shorter directed path from the collider to a vertex in Z via the replacement described. It is easy to see that if there is an m-connecting path between x and y given Z then there is a closest m-connecting path to Z. A directed path is said to be *minimal* if no subsequence of vertices with the same endpoints forms a directed path.

Proposition 3.16 *If* $\boldsymbol{\pi}$ *is an m-connecting path between* x *and* y *given* Z *then there exists a path* $\boldsymbol{\pi}^*$ *that is both minimal and closest to* Z.

Proof: For a given path $\boldsymbol{\pi}'$, suppose $c_1, \ldots, c_k$ are the colliders on $\boldsymbol{\pi}'$ and, for each i $(1 \leq i \leq k)$, $\boldsymbol{\delta}_i$ is a minimal directed path (possibly of length 0) from c_i to some vertex $z_i \in Z$. Let

$$\phi(\boldsymbol{\pi}', Z) = |\boldsymbol{\pi}'| + \sum_{i=1}^{k} |\boldsymbol{\delta}_i|$$

where $|\cdot|$ counts the number of edges on a given path; thus $\phi(\boldsymbol{\pi}', Z)$ simply counts the number of edges on $\boldsymbol{\pi}$ and on directed paths from colliders on $\boldsymbol{\pi}$ to vertices in Z.

We now perform the following procedure:

(0) Let $i = 0$, $\boldsymbol{\pi}_0 = \boldsymbol{\pi}$.

(1) If $\boldsymbol{\pi}_i$ is not minimal, then let $\boldsymbol{\pi}_i^*$ be a minimal m-connecting path between x and y given Z, formed from a subsequence of vertices on $\boldsymbol{\pi}_i$, otherwise let $\boldsymbol{\pi}_i^* = \boldsymbol{\pi}_i$.

(2) If $\boldsymbol{\pi}_i^*$ is not a closest path to Z then let $\boldsymbol{\pi}_{i+1}$ be a closest path to Z formed from $\boldsymbol{\pi}_i^*$ by repeated replacements of the type described in Definition 3.7, otherwise let $\boldsymbol{\pi}_{i+1} = \boldsymbol{\pi}_i^*$.

(3) If $\boldsymbol{\pi}_{i+1} = \boldsymbol{\pi}_i$ then step; otherwise increase i by 1 and go to step (1).

(Note that although $\boldsymbol{\pi}_i^*$ is minimal, there is no guarantee that $\boldsymbol{\pi}_{i+1}$ will be minimal, hence several iterations may be required.)

Clearly if the procedure terminates then we have a minimal m-connecting path that is also a closest path to Z. To see that the procedure will always terminate it is sufficient to notice that $\phi(\boldsymbol{\pi}_i, Z) \geq \phi(\phi(\boldsymbol{\pi}_{i+1}, Z) > 0$ with equality only if $\boldsymbol{\pi}_{i+1} = \boldsymbol{\pi}_i$. □

Lemma 3.17 *If $\boldsymbol{\pi}$ is a minimal m-connecting path between x and y which is closest to Z, then if $\langle \alpha_1, \beta, \gamma_1 \rangle$ is a collider on $\boldsymbol{\pi}$, and β^* is the first vertex after β on a directed path from β to a vertex in Z then at least one of the non-colliders $\langle \alpha_1, \beta, \beta^* \rangle$, $\langle \gamma_1, \beta, \beta^* \rangle$ has order.*

Proof: Suppose for a contradiction that both $\langle \alpha_1, \beta, \beta^* \rangle$ and $\langle \gamma_1, \beta, \beta^* \rangle$ do not have order. Note that α_1 and γ_1 are both adjacent to β^*: otherwise one of the non-colliders is unshielded, and hence has order. By Lemma 2.1, we further have $\alpha_1 ?\!\rightarrow \beta^*$ and $\gamma_1 ?\!\rightarrow \beta^*$.

