

On Efficient Adjustment in Causal Graphs

Janine Witte

WITTE@LEIBNIZ-BIPS.DE

*Leibniz Institute for Prevention Research and Epidemiology—BIPS, Bremen, Germany
and Faculty of Mathematics and Computer Science, University of Bremen, Germany*

Leonard Henckel

HENCKEL@STAT.MATH.ETHZ.CH

Seminar for Statistics, ETH Zurich, Switzerland

Marloes H. Maathuis

MAATHUIS@STAT.MATH.ETHZ.CH

Seminar for Statistics, ETH Zurich, Switzerland

Vanessa Didelez

DIDELEZ@LEIBNIZ-BIPS.DE

*Leibniz Institute for Prevention Research and Epidemiology—BIPS, Bremen, Germany
and Faculty of Mathematics and Computer Science, University of Bremen, Germany*

Editor: Peter Spirtes

Abstract

We consider estimation of a total causal effect from observational data via covariate adjustment. Ideally, adjustment sets are selected based on a given causal graph, reflecting knowledge of the underlying causal structure. Valid adjustment sets are, however, not unique. Recent research has introduced a graphical criterion for an ‘optimal’ valid adjustment set (**O**-set). For a given graph, adjustment by the **O**-set yields the smallest asymptotic variance compared to other adjustment sets in certain parametric and non-parametric models. In this paper, we provide three new results on the **O**-set. First, we give a novel, more intuitive graphical characterisation: We show that the **O**-set is the parent set of the outcome node(s) in a suitable latent projection graph, which we call the forbidden projection. An important property is that the forbidden projection preserves all information relevant to total causal effect estimation via covariate adjustment, making it a useful methodological tool in its own right. Second, we extend the existing IDA algorithm to use the **O**-set, and argue that the algorithm remains semi-local. This is implemented in the R-package `pcalg`. Third, we present assumptions under which the **O**-set can be viewed as the target set of popular non-graphical variable selection algorithms such as stepwise backward selection.

Keywords: causal discovery, causal inference, confounder selection, confounding, efficiency, graphical models, IDA algorithm, model selection, sufficient adjustment set

1. Introduction

In typical analyses of observational data, we wish to estimate the total causal effect of a (possibly multivariate) treatment or exposure \mathbf{X} on a (possibly multivariate) outcome \mathbf{Y} . Ideally, we can fully specify the underlying causal directed acyclic graph (DAG). We can then use a graphical adjustment criterion, e.g. Pearl’s back-door criterion (Pearl, 2009) or the generalised adjustment criterion (Perković et al., 2015, 2018; Shpitser et al., 2010), to check whether a set of covariates is valid for adjustment. However, there may be more than one valid adjustment set. Although all resulting estimators are then consistent, their variances may differ considerably.

There are several approaches to choose an adjustment set among all valid adjustment sets. For example, one can pick a minimal adjustment set (de Luna et al., 2011; Textor and Liškiewicz, 2011). An alternative strategy is to aim at decreasing the causal effect estimator’s variance by including variables associated with the outcome (e.g. Brookhart et al., 2006; Lunceford and Davidian, 2004; Shortreed and Ertefaie, 2017). Witte and Didelez (2019) referred to this strategy as the ‘outcome-oriented’ approach. It is especially popular when little graphical knowledge is available. A major advancement for the outcome-oriented approach was the graphical characterisation of the ‘optimal’ adjustment set (**O**-set) by Henckel et al. (2019) (HPM19). They showed that under a linear model, adjusting for the **O**-set yields the smallest asymptotic variance for the causal effect estimator compared to all other valid adjustment sets, under assumptions detailed below. Strengthening this result, Rotnitzky and Smucler (2020) (RS20) recently showed that the minimal variance property of the **O**-set is retained for a class of non-parametric estimators. All these results apply to DAGs, as well as so-called amenable completed partially directed acyclic graphs (CPDAGs; see e.g. Andersson et al., 1997) and amenable maximally oriented partially directed acyclic graphs (maxPDAGs; see Perković et al., 2017). These are larger classes of graphs allowing for undirected edges where the direction cannot be decided. Amenability implies that despite the undirected edges, an adjustment set can be identified from the CPDAG (or maxPDAG) so that this set is valid for adjustment in all DAGs in the equivalence class. If a CPDAG (or maxPDAG) is not amenable, no common adjustment set for all DAGs in the equivalence class exists (Perković et al., 2018), and hence different DAGs may imply different true causal effects of \mathbf{X} on \mathbf{Y} . However, it is then still possible to estimate a multiset of possible causal effects (meaning that all effects in the multiset are compatible with the non-amenable graph) using the IDA algorithm by Maathuis et al. (2009, 2010).

In this paper, we provide three new results on efficient causal effect estimation. First, after briefly reviewing the results of HPM19 and RS20 (Section 2), we provide an alternative, intuitive characterisation of the **O**-set. This is based on the new concept of a forbidden projection, which has many interesting properties regarding adjustment for confounding (Section 3). Second, we extend the application of the **O**-set to non-amenable CPDAGs and maxPDAGs, by incorporating optimal adjustment into the IDA algorithm (Section 4). Third, we discuss how and under what assumptions the **O**-set can be viewed as the target set of data-driven variable selection methods such as backward model selection (Section 5).

2. Optimal Adjustment for Known Causal Structure

We begin by clarifying our setting and defining the **O**-set, before proposing an alternative definition in Section 3. We defer most of the terminology and formal definitions to Appendix A; here we only state some key concepts.

(Possibly) causal nodes and forbidden nodes. Let \mathcal{G} be a causal DAG, CPDAG or maxPDAG. A path (V_1, \dots, V_m) in \mathcal{G} is called *causal* from V_1 to V_m if $V_i \rightarrow V_{i+1}$ for all $i \in \{1, \dots, m-1\}$. It is called *possibly causal* if there are no $i, j \in \{1, \dots, m\}$, $i < j$, such that $V_i \leftarrow V_j$. Otherwise it is called *non-causal* from V_1 to V_m . A path from \mathbf{X} to \mathbf{Y} is *proper* if only its first node is in \mathbf{X} . If there is a causal path from V_1 to V_m in \mathcal{G} , then V_m is called a *descendant* of V_1 in \mathcal{G} . Analogously, if there is a possibly causal path from V_1 to V_m in \mathcal{G} , then V_m is called a *possible descendant* of V_1 in \mathcal{G} . The set of all

descendants of V_1 in \mathcal{G} is denoted by $\text{de}(V_1, \mathcal{G})$, and the set of all possible descendants by $\text{possde}(V_1, \mathcal{G})$. The *causal nodes* with respect to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} , denoted by $\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$, are the nodes on proper causal paths from \mathbf{X} to \mathbf{Y} , excluding \mathbf{X} itself. The *possibly causal nodes* $\text{posscn}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$ are defined analogously. The *forbidden set* with respect to (\mathbf{X}, \mathbf{Y}) and \mathcal{G} is defined as $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \text{possde}(\text{posscn}(\mathbf{X}, \mathbf{Y}, \mathcal{G}), \mathcal{G}) \cup \mathbf{X}$. In a DAG, this simplifies to $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \text{de}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{G}), \mathcal{G}) \cup \mathbf{X}$. The nodes in the forbidden set are called *forbidden nodes*. It can be shown that valid adjustment sets never contain forbidden nodes (Perković et al., 2018).

Valid adjustment sets. We consider a set of treatments \mathbf{X} and a set of outcomes \mathbf{Y} . A (possibly empty) set \mathbf{Z} is a *valid adjustment set* relative to (\mathbf{X}, \mathbf{Y}) if the interventional distribution $f(\mathbf{y} \mid \text{do}(\mathbf{x}))$ of \mathbf{Y} , given we set \mathbf{X} to \mathbf{x} by intervention, factorises as follows:

$$f(\mathbf{y} \mid \text{do}(\mathbf{x})) = \begin{cases} f(\mathbf{y} \mid \mathbf{x}) & \text{if } \mathbf{Z} = \emptyset, \\ \int_{\mathbf{z}} f(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) f(\mathbf{z}) d\mathbf{z} & \text{otherwise.} \end{cases}$$

Valid adjustment sets can be read off from a given causal DAG, CPDAG or maxPDAG \mathcal{G} using the *generalised adjustment criterion* (Perković et al., 2017, 2018; Shpitser et al., 2010), which generalises Pearl’s back-door criterion (Pearl, 2009): \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} if and only if the following three conditions hold: (a) every proper possibly causal path from $(\mathbf{X}$ to $\mathbf{Y})$ starts with a directed edge out of \mathbf{X} , (b) $\mathbf{Z} \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \emptyset$, (c) all proper non-causal definite-status paths from \mathbf{X} to \mathbf{Y} are blocked by \mathbf{Z} . Property (a) is called *amenability*. See Appendix A for the definition of a definite-status path. In a DAG, all paths are of definite status.

We consider two model classes and corresponding strategies for estimating causal effects when a valid adjustment set is available: (i) the causal linear model with possibly non-Gaussian error terms, where causal effects are estimated using linear regression (HPM19), and (ii) the more general non-parametric causal model, where estimation proceeds non-parametrically (RS20). In both settings, we assume an underlying causal DAG, and that we observe all variables displayed as nodes in the DAG, i.e. there are no latent variables.

Causal linear models (HPM19). A causal linear model is a causal DAG where every edge represents a linear causal effect. In a causal linear model, the (joint) causal effect of $\mathbf{X} = \{X_1, \dots, X_{k_x}\}$ on $\mathbf{Y} = \{Y_1, \dots, Y_{k_y}\}$ is defined as the matrix $\boldsymbol{\tau}_{\mathbf{y}\mathbf{x}}$ with elements

$$\begin{aligned} (\boldsymbol{\tau}_{\mathbf{y}\mathbf{x}})_{j,i} &= \frac{\partial}{\partial x_i} \text{E}(Y_j \mid \text{do}(x_1, \dots, x_{k_x})) \\ &= \text{E}(Y_j \mid \text{do}(x_1, \dots, x_i + 1, \dots, x_{k_x})) - \text{E}(Y_j \mid \text{do}(x_1, \dots, x_{k_x})), \end{aligned}$$

where element $(\boldsymbol{\tau}_{\mathbf{y}\mathbf{x}})_{j,i}$ corresponds to the controlled direct effect (Robins and Greenland, 1992; Pearl, 2001) of X_i on Y_j relative to \mathbf{X} . In other words, $(\boldsymbol{\tau}_{\mathbf{y}\mathbf{x}})_{j,i}$ is the difference in $\text{E}(Y_j)$ when \mathbf{X} is set to $(x_1, \dots, x_i + 1, \dots, x_{k_x})$ by intervention, compared to when \mathbf{X} is set to (x_1, \dots, x_{k_x}) by intervention. We can compute the effect of more general interventions as functions of the elements of $\boldsymbol{\tau}_{\mathbf{y}\mathbf{x}}$; for example, the sum of the first row corresponds to the effect on Y_1 of increasing all elements of (x_1, \dots, x_{k_x}) by one. Given a valid adjustment set \mathbf{Z} for the effect of \mathbf{X} on \mathbf{Y} , $\boldsymbol{\tau}_{\mathbf{y}\mathbf{x}}$ can be rewritten as a matrix of regression coefficients as follows: Denote by $\boldsymbol{\beta}_{\mathbf{y}\mathbf{x}, \mathbf{z}}$ the $(k_y \times k_x)$ -matrix whose (j, i) -th element is the regression coefficient $\beta_{y_j x_i \cdot \mathbf{x}_{-i} \mathbf{z}}$ of X_i in a linear regression of Y_j on X_i and $\mathbf{Z} \cup \mathbf{X}_{-i}$, where $\mathbf{X}_{-i} = \mathbf{X} \setminus \{X_i\}$. Then

$\tau_{\mathbf{y}\mathbf{x}} = \beta_{\mathbf{y}\mathbf{x},\mathbf{z}}$. The ordinary least squares (OLS) estimator $\hat{\beta}_{\mathbf{y}\mathbf{x},\mathbf{z}}$ is a consistent estimator of $\beta_{\mathbf{y}\mathbf{x},\mathbf{z}}$. We denote the asymptotic variance of $\hat{\beta}_{y_j x_i, \mathbf{x}_{-i}\mathbf{z}}$ by $a.var(\hat{\beta}_{y_j x_i, \mathbf{x}_{-i}\mathbf{z}})$.

Non-parametric estimation of causal effects (RS20). In the more general setting of a causal DAG without linearity or other assumptions on the functional form, we define the causal effect of \mathbf{X} on \mathbf{Y} as follows. Let \mathcal{X} be the set of values that \mathbf{X} can take. For a pair of vectors $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, the causal effect of intervening to set \mathbf{X} to \mathbf{x} vs. \mathbf{x}' is the vector $\Delta_{\mathbf{y}\mathbf{x}\mathbf{x}'}$ with elements

$$(\Delta_{\mathbf{y}\mathbf{x}\mathbf{x}'})_j = E(Y_j | do(\mathbf{x})) - E(Y_j | do(\mathbf{x}')).$$

Note that in the non-parametric case, it is not possible to compactly represent the causal effect of \mathbf{X} on \mathbf{Y} in a $(k_y \times k_x)$ -matrix. RS20 considered the class of regular asymptotically linear estimators for the non-parametric estimation of $\Delta_{\mathbf{y}\mathbf{x}\mathbf{x}'}$. This class includes inverse probability weighting by a non-parametrically estimated propensity score (Hirano et al., 2003), non-parametric outcome regression (Hahn, 1998), and double machine learning (Chernozhukov et al., 2018; Smucler et al., 2019). We use $\hat{\Delta}_{\mathbf{y}\mathbf{x}\mathbf{x}',\mathbf{z}}$ to denote an estimator from this class that estimates $\Delta_{\mathbf{y}\mathbf{x}\mathbf{x}'}$ adjusting for a valid adjustment set \mathbf{Z} . Under a causal DAG model and certain smoothness and complexity restrictions, $\hat{\Delta}_{\mathbf{y}\mathbf{x}\mathbf{x}',\mathbf{z}}$ is a consistent estimator of $\Delta_{\mathbf{y}\mathbf{x}\mathbf{x}'}$. For given \mathbf{y} , \mathbf{x} and \mathbf{x}' , the asymptotic distribution of estimators from this class depends only on \mathbf{Z} , therefore we do not further distinguish between the estimators. We denote the asymptotic variance of $(\hat{\Delta}_{\mathbf{y}\mathbf{x}\mathbf{x}',\mathbf{z}})_j = \hat{\Delta}_{y_j \mathbf{x}\mathbf{x}',\mathbf{z}}$ by $a.var(\hat{\Delta}_{y_j \mathbf{x}\mathbf{x}',\mathbf{z}})$. See RS20 and the references therein for more details on regular asymptotically linear estimators.

Definition 1 (O-set; HPM19 Definition 3.8) Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a DAG, CPDAG or *max*PDAG \mathcal{G} . Then $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$ is defined as:

$$\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \text{pa}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{G}), \mathcal{G}) \setminus \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}).$$

An example is given in Figure 1. It shows the causal relations between 12 symptoms of prodromal schizophrenia as measured by the *Schizotypic Syndrome Questionnaire* (van Kampen, 2006). The DAG was constructed using a combination of expert knowledge and data-driven structure learning (van Kampen, 2014). For illustration, we here take this given DAG as ground truth. Suppose we are interested in the causal effect of Alienation (ALN) on Delusional Thinking (DET). The bold edges indicate the causal paths with causal nodes {PER, SUS, FTW, DET} (circles). The parents of the causal nodes are {ALN, PER, SUS, FTW, AIS, CDR}, the forbidden set is {ALN, PER, SUS, FTW, DET, HOS, EGC} and the \mathbf{O} -set is {ALN, PER, SUS, FTW, AIS, CDR} \setminus {ALN, PER, SUS, FTW, DET, HOS, EGC} = {AIS, CDR} (shown in boxes). Other valid adjustment sets are, for example, {AFF, SAN}, {AIS, CDR, AFF} and {AFF, APA, AIS, CDR, SAN}. This can be checked using the generalised adjustment criterion stated above.

Note that in many applications it might be possible to augment a causal graph e.g. with further parents of Y that are marginally independent of all other non-descendants of Y . This induces a different \mathbf{O} -set illustrating that this set depends on what variables are included in the graph. Note also that the \mathbf{O} -set is defined even if no valid adjustment set exists, but this case will rarely be of interest.




Figure 1: DAG from van Kampen (2014) illustrating the assumed causal relations between 12 prodromal symptoms of schizophrenia: AFF=Affective Flattening, AIS=Active Isolation, ALN=Alienation, APA=Apathy, CDR=Cognitive Derailment, DET=Delusional Thinking, EGC=Egocentrism, FTW=Living in a Fantasy World, HOS=Hostility, PER=Perceptual Aberrations, SAN=Social Anxiety, SUS=Suspiciousness. We are interested in the causal effect of ALN on DET, both shown in grey circles. Bold arrows show the causal paths from ALN to DET. The forbidden nodes are shown as circles, nodes in the \mathbf{O} -set are shown as boxes.