There exists a vertex ν between x and α_1 along $\boldsymbol{\pi}$ such that either:

(i) ν is not adjacent to β^*,

(ii) ν is a child of β^*, or

(iii) if $\boldsymbol{\pi}^*$ is the path formed by concatenating the subpath $\boldsymbol{\pi}(x, \nu)$ and the (ν, β^*) edge then ν is either (a) a collider on both $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^*$ or (b) a non-collider on both paths, or (c) $\nu = x$.

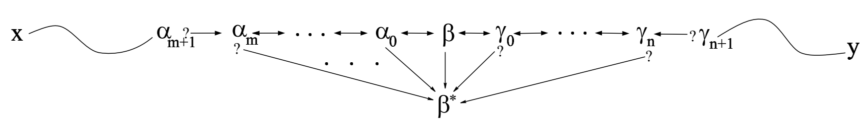

Figure 15: Diagram for the proof of Lemma 3.17: either $\langle x, \ldots, \alpha_m, \beta^*, \gamma_n, \ldots, y \rangle$ is an m-connecting path closer to Z, or at least one of the non-colliders $\langle \alpha_0, \beta, \beta^* \rangle$ and $\langle \gamma_0, \beta, \beta^* \rangle$ has order.

Such a vertex is guaranteed to exist because x satisfies either clause (i) or (iii)(c). (Note that $\nu - \beta^*$ is ruled out since $\beta \to \beta^*$.) Let α_{m+1} be the vertex closest to β^* on the path $\boldsymbol{\pi}(x, \beta)$ which satisfies these conditions, and let $\langle \alpha_m, \ldots, \alpha_1 \rangle$ be the non-endpoint vertices on the sub-path $\boldsymbol{\pi}(\alpha_{m+1}, \beta)$. See Figure 15. By construction, every vertex α_j $(1 \leq j \leq m)$ fails to satisfy any of these conditions, hence each α_j is adjacent to β^*, is not a child of β^*, and is either a collider on $\boldsymbol{\pi}$ and a non-collider on $\boldsymbol{\pi}^*$ or vice versa.

We show by induction that for $1 \leq j \leq m$, $\alpha_j \to \beta^*$ hence α_j is a collider on $\boldsymbol{\pi}$. By hypothesis, there is an edge $\beta \to \beta^*$ and β is a collider on $\boldsymbol{\pi}$ in $\mathcal{G}$, hence we have $\alpha_1 ?\!\!\to \beta \to \beta^*$. By construction, α_j $(j \leq m)$ is not of the same type on both $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^*$, hence by Lemma 2.1 we have $\alpha_1 \to \beta^*$ and α_1 is a collider on $\boldsymbol{\pi}$. If there is an edge $\alpha_j \to \beta^*$ $(j < m)$, and α_j is a collider on $\boldsymbol{\pi}$, then again by the construction and Lemma 2.1 we have $\alpha_{j+1} \to \beta^*$ and α_{j+1} is a collider on $\boldsymbol{\pi}$.

Now consider the vertex α_{m+1}: by construction α_{m+1} satisfies one of the conditions (i)-(iii). Since we have either $\alpha_1 ?\!\!\to \beta \to \beta^*$ (if $m + 1 = 1$) or $\alpha_{m+1} ?\!\!\to \alpha_m \to \beta^*$ (if $m + 1 > 1$), it follows that $\alpha_{m+1} \leftarrow \beta^*$ is ruled out: otherwise $\mathcal{G}$ would not be ancestral. Thus (ii) is ruled out.

If α_{m+1} is not adjacent to β^* then the path formed by the sequence of vertices $\langle \alpha_{m+1}, \ldots, \alpha_1, \beta, \beta^* \rangle$ either forms an unshielded non-collider ($m = 1$), or forms a discriminating path for $\langle \alpha_1, \beta, \beta^* \rangle$ $(m > 1)$. Further, since $\boldsymbol{\pi}$ is minimal, by Lemma 3.14 every one of the colliders $\langle \alpha_{m+1}, \alpha_m, \alpha_{m-1} \rangle$, $\ldots$, $\langle \alpha_2, \alpha_1, \beta \rangle$ has order. Consequently $\langle \alpha_1, \beta, \beta^* \rangle$ has order, which is a contradiction. Hence (i) is ruled out.