Proposition 2 (HPM19 Theorem 3.10 (1)) *Let \mathbf{X} and \mathbf{Y} be disjoint subsets of the node set \mathbf{V} of a causal DAG, CPDAG or \max PDAG \mathcal{G} . The set $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$ is a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} if (i) $\mathbf{Y} \subseteq \text{possde}(\mathbf{X}, \mathcal{G})$ and (ii) a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} exists.*

Condition (i) can be checked using a simple query on \mathcal{G} . If $\mathbf{Y} \not\subseteq \text{possde}(\mathbf{X}, \mathcal{G})$, we know that the causal effect of \mathbf{X} on $\mathbf{Y} \setminus \text{possde}(\mathbf{X}, \mathcal{G})$ is zero. Hence, without loss of generality, we can consider the set of outcome variables $\mathbf{Y} \cap \text{possde}(\mathbf{X}, \mathcal{G})$ instead of \mathbf{Y} . Condition (ii) is satisfied if $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$ or any other subset of $\mathbf{V} \setminus \{\mathbf{X}, \mathbf{Y}\}$ fulfils the generalised adjustment criterion stated above. For the DAG in Figure 1, it can easily be seen that $\text{DET} \in \text{de}(\text{ALN})$, hence condition (i) is satisfied. Under condition (i), condition (ii) is always satisfied for univariate treatment and outcome in a DAG, because the parents of treatment then form a valid adjustment set (see Pearl, 2009, p. 72f.).

The following proposition, which builds on earlier work by Kuroki and Miyakawa (2003) and Kuroki and Cai (2004), establishes the optimality of the \mathbf{O} -set in terms of the asymptotic variance in the linear and in the non-parametric setting.

Proposition 3 *Let \mathbf{X} and \mathbf{Y} be disjoint subsets of the node set \mathbf{V} of a causal DAG, CPDAG or maxPDAG \mathcal{G} , such that $\mathbf{Y} \subseteq \text{possde}(\mathbf{X}, \mathcal{G})$. Let \mathbf{Z} be a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} and let $\mathbf{O} = \mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$.*

- (a) *(HPM19 Theorem 3.10 (2)) If the variables \mathbf{V} follow a linear causal model compatible with \mathcal{G} , then, for every $X_i \in \mathbf{X}$ and $Y_j \in \mathbf{Y}$, $a.\text{var}(\hat{\beta}_{y_j x_i, \mathbf{x}_{-\mathbf{i}} \mathbf{o}}) \leq a.\text{var}(\hat{\beta}_{y_j x_i, \mathbf{x}_{-\mathbf{i}} \mathbf{z}})$.*
- (b) *(RS20 Theorem 2) For every $Y_j \in \mathbf{Y}$ and pair of vectors $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, $a.\text{var}(\hat{\Delta}_{y_j \mathbf{x} \mathbf{x}' \mathbf{o}}) \leq a.\text{var}(\hat{\Delta}_{y_j \mathbf{x} \mathbf{x}' \mathbf{z}})$.*

In other words, for a given causal linear model, the \mathbf{O} -set yields the smallest asymptotic variance for the OLS estimator among all valid adjustment sets. If linearity cannot be assumed, the \mathbf{O} -set yields the smallest variance for regular asymptotically linear estimators. Thus, assume that Figure 1 represents a causal linear model. Proposition 3 then implies that if we estimate the effect of ALN on DET by linearly regressing DET on ALN and the \mathbf{O} -set $\{\text{AIS}, \text{CDR}\}$, then the estimator will have a smaller asymptotic variance than if we regress DET on ALN and a different valid adjustment set, say the parent set of ALN, which equals $\{\text{AFF}, \text{SAN}\}$. Moreover, when we relax linearity, non-parametric adjustment for the \mathbf{O} -set $\{\text{AIS}, \text{CDR}\}$ is more efficient than non-parametric adjustment for any other valid adjustment set, provided the estimator is in the class of regular asymptotically linear estimators.

3. The \mathbf{O} -Set via Forbidden Projection

In this section we provide an alternative, intuitive construction of the \mathbf{O} -set. For the sake of clarity, we restrict ourselves to DAGs; generalisations to amenable maxPDAGs are given in Appendix C.

To motivate our alternative construction, we posit that a useful adjustment set should be i) valid, ii) easy to compute, and iii) efficient. Consider singleton treatment X and outcome Y , where the latter is not an ancestor of X . The parents of X are easy to determine and guaranteed to be valid (see Pearl, 2009, p. 72f.). However, it is well-known that adjusting for variables strongly associated with treatment tends to reduce the efficiency of OLS and other estimators of the treatment effect. Hence, adjusting for the parents of treatment is typically inefficient compared to other valid adjustment sets. In contrast, it is also well-known that regression adjustment for variables strongly associated with the outcome tends to improve the efficiency of OLS and other estimators. Hence, the parents of the outcome would appear a natural, easy to determine and more efficient alternative for adjustment. However, the parents of Y are not guaranteed to be a valid adjustment set; they may contain forbidden nodes, specifically mediators between treatment and outcome. For example, in Figure 1, FTW is a parent of the outcome DET, but a descendant of the treatment ALN and hence cannot be used for adjustment. Simply omitting such nodes from the parents of Y does not generally lead to a valid adjustment set either. For example, CDR alone does

not form a valid adjustment set in Figure 1, since there are open confounding paths, e.g. $ALN \leftarrow SAN \rightarrow AIS \rightarrow SUS \rightarrow FTW \rightarrow DET$.

Nonetheless, the intuition of using the parents of Y is correct if applied to a modified graph. As we show below, marginalising out, i.e. projecting over, the forbidden nodes results in a graph where the parent set of Y indeed coincides with the \mathbf{O} -set, and is thus guaranteed to yield an estimator with minimal asymptotic variance in the settings we consider, see Proposition 3. This characterization of the \mathbf{O} -set thus combines validity, graphical simplicity and efficiency. We will now explain this formally.

Consider again the case of a DAG \mathcal{D} containing sets \mathbf{X} and \mathbf{Y} . We first need the concept of latent projection, used to *marginalise* or *collapse over* latent, i.e. unobserved nodes, while preserving the remaining causal relations and (in)dependencies between the observed nodes.

Definition 4 (Latent projection; Verma and Pearl, 1990; Shpitser et al., 2014)

Let \mathcal{D} be a DAG with node set $\mathbf{W} \cup \mathbf{L}$ and $\mathbf{W} \cap \mathbf{L} = \emptyset$. The latent projection $\mathcal{D}(\mathbf{W})$ over \mathbf{L} on \mathbf{W} is a graph with node set \mathbf{W} and edges as follows: For distinct nodes $W_i, W_j \in \mathbf{W}$,

1. $\mathcal{D}(\mathbf{W})$ contains a directed edge $W_i \rightarrow W_j$ if and only if \mathcal{D} contains a directed path $W_i \rightarrow \dots \rightarrow W_j$ on which all non-endpoint nodes are in \mathbf{L} ,
2. $\mathcal{D}(\mathbf{W})$ contains a bi-directed edge $W_i \leftrightarrow W_j$ if and only if \mathcal{D} contains a path, with at least one non-endpoint node, of the form $W_i \leftarrow \dots \rightarrow W_j$ on which all non-endpoint nodes are non-colliders and in \mathbf{L} .

In the latent projection $\mathcal{D}(\mathbf{W})$, two nodes may be connected by a directed and a bi-directed edge at the same time. (In)dependence relations can be read off from a latent projection using the m -separation criterion (Richardson, 2003). For disjoint $\mathbf{A}, \mathbf{B}, \mathbf{C} \subset \mathbf{W}$, \mathbf{A} and \mathbf{B} are d -separated given \mathbf{C} in \mathcal{D} if and only if \mathbf{A} and \mathbf{B} are m -separated given \mathbf{C} in $\mathcal{D}(\mathbf{W})$ (Richardson et al., 2017).

For our definition of the \mathbf{O} -set, we project over the forbidden nodes, save \mathbf{X} and \mathbf{Y} , which motivates the following definition:

Definition 5 (Forbidden projection) Let \mathcal{D} be a DAG with node set \mathbf{V} and let \mathbf{X} and \mathbf{Y} be disjoint subsets of \mathbf{V} . We call the graph $\mathcal{D}^{\mathbf{X}\mathbf{Y}} = \mathcal{D}((\mathbf{V} \setminus \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})) \cup \mathbf{X} \cup \mathbf{Y})$ the forbidden projection of \mathcal{D} with respect to (\mathbf{X}, \mathbf{Y}) .

Figures 2 and 3 show some examples, where the forbidden nodes are shown as circles. In panels **A** and **B** of Figure 2, the forbidden sets only contain nodes in $\mathbf{X} \cup \mathbf{Y}$, hence nothing is projected over. Panels **E** and **F** show DAGs where the forbidden projection has bi-directed edges, which will become relevant in Proposition 6.

While we primarily introduce the forbidden projection to provide an alternative characterisation of the \mathbf{O} -set, it is a useful tool in its own right. In particular, as we show next, the forbidden projection of a causal DAG preserves all information relevant to the estimation of a causal effect via adjustment. All proofs are given in Appendix B and generalised to maxPDAGs in Appendix C.

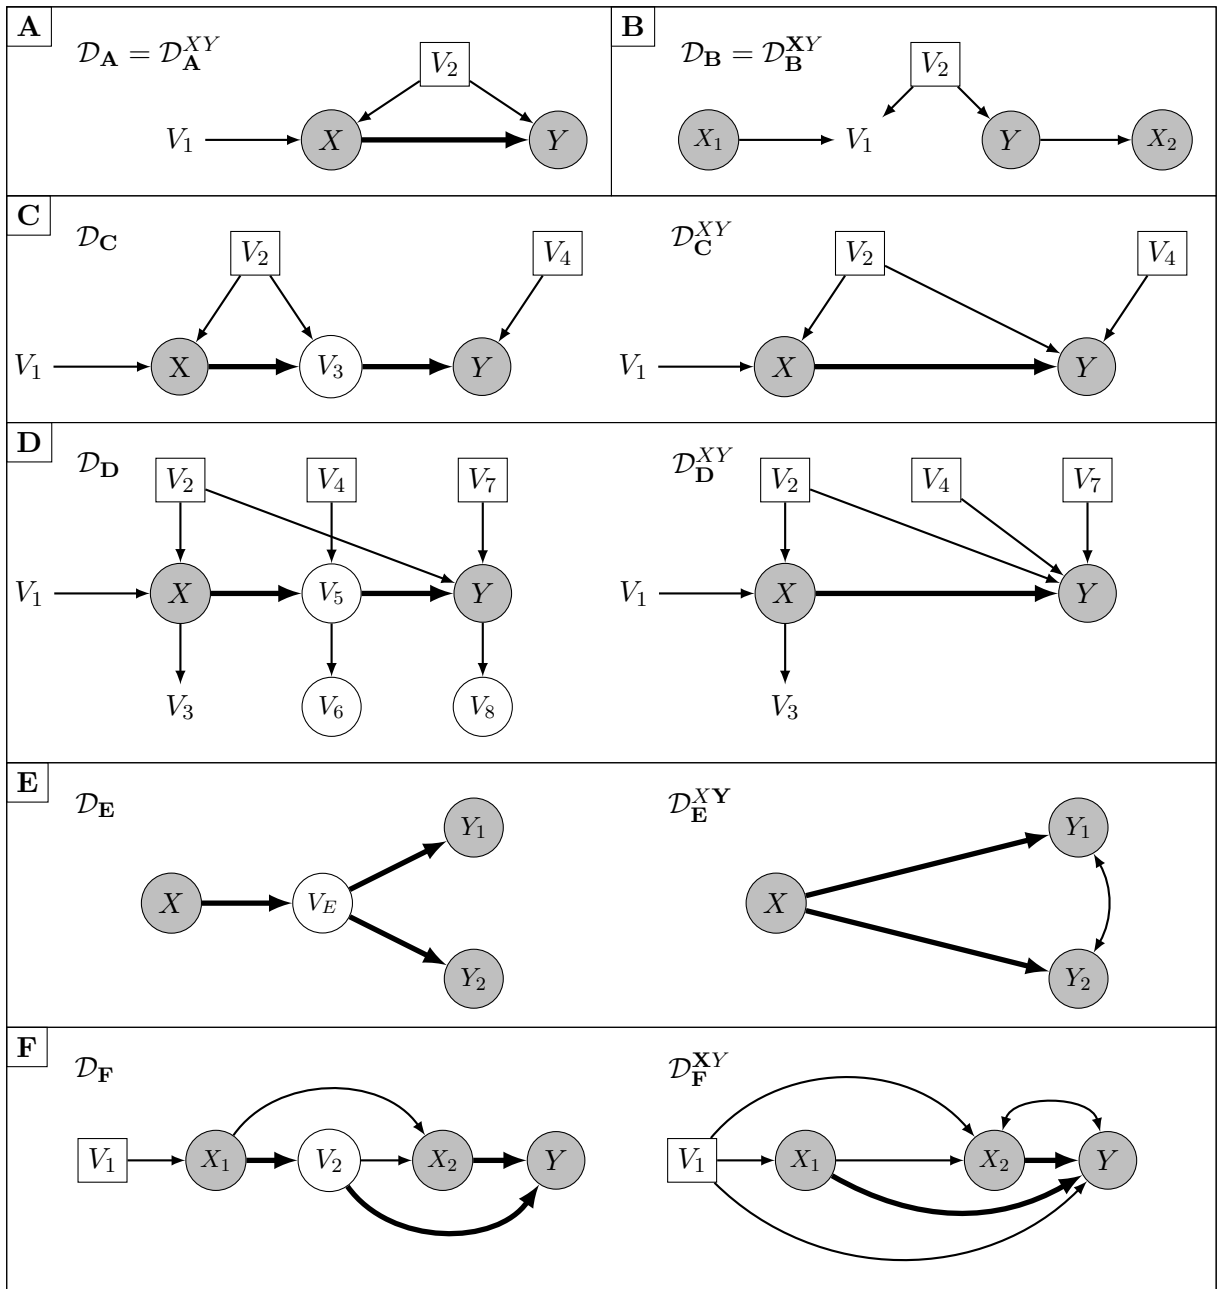


Figure 2: Example DAGs with their forbidden projections. The forbidden nodes are shown as circles, nodes in the \mathbf{O} -set are shown as boxes. The bold arrows show the causal paths from \mathbf{X} to \mathbf{Y} . In panels **A** and **B**, the original DAGs and their forbidden projections are identical. In panel **E**, the empty set is a valid adjustment set and also the \mathbf{O} -set. In panel **F**, the bi-directed edge between X_2 and Y indicates that the effect of $\mathbf{X} = \{X_1, X_2\}$ on Y is not identified via adjustment.

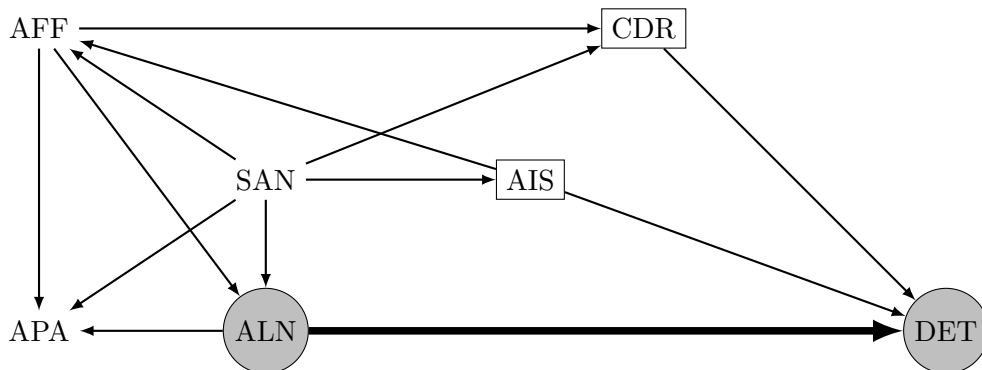


Figure 3: Forbidden projection of the DAG in Figure 1 with respect to $\mathbf{X} = \{\text{ALN}\}$ and $\mathbf{Y} = \{\text{DET}\}$. The forbidden nodes are shown as circles, nodes in the \mathbf{O} -set (parents of DET) are shown as boxes. The bold arrow shows the causal path from ALN to DET.

First, the forbidden projection can be used to check whether a valid adjustment set exists relative to given sets of nodes \mathbf{X} and \mathbf{Y} :

Proposition 6 *Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a causal DAG \mathcal{D} such that $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$. Then a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} exists if and only if there is no bi-directed edge between any $X \in \mathbf{X}$ and $Y \in \mathbf{Y}$ in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.*

In Figure 2, valid adjustment sets with respect to \mathbf{X} and \mathbf{Y} exist in all panels except for panel **F**. The effect of $\mathbf{X} = \{X_1, X_2\}$ on Y in panel **F** is, however, identified e.g. by the more general G-formula (Robins, 1986; Dawid and Didelez, 2010), the algorithm in Tian and Pearl (2003), or the methods in Nandy et al. (2017). See Guo and Perković (2020) and RS20 for results on efficient adjustment in the linear and non-parametric case, respectively. The bi-directed edge between Y_1 and Y_2 in panel **E** has no relevance in defining or determining a valid adjustment set.

For singleton Y such that a valid adjustment set with respect to (\mathbf{X}, Y) exists, the forbidden projection is particularly easy to interpret, as it is itself a causal DAG.

Proposition 7 *Let \mathbf{X} and $\{Y\}$ be disjoint node sets in a causal DAG \mathcal{D} such that $Y \in \text{de}(\mathbf{X}, \mathcal{D})$. Then $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ is a causal DAG if and only if there exists a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} .*

Further, an adjustment set that is valid in the original graph is also valid in the forbidden projection and vice versa:

Proposition 8 *Let \mathbf{X} , \mathbf{Y} and \mathbf{Z} be disjoint node sets in a causal DAG \mathcal{D} . Then \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} if and only if \mathbf{Z} is also a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.*

Using the forbidden projection, we now define the \mathbf{O}^* -set and prove that it is equal to the \mathbf{O} -set.