Thus condition (iii) holds, so α_{m+1} is either an endpoint, or of the same type on both $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^*$. Further, $\alpha_{m+1} ?\!\!\to \beta^*$ since (ii) is ruled out.

By a symmetric argument we may show that there is a vertex γ_{n+1} on $\boldsymbol{\pi}(\beta, y)$, such that $\gamma_{n+1} ?\!\!\to \beta^*$ and γ_{n+1} is either an endpoint, or is of the same type on both $\boldsymbol{\pi}$ and the path formed by concatenating the $\beta^* \leftarrow ?\gamma_{n+1}$ edge and the subpath $\boldsymbol{\pi}(\gamma_{n+1}, y)$.

Let $\boldsymbol{\mu}$ be the path formed by concatenating the subpaths $\boldsymbol{\pi}(x, \alpha_{m+1})$, $\alpha_{m+1}? \rightarrow \beta^* \leftarrow ?\gamma_{n+1}$ and $\boldsymbol{\pi}(\gamma_{n+1}, y)$ (if $x = \alpha_{m+1}$, or $\gamma_{n+1} = y$ then omit the relevant subpath). $\boldsymbol{\mu}$ forms an m-connecting path given Z because α_{m+1} and γ_{n+1} are of the same type as they are on $\boldsymbol{\pi}$, and β^* is an ancestor of Z. However $\boldsymbol{\mu}$ is closer to Z than $\boldsymbol{\pi}$, which is a contradiction. □

Lemma 3.18 *If $\boldsymbol{\delta}$ is a minimal directed path from γ to z in an ancestral graph $\mathcal{G}$ then every non-collider on $\boldsymbol{\delta}$ is unshielded and hence has order.*

Proof: Suppose that $\langle a, b, c\rangle$ is a non-collider on $\boldsymbol{\delta}$. In this case we have $a \rightarrow b \rightarrow c$. If a and c are adjacent then, by the ancestral property, we have $a \rightarrow c$, which contradicts the minimality of $\boldsymbol{\delta}$. □

Though we do not need it here, in fact no non-consecutive vertices on $\boldsymbol{\delta}$ are adjacent.

Corollary 3.19 *Suppose that $\mathcal{G}_1$ and $\mathcal{G}_2$ are maximal ancestral graphs with the same adjacencies, and the same colliders with order. If in $\mathcal{G}_1$: $\boldsymbol{\pi}$ is a minimal m-connecting path between x and y given Z, $\boldsymbol{\pi}$ is a closest path to Z, $\langle \alpha, \beta, \gamma\rangle$ is a collider on $\boldsymbol{\pi}$, and $\boldsymbol{\delta}$ forms a minimal directed path from β to a vertex $z \in Z$, then the corresponding path $\boldsymbol{\delta}^*$ is a directed path in $\mathcal{G}_2$.*

Proof: Let $\langle \beta = \nu_0, \nu_1, \ldots, \nu_n = z\rangle$ be the vertices on δ. The proof is by induction on the edges (ν_i, ν_{i+1}) of the path $\boldsymbol{\delta}^*$.

Base case: (i=1) Since $\langle \alpha, \beta\gamma\rangle$ is a collider on $\boldsymbol{\pi}$ in $\mathcal{G}_1$, and $\boldsymbol{\pi}$ is minimal, by Corollary 3.15 $\langle \alpha, \beta, \gamma\rangle$ also forms a collider in $\mathcal{G}_2$. By Lemma 3.17 at least one of the non-colliders, $\langle \alpha, \beta, \nu_1\rangle$, $\langle \gamma, \beta, \nu_1\rangle$ has order in $\mathcal{G}_1$, and by Proposition 3.7 is also a non-collider in $\mathcal{G}_2$. It then follows that $\beta \rightarrow \nu_1$ in $\mathcal{G}_2$ as required (in fact, both non-colliders, are present in $\mathcal{G}_2$).