Definition 9 (O*-set) Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a DAG \mathcal{D} . We define $\mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D})$ as:

$$\mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \text{pa}(\mathbf{Y}, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) \setminus (\mathbf{X} \cup \mathbf{Y}).$$

In words, the \mathbf{O}^* -set is the set of parents of \mathbf{Y} in the forbidden projection $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$, excluding treatment nodes and outcome nodes. The next proposition states our key result.

Proposition 10 Let \mathbf{X} and \mathbf{Y} be disjoint subsets of the node set \mathbf{V} of a DAG \mathcal{D} such that $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$. Then $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D})$.

It now follows trivially that the statements about the \mathbf{O} -set in Proposition 3 are true for the \mathbf{O}^* -set as well.

Again, $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$ in Proposition 10 is not a severe restriction, because if $\mathbf{Y} \not\subseteq \text{de}(\mathbf{X}, \mathcal{D})$, we can instead consider the effect on $\mathbf{Y} \cap \text{de}(\mathbf{X}, \mathcal{D})$, as we know that the effect on $\mathbf{Y} \setminus \text{de}(\mathbf{X}, \mathcal{D})$ is zero.

Figure 3 shows the forbidden projection with respect to ALN and DET of the DAG in Figure 1. The \mathbf{O} -set $\{\text{AIS}, \text{CDR}\}$ (in boxes) is the parent set of DET. All other valid adjustment sets are less efficient, for example the afore-mentioned sets $\{\text{AFF}, \text{SAN}\}$, $\{\text{AIS}, \text{CDR}, \text{AFF}\}$ and $\{\text{AFF}, \text{APA}, \text{AIS}, \text{CDR}, \text{SAN}\}$. Due to Proposition 8, the validity of all of these sets can be confirmed by using the generalised adjustment criterion stated above on either the original DAG (Figure 1) or its forbidden projection (Figure 3). See Figure 2 for further examples.

To summarise, the forbidden projection can be used as follows: First, check in the original graph if $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$. Next, construct the forbidden projection $\mathcal{G}^{\mathbf{X}\mathbf{Y}}$ and check for bi-directed edges. If there is a bi-directed edge between a node in \mathbf{X} and a node in \mathbf{Y} , then the causal effect of interest is not identified via adjustment (Proposition 6). Else, $\mathcal{G}^{\mathbf{X}\mathbf{Y}}$ contains all information necessary to determine a valid adjustment set (Proposition 8), and in particular the \mathbf{O} -set, which is then the set of parents of \mathbf{Y} (Definition 9, Proposition 10). If \mathbf{Y} contains only one node, then $\mathcal{G}^{\mathbf{X}\mathbf{Y}}$ is a causal DAG itself and hence straightforward to interpret (Proposition 7).

4. Optimal Adjustment in the IDA Algorithm

In Sections 2 and 3, we considered optimal adjustment in DAGs, and in Appendix C we generalised the results to amenable maxPDAGs, which include amenable CPDAGs. As a reminder, a maxPDAG is said to be amenable relative to (\mathbf{X}, \mathbf{Y}) if every proper possibly causal path from \mathbf{X} to \mathbf{Y} starts with a directed edge out of \mathbf{X} . In this section, we consider non-amenable CPDAGs and maxPDAGs.

CPDAGs and maxPDAGs are of interest because they are the output of popular causal search algorithms, i.e. algorithms that attempt to learn a graph from data. Under the linear model with Gaussian error terms, which we focus on in this section, it is generally not possible to learn a unique DAG. Even under the additional assumptions of causal sufficiency and faithfulness (see e.g. Spirtes et al., 2000), one can at best learn a Markov equivalence class of DAGs, uniquely represented by a CPDAG (see e.g. Andersson et al., 1997). Given additional knowledge of some causal relationships between variables, access to interventional data, or other model restrictions, one can obtain a refinement of this class,

uniquely represented by a maxPDAG (Meek, 1995; Perković et al., 2017). For a CPDAG or maxPDAG \mathcal{G} , we use $[\mathcal{G}]$ to denote the set of DAGs that it represents. The interpretation of edges in a CPDAG or maxPDAG \mathcal{G} is as follows: A directed edge $A \rightarrow B$ means that this edge is present in all DAGs in $[\mathcal{G}]$. An undirected edge $A - B$ means that A and B are adjacent in every DAG in $[\mathcal{G}]$ and there is at least one DAG in $[\mathcal{G}]$ with $A \rightarrow B$ and at least one with $A \leftarrow B$.

We suppose in this section that we are interested in a univariate exposure X and a univariate outcome Y . For a given CPDAG or maxPDAG \mathcal{G} , the true causal effect of X on Y may differ across the DAGs in $[\mathcal{G}]$. In particular, Perković (2020) (Proposition 4.2) showed that assuming $Y \notin \text{pa}(X, \mathcal{G})$, the true causal effect of X on Y differs across DAGs in $[\mathcal{G}]$ if and only if \mathcal{G} is non-amenable relative to (X, Y) , i.e. there is a possibly causal path from X to Y that starts with an undirected edge. Hence, when \mathcal{G} is non-amenable relative to (X, Y) , we can at best determine a multiset of possible causal effects $(\tau_{yx}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}$, one for each DAG in $[\mathcal{G}]$. (A multiset $(\tau_{yx}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}$ may contain the same entry multiple times, e.g. if $[\mathcal{G}]$ contains five DAGs, of which three imply an effect of 0 and two imply an effect of 1.2, then $(\tau_{yx}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]} = \{0, 0, 0, 1.2, 1.2\}$.) While obviously less informative than a single number, this multiset of possible causal effects may still yield useful statistics. The minimum absolute value, for example, is a lower bound for the size of the causal effect. However, enumerating all DAGs in $[\mathcal{G}]$ is computationally very expensive even for moderately sized \mathcal{G} when there are many undirected edges.

Maathuis et al. (2009) proposed to reduce the complexity of this problem as follows. Consider two DAGs $\mathcal{D}, \mathcal{D}' \in [\mathcal{G}]$ such that $\text{pa}(X, \mathcal{D}) = \text{pa}(X, \mathcal{D}') = \mathbf{P}$ and $Y \notin \mathbf{P}$. As the parents of X form a valid adjustment set (Pearl, 2009, p. 72f.), $\tau_{yx}(\mathcal{D}) = \tau_{yx}(\mathcal{D}') = \tau_{yx}(\mathbf{P})$, where $\tau_{yx}(\mathbf{P})$ denotes the coefficient of X in the linear regression of Y on X and \mathbf{P} , i.e. $\beta_{yx.\mathbf{P}}$. Let $\mathbb{P} = \{\text{pa}(X, \mathcal{D}) \mid \mathcal{D} \in [\mathcal{G}]\}$ denote the set of all possible parent sets of X compatible with \mathcal{G} . Then $(\tau_{yx}(\mathbf{P}))_{\mathbf{P} \in \mathbb{P}}$ contains the same distinct values as $(\tau_{yx}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}$, while $|\mathbb{P}| \leq |[\mathcal{G}]|$. Maathuis et al. (2009) showed that it is possible to determine \mathbb{P} locally from the CPDAG \mathcal{G} without enumerating all DAGs in \mathcal{G} . They hence proposed a simple local procedure for calculating $(\hat{\tau}_{yx}(\mathbf{P}))_{\mathbf{P} \in \mathbb{P}}$, which is called ‘local IDA’ (local Intervention Calculus when the DAG is Absent). Perković et al. (2017) proposed a semi-local generalisation to maxPDAGs (‘semi-local IDA’).

The semi-local IDA algorithm for a maxPDAG is given in Algorithm 1. Let $\text{sib}(X, \mathcal{G})$ denote the set of nodes sharing an undirected edge with X in \mathcal{G} . Semi-local IDA loops over all subsets $\mathbf{S} \subseteq \text{sib}(X, \mathcal{G})$. It first constructs a graph \mathcal{G}' such that $\text{pa}(X, \mathcal{G}') = \mathbf{P} = \text{pa}(X, \mathcal{G}) \cup \mathbf{S}$. Here the complexity reduction becomes apparent: only the edges adjacent to X need to be oriented. To verify whether the added orientations are compatible with the original graph \mathcal{G} , the algorithm attempts to extend the graph to a maxPDAG by applying Meek’s orientation rules (ConstructMaxPDAG algorithm; Meek, 1995; Perković et al., 2017; see Figure 8 in Appendix A). This step is semi-local as edges not adjacent to X need to be oriented. If successful, $\hat{\beta}_{yx.\mathbf{P}}$ is added as a possible causal effect estimate, where $\mathbf{P} = \text{pa}(X, \mathcal{G}') = \mathbf{S} \cup \text{pa}(X, \mathcal{G})$.

Nandy et al. (2017) further generalised semi-local IDA to sets \mathbf{X} and \mathbf{Y} . However, this procedure does not use regression adjustment for possible causal effect estimation and is therefore not directly related to our results.

Algorithm 1 Local or semi-local IDA (Maathuis et al., 2009; Perković et al., 2017).
When the input is a CPDAG, line 7 can be simplified and the algorithm becomes fully local.

Require: CPDAG or maxPDAG \mathcal{G} with node set $\mathbf{V} = \{V_1, \dots, V_p, X, Y\}$, i.i.d. observations for V_1, \dots, V_p, X, Y

Ensure: multiset of estimates $\hat{\Theta}$

```

1:  $\hat{\Theta} \leftarrow \emptyset$ 
2:  $\text{sib}(X, \mathcal{G}) \leftarrow \{V \in \mathbf{V} : X - V \text{ in } \mathcal{G}\}$ 
3: for all  $\mathbf{S} \subseteq \text{sib}(X, \mathcal{G})$  do
4:   LocalBg  $\leftarrow \emptyset$ 
5:   for all  $S \in \mathbf{S}$ , add  $\{S \rightarrow X\}$  to LocalBg
6:   for all  $S \in \text{sib}(X, \mathcal{G}) \setminus \mathbf{S}$ , add  $\{S \leftarrow X\}$  to LocalBg
7:    $\mathcal{G}' \leftarrow \text{ConstructMaxPDAG}(\mathcal{G}, \text{LocalBg})$ 
8:   if  $\mathcal{G}' \neq \text{"FAIL"}$  then
9:     if  $Y \notin \text{pa}(X, \mathcal{G}')$  then
10:      regress  $Y$  on  $X \cup \text{pa}(X, \mathcal{G}')$  and add the estimated coefficient of  $X$  to  $\hat{\Theta}$ 
11:     else
12:       add 0 to  $\hat{\Theta}$ 
13:     end if
14:   end if
15: end for
16: return  $\hat{\Theta}$ 

```

Algorithm 2 Optimal IDA.

Require: CPDAG or maxPDAG \mathcal{G} with node set $\mathbf{V} = \{V_1, \dots, V_p, X, Y\}$, i.i.d. observations for V_1, \dots, V_p, X, Y

Ensure: multiset of estimates $\hat{\Theta}$

```

1:  $\hat{\Theta} \leftarrow \emptyset$ 
2:  $\text{sib}(X, \mathcal{G}) \leftarrow \{V \in \mathbf{V} : X - V \text{ in } \mathcal{G}\}$ 
3: for all  $\mathbf{S} \subseteq \text{sib}(X, \mathcal{G})$  do
4:   LocalBg  $\leftarrow \emptyset$ 
5:   for all  $S \in \mathbf{S}$ , add  $\{S \rightarrow X\}$  to LocalBg
6:   for all  $S \in \text{sib}(X, \mathcal{G}) \setminus \mathbf{S}$ , add  $\{S \leftarrow X\}$  to LocalBg
7:    $\mathcal{G}' \leftarrow \text{ConstructMaxPDAG}(\mathcal{G}, \text{LocalBg})$ 
8:   if  $\mathcal{G}' \neq \text{"FAIL"}$  then
9:     if  $Y \in \text{possde}(X, \mathcal{G}')$  then
10:      regress  $Y$  on  $X \cup \mathbf{O}(X, Y, \mathcal{G}')$  and add the estimated coefficient of  $X$  to  $\hat{\Theta}$ 
11:     else
12:       add 0 to  $\hat{\Theta}$ 
13:     end if
14:   end if
15: end for
16: return  $\hat{\Theta}$ 

```

4.1 Optimal IDA

HPM19 established that the parents of X , as used for adjustment by semi-local IDA, form one of the least efficient valid adjustment sets. It therefore seems a good idea to replace $\text{pa}(X, \mathcal{D})$ by the \mathbf{O} -set within the IDA algorithm to improve estimation precision. The key question is, however, whether the possible \mathbf{O} -sets can still be determined semi-locally. More formally, our aim is to estimate the multiset $(\tau_{yx}(\mathbf{O}))_{\mathbf{O} \in \mathbb{O}}$, $\mathbb{O} = \{\mathbf{O}(X, Y, \mathcal{D}) \mid \mathcal{D} \in [\mathcal{G}]\}$, where with a slight abuse of notation we define $\tau_{yx}(\mathbf{O}) = 0$ if $Y \notin \text{posde}(X, \mathcal{D})$. As before, for two DAGs \mathcal{D} and \mathcal{D}' with the same valid \mathbf{O} -set $\mathbf{O}(X, Y, \mathcal{D}) = \mathbf{O}(X, Y, \mathcal{D}') = \mathbf{O}$, we have $\tau_{yx}(\mathcal{D}) = \tau_{yx}(\mathcal{D}') = \tau_{yx}(\mathbf{O})$.

At first glance, it appears impossible to determine \mathbb{O} locally or semi-locally, as by Definitions 1 and 9 the causal nodes, their parents and the forbidden nodes, or the forbidden projection, are required to find the \mathbf{O} -set. However, it turns out that \mathbb{O} can be determined semi-locally almost in the same manner as \mathbb{P} . This is because once the directions of all edges involving X are given, i.e. for given \mathbf{P} , application of Meek’s rules reveals all descendants of X and, in consequence, all causal nodes, their parents and the forbidden nodes (cf. Lemma 18 in Appendix C). Hence, via Meek’s rules there exists a correspondence between possible parent sets and possible \mathbf{O} -sets. We therefore propose Algorithm 2, which we call optimal IDA. It is implemented in the R package `pcalg` (Kalisch et al., 2012, 2019).

Algorithm 2 does not specify whether $\mathbf{O}(X, Y, \mathcal{G}')$ is determined from \mathcal{G}' or from the forbidden projection. We expect this choice to be of limited relevance to the algorithm’s runtime. In our implementation, we determine $\mathbf{O}(X, Y, \mathcal{G}')$ directly from \mathcal{G}' . Note also that different possible parent sets can correspond to the same \mathbf{O} -set. Hence, optimal IDA could be modified to collect all sets in \mathbb{O} first, remove duplicates, and only then estimate regression coefficients.

In the following, we first state formally what can be said about the efficiency of the estimates output by optimal IDA, showing that it is worthwhile to replace the parents of X by the \mathbf{O} -set. Subsequently we compare the computational burden of the two algorithms.

Proposition 11 *Let X and Y be nodes in a causal CPDAG or maxPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, such that \mathbf{V} follows a causal linear model compatible with \mathcal{G} with Gaussian errors. Let $\widehat{\Theta}^{\mathbf{P}}$ and $\widehat{\Theta}^{\mathbf{O}}$ be the multisets returned by semi-local IDA and optimal IDA, respectively, applied to X, Y and \mathcal{G} , with the subsets of $\text{sib}(X, \mathcal{G})$ considered in the same order for both. Then, for $i \in \{1 \dots, k\}$, with $k = |\widehat{\Theta}^{\mathbf{P}}| = |\widehat{\Theta}^{\mathbf{O}}|$,*

1. $\mathbb{E}[\widehat{\Theta}_i^{\mathbf{P}}] = \mathbb{E}[\widehat{\Theta}_i^{\mathbf{O}}]$ and
2. $a.\text{var}(\widehat{\Theta}_i^{\mathbf{P}}) \geq a.\text{var}(\widehat{\Theta}_i^{\mathbf{O}})$.

The proof is given in Appendix D. Note that if we do not assume Gaussianity in Proposition 11, then $a.\text{var}(\widehat{\beta}_{yx.\mathbf{O}}) \leq a.\text{var}(\widehat{\beta}_{yx.\mathbf{Z}})$ can only be guaranteed if (i) \mathbf{Z} is a valid adjustment set in the true DAG, and (ii) \mathbf{O} is the \mathbf{O} -set of the true DAG. This is because in a causal linear model with non-Gaussian errors, a variable is only required to be linear in its parents, and is not necessarily linear given another node set (cf. Nandy et al., 2017). However, if we are willing to assume that all errors in the underlying causal model are non-Gaussian, alternative causal search approaches exist which output a DAG instead of an equivalence class, e.g. algorithms such as LiNGAM (Shimizu et al., 2006).