Inductive step: $(1 < i < n)$ Assume that the subpath $\boldsymbol{\delta}^*(\beta, \nu_i)$ forms a directed path from β to ν_i. By Lemma 3.18 the non-collider $\langle \nu_{i-1}, \nu_i, \nu_{i+1} \rangle$ has order, and hence is a non-collider in $\mathcal{G}_2$. By the induction hypothesis we have $\nu_{i-1} \to \nu_i$ in $\mathcal{G}_2$, hence $\nu_i \to \nu_{i+1}$ in $\mathcal{G}_2$ as required. □

3.9 Characterization of Markov equivalence

We now state and prove the main result of this paper.

Theorem 3.20 *Maximal ancestral graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ are Markov equivalent if and only if $\mathcal{G}_1$ and $\mathcal{G}_2$ have the same adjacencies and the same colliders with order.*

Proof: $(\Rightarrow)$ Since $\mathcal{G}_1$ and $\mathcal{G}_2$ have the same adjacencies and colliders with order, by Proposition 3.7, $\mathcal{G}_1$ and $\mathcal{G}_2$ also have the same non-colliders with order.

If x and y are m-connected given Z in $\mathcal{G}_1$, then by Proposition 3.16 there exists a minimal m-connecting path $\boldsymbol{\pi}$, between x and y which is closest to Z in $\mathcal{G}_1$. By Corollary 3.15 every triple on $\boldsymbol{\pi}$ is of the same type on the corresponding path $\boldsymbol{\pi}^*$ in $\mathcal{G}_2$. Hence every non-collider on $\boldsymbol{\pi}^*$ is not in Z. Since $\boldsymbol{\pi}$ is m-connecting, every collider γ on $\boldsymbol{\pi}$ is an ancestor of Z, hence if $\gamma \notin Z$ then there exists a minimal directed path $\boldsymbol{\delta}_\gamma$ from γ to some vertex $z_\gamma \in Z$. By Corollary 3.19 the corresponding path $\boldsymbol{\delta}^*_\gamma$ forms a directed path from γ to z_γ in $\mathcal{G}_2$. Thus every collider on $\boldsymbol{\pi}^*$ is an ancestor of Z in $\mathcal{G}_2$. Thus $\boldsymbol{\pi}^*$ m-connects x and y given Z in $\mathcal{G}_2$.

Likewise, it is easy to see (by symmetry) that an m-connecting path in $\mathcal{G}_2$ implies that there is an m-connecting path in $\mathcal{G}_1$, hence $\mathcal{G}_1$ and $\mathcal{G}_2$ are Markov equivalent.

$(\Leftarrow)$ Conversely, if $\mathcal{G}_1$ and $\mathcal{G}_2$ are Markov equivalent then by Proposition 3.3 they have the same adjacencies, and by Proposition 3.6 they have the same colliders with order. □

4 Discussion and relation to other work

In prior work, Spirtes and Richardson (1997), provide the following characterization of Markov equivalence for ancestral graphs:

Theorem 4.1 *Two maximal ancestral graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ are Markov equivalent if and only if:*

(i) *$\mathcal{G}_1$ and $\mathcal{G}_2$ have the same adjacencies;*

(ii) *$\mathcal{G}_1$ and $\mathcal{G}_2$ have the same unshielded colliders; and*

(iii) *If $\boldsymbol{\pi}$ forms a discriminating path for β in $\mathcal{G}_1$ and $\mathcal{G}_2$, then β is a collider on $\boldsymbol{\pi}$ in $\mathcal{G}_1$ if and only if it is a collider on $\boldsymbol{\pi}$ in $\mathcal{G}_2$.*

Theorem 3.20 is a stronger result since it requires verifying the presence of only those triples with order, whereas clause (iii) of this result requires us to verify that if there is a discriminating path in both $\mathcal{G}_1$ and $\mathcal{G}_2$ then the triple discriminated is a collider or non-collider in both. However, as shown in Figure 7 it is possible for a discriminating path to contain triples without order, and thus also to discriminate triples without order, which Theorem 3.20 demonstrates is clearly unnecessary.