Remark 12 (1) *In terms of the computational burden, semi-local and optimal IDA are very similar for maxPDAGs. The key difference is that optimal IDA adjusts for the \mathbf{O} -set instead of the parent set of X (line 10), where the \mathbf{O} -set is straightforward to determine from \mathcal{G}' . However, optimal IDA crucially relies on the construction of the maxPDAG in line 7 to determine the \mathbf{O} -set, while in semi-local IDA this step can be replaced by a simple local query when the input is known to be a CPDAG. Hence, for the special case of a CPDAG, semi-local IDA can be made fully local by simplifying line 7, whereas optimal IDA cannot.*

(2) *A further minor difference between semi-local and optimal IDA is the \mathbf{if} -statement in line 9. Semi-local IDA only checks whether $Y \notin \text{pa}(X, \mathcal{G}')$, whereas optimal IDA checks the stronger condition $Y \in \text{possde}(X, \mathcal{G}')$. Both conditions ensure that the considered adjustment sets $\text{pa}(X, \mathcal{G}')$ and $\mathbf{O}(X, Y, \mathcal{G}')$, respectively, are valid adjustment sets. Moreover, if $Y \notin \text{possde}(X, \mathcal{G}')$, then $\tau_{yx}(\mathcal{D}) = 0$ for any $\mathcal{D} \in [\mathcal{G}']$. The 0 estimate of optimal IDA in this case is therefore the most efficient estimate. Alternatively, we could also insist on $Y \in \text{possde}(X, \mathcal{G}')$ in semi-local IDA and return 0 otherwise. As discussed in the appendix of Maathuis et al. (2009), this is only recommended if the input graph is thought to be reliable, but can lead to the amplification of errors if the input graph is not accurate.*

Remark 13 *Proposition 11 concerns the asymptotic variance when the true CPDAG or a true maxPDAG is given. When the graph is estimated on the same data as used for IDA, the naive standard errors from the adjusted linear regressions are invalid. Although considerable progress has been made in the area of post-selection inference (e.g. Berk et al., 2013; Belloni et al., 2014; Rinaldo et al., 2019), no method has been proposed specifically for estimating standard errors of causal effect estimates after causal search.*

It is straightforward to extend optimal IDA to situations where \mathbf{X} and \mathbf{Y} are sets. However, as noted earlier, in this case joint causal effect estimation via regression adjustment is not always possible. Optimal IDA will then not return an estimate. The estimation procedures used by joint IDA (Nandy et al., 2017) provide an alternative.

4.2 Illustration

We now illustrate optimal IDA (Algorithm 2) using a toy example. Consider the CPDAG \mathcal{G} shown in Figure 4(a) and suppose we are interested in the causal effect of X on Y . Clearly, \mathcal{G} is not amenable relative to (X, Y) and thus it is sensible to apply optimal IDA.

The set $\text{sib}(X, \mathcal{G})$ contains 3 nodes, hence there are 8 potential orientations of the undirected edges with endpoint X . From these 8, 3 imply new v-structures and are thus not compatible with \mathcal{G} . The other 5 can be extended to the maxPDAGs shown in Figure 4(b-f), where the bold arrows indicate orientations derived by Meek’s rules (see Figure 8 in Appendix A). For example, in 4(b) it follows from $V_1 \rightarrow X \rightarrow V_3$ that $V_1 \rightarrow V_3$ by Meek’s Rule 2. By Rule 1, it then follows that $V_3 \rightarrow V_5$. The compatibility check and the application of Meek’s rules are carried out in line 7 of optimal IDA.

Next, optimal IDA checks for each maxPDAG \mathcal{G}' , whether $Y \in \text{possde}(X, \mathcal{G}')$. Here, this is the case for all maxPDAGs except 4(c). For the other four graphs, $\mathbf{O} = \mathbf{O}(X, Y, \mathcal{G}')$ is determined and used to compute $\hat{\beta}_{yx, \mathbf{O}}$. We indicate $\mathbf{O}(X, Y, \mathcal{G}')$ by boxes in the Figures 4(b) and 4(d)-(f). For (c), an effect estimate of zero is returned.




Figure 4: A CPDAG \mathcal{G} (a) and the five maxPDAGs (b-f) corresponding to the five valid orientations of the neighbourhood of X . The bold edges have been obtained by applying Meek's rules. For each maxPDAG \mathcal{G}' , the boxes \square indicate $\mathbf{O}(X, Y, \mathcal{G}')$, while the diamonds \diamond indicate $\text{pa}(X, \mathcal{G})$. In (c), optimal IDA returns 0, as there is no possibly causal path from X to Y .

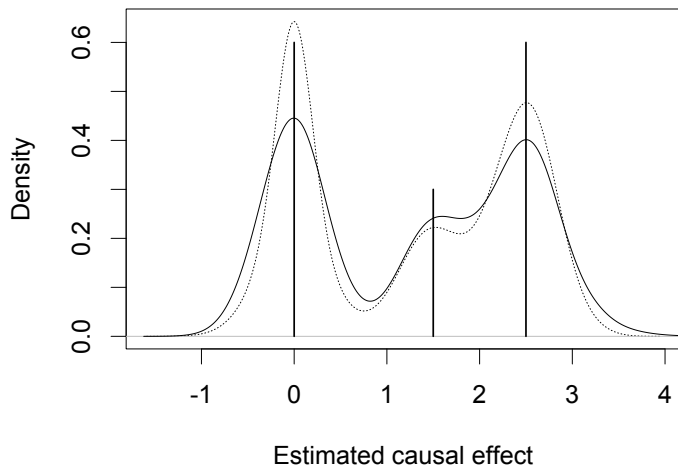


Figure 5: IDA density plot in the style of Maathuis et al. (2009). Shown are density curves for the estimated possible causal effects returned by local IDA (solid) and optimal IDA (dotted). The true possible causal effects are 0, 1.5 and 2.5 (vertical lines; height indicates relative frequency: 0 and 2.5 each occur in two of the five maxPDAGs in Figure 4).

For comparison, the diamonds in Figures 4(b-f) show the adjustment sets in local IDA (Algorithm 1), i.e. $\text{pa}(X, \mathcal{G}')$. In (b) and (e), $\text{pa}(X, \mathcal{G}') = \mathbf{O}(X, Y, \mathcal{G}')$. In (c), optimal IDA returns zero (Algorithm 2, line 12), while local IDA returns $\beta_{yx, \mathbf{P}}$ with $\mathbf{P} = \{V_1, V_3\}$, which converges to $\beta_{yx, \mathbf{P}} = 0$. The main advantage of optimal IDA becomes apparent in cases (d) and (f): In (d), $\mathbf{O}(X, Y, \mathcal{G}') = \emptyset$, whereas $\text{pa}(X, \mathcal{G}') = \{V_4\}$ which is guaranteed to reduce efficiency. In (f), $\text{pa}(X, \mathcal{G}') = \{V_3\}$ and $\mathbf{O}(X, Y, \mathcal{G}') = \{V_3, V_5\}$, where the latter improves efficiency.

For further illustration, we carried out a small simulation study in which we generated data according to a causal linear model compatible with Figure 4(b). 1 000 datasets with 40 observations each were generated and given as input to local IDA and optimal IDA, together with the CPDAG in Figure 4(a). The true possible causal effects are 0, 1.5 and 2.5, visualised as vertical lines in Figure 5. The plot shows smoothed density curves for the estimates returned by local IDA (solid) and optimal IDA (dotted). The density plot for optimal IDA is clearly narrower around the values 0 and 2.5. The difference between the algorithms is even more pronounced for graphs with more nodes and longer paths (not shown). The R-code (R Core Team, 2019) for reproducing Figure 5 is available in the Online Supplement.

4.3 Simulation

In order to compare the performance of optimal versus local IDA in finite sample settings, we carried out a more extensive simulation study. The design was chosen to reflect a typical situation where IDA is used, i.e. interest lies in the causal effect of X on Y in a (known or estimated) CPDAG \mathcal{G} that is non-amenable relative to (X, Y) . Non-amenable implies that the multiset $(\tau_{xy}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}$ of possible causal effects of X on Y compatible with \mathcal{G} contains more than one distinct value (for almost all parameters values of the causal linear model) (Perković, 2020, Proposition 4.2). A useful summary of $(\tau_{xy}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}$ is the minimum absolute value, $\min(\text{abs}((\tau_{xy}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}))$, because when this value is non-zero, we know that X has *some* effect on Y . The aim of our simulation study was to compare how well $\min(\text{abs}((\tau_{xy}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}))$ is estimated by optimal versus local IDA, in terms of the Monte-Carlo mean squared error (MSE).