More significantly, to verify the conditions of Theorem 4.1, in principle we need to find every discriminating path for a given triple: otherwise it is possible that, although a triple is discriminated by *some* path in $\mathcal{G}_1$ and *some* path in $\mathcal{G}_2$, in fact there is no discriminating path that is common to both graphs. Since there may be exponentially many such paths, finding them all would not be feasible in polynomial-time.

By contrast, verifying the conditions of Theorem 3.20 requires $O(n^3e^2)$ calculations, where the graphs have n vertices, and e edges; see Spirtes and Richardson (1997). (The algorithm presented in Spirtes and Richardson (1997) verifies the conditions of Theorem 3.20, rather than Theorem 4.1, though the paper only proves the latter result.)

4.1 Summary graphs and MC graphs

The problem of constructing graphical representations for the independence structure of DAGs under marginalizing *and* conditioning was originally posed by N. Wermuth in 1994 in a lecture at CMU. Wermuth, Cox and Pearl developed an approach to this problem based on *summary graphs* (see Wermuth et al., 1999, Cox and Wermuth, 1996, Wermuth and Cox, 2000).

For a given summary graph $\mathcal{H}$ it is always possible to construct a DAG $\mathcal{D}(\mathcal{H})$ with additional variables such that the DAG is Markov equivalent to $\mathcal{H}$ after marginalizing and conditioning. Consequently, it is always possible to transform a summary graph into an ancestral graph, via the transformation described in §2.6. Hence, via this transformation, the results in this paper also provide an algorithm for determining the Markov equivalence of two summary graphs.

More recently J. Koster has introduced another class of graphs, called *MC-graphs*, together with an operation of marginalizing and conditioning. (See Koster, 2002, 1999a,b). For MC-graphs it is not always the case that there exists some DAG which is Markov equivalent to the MC-graph under marginalizing and conditioning. However, for the subclass of MC-graphs which are Markov equivalent to DAGs with additional variables under marginalizing and conditioning, we may again apply the results of this paper to establish Markov equivalence.

4.2 Future Work

Meek (1995), and Andersson et al. (1997) identified those edges which were present in every graph in a Markov equivalence class of DAGs. To date, no such characterization has been provided for ancestral graphs, though partial characterizations of Markov equivalence classes for ancestral graphs have been obtained using POIPGs and PAGs by Richardson and Spirtes (2003) and Spirtes et al. (1993). Ali and Richardson (2002) represented Markov equivalence classes of maximal ancestral graphs in a way analogous to that given by Andersson et al. (1997) for DAGs, but a full characterization of the resulting graph is still work in progress. Chickering (1995) and Chickering (2002b) provided transformational characterizations of Markov equivalence and the submodel relation for DAGs. Finding such characterizations for ancestral graphs remains an open problem.

Acknowledgments

This research was supported by the National University of Singapore, the U.S. National Science Foundation under grant DMS-9972008, the U.S. Environmental Protection Agency and the William and Flora Hewlett Foundation.

References

Ali, R. A. and T. Richardson (2002). Markov equivalence classes for maximal ancestral graphs. In A. Darwiche and N. Friedman (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the* 18^{th} *Conference*, San Francisco, pp. 1–8. Morgan Kaufmann.

Andersson, S. A., D. Madigan, and M. D. Perlman (1997). A characterization of Markov equivalence classes for acyclic digraphs. *Ann. Statist. 25*, 505–541.

Chickering, D. (1995). A transformational characterization of equivalent bayesian networks. In P. Besnard and S. Hanks (Eds.), *Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence*, San Mateo, CA, pp. 87–98. Morgan Kaufmann.

Chickering, D. (2002a). Learning equivalence classes of Bayesian network structures. *Journal of Machine Learning Research 2*, 445–498.

Chickering, D. (2002b). Optimal structure identification with greedy search. *Journal of Machine Learning Research 3*, 507–554.

Cooper, G. F. (1995). Causal discovery from data in the presence of selection bias. In D. Fisher (Ed.), *Preliminary papers of the Fifth International Workshop on AI and Statistics, January 4-7, Fort Lauderdale, Florida*, pp. 140–150.

Cox, D. R. and N. Wermuth (1996). *Multivariate Dependencies: Models, Analysis and Interpretation.* London: Chapman and Hall.

Dawid, A. (1980). Conditional independence for statistical operations. *Ann. Statist. 8*, 598–617.

Drton, M. (2004). Maximum likelihood estimation in seemingly unrelated regressions using algebraic geometry. *Journal of Symbolic Computation*. submitted.

Drton, M., S. A. Andersson, and M. D. Perlman (2003). Conditional independence models for seemingly unrelated regressions with incomplete data. Technical Report 431, Dept. of Statistics, University of Washington.

Drton, M. and T. Richardson (2004a). Simplicial sets and orientable edges in graphical models for marginal independence. Technical report, Department of Statistics, University of Washington.

Drton, M. and T. S. Richardson (2004b). Iterative conditional fitting for Gaussian ancestral graph models. In U. Kjærulff and C. Meek (Eds.), *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pp. 130–137. San Francisco: Morgan Kaufmann.

Drton, M. and T. S. Richardson (2004c). Multimodality of the likelihood in the bivariate seemingly unrelated regressions model. *Biometrika 91*, 383–392.

Frydenberg, M. (1990). The chain graph Markov property. *Scand. J. Statist. 17*, 333–353.

Gillispie, S. and M. D. Perlman (2001). Enumerating Markov equivalence classes of acyclic digraph models. In J. Breese and D. Koller (Eds.), *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, pp. 171–177. San Francisco: Morgan Kaufmann.

Koster, J. T. A. (1999a, September). Linear structural equations and graphical models. Lecture notes, The Fields Institute, Toronto.

Koster, J. T. A. (1999b, October). On the validity of the Markov interpretation of path diagrams of Gaussian structural equation systems with correlated errors. *Scand. J. Statist. 26*(3), 413–431.

Koster, J. T. A. (2002). Marginalizing and conditioning in graphical models. *Bernoulli 8*(6), 817–840.

Lauritzen, S. (1996). *Graphical Models*. Number 17 in Oxford Statistical Science Series. Oxford, UK: Clarendon Press.

Meek, C. (1995). Causal inference and causal explanation with background knowledge. In P. Besnard and S. Hanks (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the* 11^{th} *Conference*, San Francisco, pp. 403–410. Morgan Kaufmann.

Pearl, J. (2000). *Causality: models, reasoning and inference.* Cambridge, UK: Cambridge University Press.

Richardson, T. and P. Spirtes (2003). Causal inference via ancestral graph models. In P. Green, N. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*. Oxford: Oxford University Press.

Richardson, T. S. and P. Spirtes (2002). Ancestral graph Markov models. *Annals of Statistics 30*(4), 962–1030.

Spirtes, P., C. Glymour, and R. Scheines (1993). *Causation, Prediction and Search*. Number 81 in Lecture Notes in Statistics. Springer-Verlag.

Spirtes, P. and T. S. Richardson (1997). A polynomial-time algorithm for determining DAG equivalence in the presence of latent variables and selection bias. In D. Madigan and P. Smyth (Eds.), *Preliminary papers of the Sixth International Workshop on AI and Statistics, January 4-7, Fort Lauderdale, Florida*, pp. 489–501.

Verma, T. and J. Pearl (1990). Equivalence and synthesis of causal models. In M. Henrion, R. Shachter, L. Kanal, and J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the* 6th *Conference*, Mountain View, CA, pp. 220–227. Association for Uncertainty in AI.

Verma, T. and J. Pearl (1991). Equivalence and synthesis of causal models. Technical Report R-150, Cognitive Systems Laboratory, UCLA.

Wermuth, N. and D. Cox (2000). A sweep operator for triangular matrices and its statistical applications. Technical Report 00-04, ZUMA Institute, Mannheim, Germany.

Wermuth, N., D. Cox, and J. Pearl (1999). Explanations for multivariate structures derived from univariate recursive regressions. Technical Report Revision of 94-1, University of Mainz, Mainz, Germany.

Zellner, A. (1962). An efficient method of estimating seemingly unrelated regression equations and tests for aggregation bias. *JASA 57*, 348–368.