We investigated 24 scenarios by considering all combinations of the following parameters: number of nodes $p \in \{10, 20, 50, 100\}$, expected number of neighbours per node $d \in \{2, 3, 4\}$, and sample size $n \in \{100, 1\,000\}$. In each scenario, the following was repeated 1 000 times (R code for reproducing the simulation study is available in the Online Supplement):

A DAG \mathcal{D} , with CPDAG \mathcal{G} , with p nodes and d expected neighbours per node was randomly chosen such that \mathcal{G} was non-amenable relative to two randomly chosen nodes (X, Y) and such that $\min(\text{abs}((\tau_{xy}(\mathcal{D}))_{\mathcal{D} \in [\mathcal{G}]}))$ was non-zero. (Note that the DAG with its unique ‘true’ causal effect was simulated for convenience only. Conceptually, we drew directly from the space of CPDAGs, which is why we consider the whole multiset of possible effects to be ‘the truth’.) The following was then repeated 100 times: A dataset with n observations was generated from a linear causal model on \mathcal{D} where the non-zero coefficients were randomly chosen from a uniform distribution on $[-1, -0.1] \cup [0.1, 1]$. Greedy equivalence search (Chickering, 2002) was applied to the data, yielding an estimated CPDAG \mathcal{G}^* . Optimal and local IDA were both applied to the true CPDAG \mathcal{G} and the estimated

CPDAG \mathcal{G}^* . The four output multisets of estimates were summarised by their minimum absolute values. These were compared on the basis of their Monte-Carlo MSE, i.e. the squared difference between the estimated minimum absolute value and the true minimum absolute value, averaged over the 100 repetitions. Specifically, we calculated the MSEs for the estimated minima using optimal IDA versus local IDA by computing the relative MSE (RMSE), $\text{MSE}(\text{optimal IDA})/\text{MSE}(\text{local IDA})$. This was done separately for \mathcal{G} and for \mathcal{G}^* , and denoted r and r^* , respectively. An RMSE of less than one indicates that optimal IDA is more precise than local IDA in estimating the minimum of the multiset of causal effects. In light of Remark 13, we did not consider estimated standard errors.

In addition to the above 24 scenarios, we investigated the relative performance of optimal IDA in a scenario where the graph is sparse ($d = 1$) and the sample size is moderate ($n = 100$). We considered eight settings where the number of nodes was between $p = 10$ and $p = 1000$. Such high-dimensional scenarios occur, for instance, with gene expression data. As greedy equivalence search is slow for large graphs, we reduced the number of replications from 1000 to 100 and the number of datasets per graph from 100 to 10.

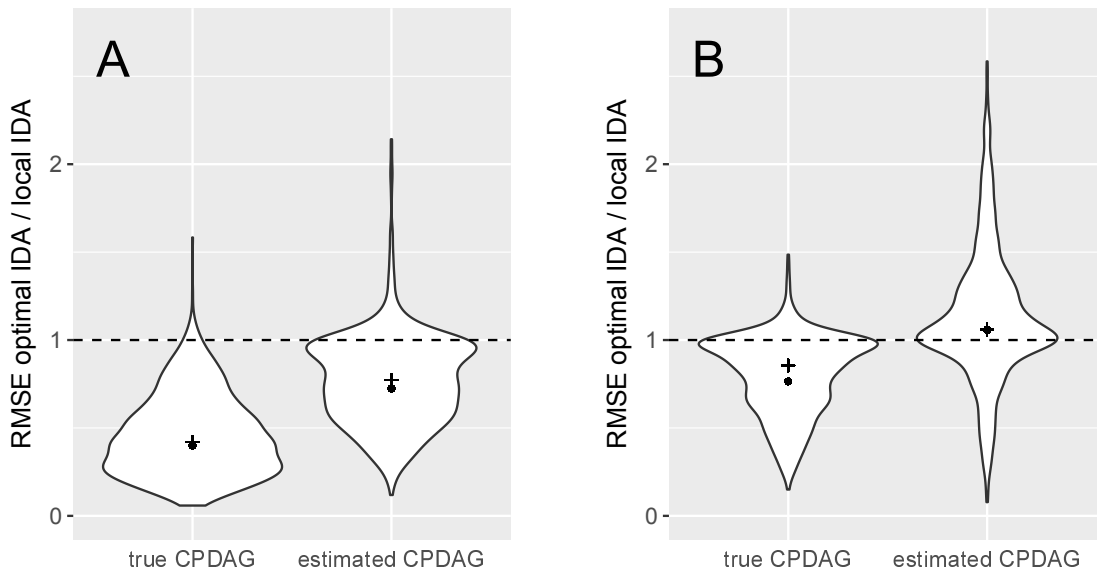


Figure 6: Violin plots of the relative mean squared errors (RMSEs) r and r^* for the true and estimated CPDAGs, respectively. Scenario **A**: $p = 100$ nodes, $d = 4$ expected neighbours per node, sample size $n = 1000$. Scenario **B**: $p = 10$, $d = 4$ and $n = 100$. The dots mark the geometric means, the plus signs the medians.

Figure 6 shows violin plots of the RMSEs r and r^* over the 1000 repetitions, together with the geometric mean and the median. Two scenarios are shown: The one where optimal IDA showed the best overall performance (scenario **A**, $p = 100$, $d = 4$, $n = 1000$), and the worst one (scenario **B**, $p = 10$, $d = 4$, $n = 100$) of all the simulation settings considered. The geometric means and medians for all scenarios are summarised in Tables 1 and 2; the

	$n = 100$			$n = 1000$		
	$d = 2$	$d = 3$	$d = 4$	$d = 2$	$d = 3$	$d = 4$
$p = 10$	0.70 (0.76)	0.72 (0.79)	0.76 (0.86)	0.69 (0.78)	0.71 (0.78)	0.75 (0.86)
$p = 20$	0.64 (0.69)	0.63 (0.68)	0.61 (0.66)	0.63 (0.68)	0.60 (0.65)	0.59 (0.65)
$p = 50$	0.60 (0.64)	0.54 (0.57)	0.51 (0.55)	0.55 (0.58)	0.50 (0.54)	0.46 (0.49)
$p = 100$	0.57 (0.61)	0.50 (0.52)	0.44 (0.46)	0.54 (0.58)	0.44 (0.46)	0.40 (0.42)

Table 1: Geometric means (in parentheses: medians) of the relative mean squared errors (RMSEs) r over 1 000 repetitions for scenarios with different numbers of nodes (p), expected number of neighbours per node (d), and sample sizes (n). Optimal and local IDA were applied to the true CPDAG \mathcal{G} .

	$n = 100$			$n = 1000$		
	$d = 2$	$d = 3$	$d = 4$	$d = 2$	$d = 3$	$d = 4$
$p = 10$	1.06 (1.01)	1.06 (1.04)	1.06 (1.06)	0.95 (0.99)	0.97 (1.00)	1.01 (1.00)
$p = 20$	0.99 (1.00)	0.99 (1.00)	0.96 (1.00)	0.88 (0.96)	0.89 (0.97)	0.94 (0.99)
$p = 50$	0.94 (0.98)	0.89 (0.93)	0.89 (0.93)	0.81 (0.90)	0.79 (0.85)	0.78 (0.86)
$p = 100$	0.97 (1.00)	0.94 (0.97)	0.90 (0.94)	0.81 (0.91)	0.73 (0.80)	0.72 (0.77)

Table 2: Geometric means (in parentheses: medians) of the relative mean squared errors (RMSEs) r^* over 1 000 repetitions for scenarios with different numbers of nodes (p), expected number of neighbours per node (d), and sample sizes (n). Optimal and local IDA were applied to the estimated CPDAG \mathcal{G}^* .

complete set of violin plots is shown in Appendix E. Optimal IDA clearly outperformed local IDA, in terms of the geometric mean and median of the RMSE, in all scenarios when applied to the true CPDAG. When the CPDAG was estimated using greedy equivalence search, optimal IDA was still superior in the majority of scenarios, but r^* was notably larger than r in all scenarios, i.e. the relative performance of optimal IDA was worse with an estimated CPDAG than with a known CPDAG. As an estimated graph inevitably contains some errors regarding the presence and direction of edges, this result may indicate that estimation adjusting for the \mathbf{O} -set suffers more from such errors than adjusting for the set of parents of X .

Small n and small p do not entail much advantage of using optimal IDA: In graphs with only a few nodes, the \mathbf{O} -set and the set of parents of X are often similar or even coincide, so that the gain in efficiency when using the \mathbf{O} -set is less pronounced. A smaller sample size leads to more errors in the estimated graph, which affects estimation of the \mathbf{O} -set more than estimation of the set of parents of X , as we conjectured above. However, optimal IDA seems to have a slight advantage for larger d when p is also larger.




Figure 7: Violin plots of the relative mean squared errors (RMSEs) r and r^* for the true and estimated CPDAGs, respectively. All graphs have $d = 1$ expected neighbour per node and graphs were estimated from $n = 100$ observations; the number of nodes p varies. The dots mark the geometric means, the plus signs the medians.

	$p = 10$	$p = 20$	$p = 50$	$p = 100$
true CPDAG	0.71 (0.77)	0.69 (0.76)	0.66 (0.69)	0.66 (0.71)
estimated CPDAG	1.04 (1.06)	1.04 (1.01)	0.99 (1.02)	0.93 (1.00)
	$p = 250$	$p = 500$	$p = 750$	$n = 1000$
true CPDAG	0.60 (0.68)	0.67 (0.69)	0.67 (0.77)	0.69 (0.75)
estimated CPDAG	0.73 (1.02)	1.16 (1.09)	0.94 (1.04)	0.79 (1.00)

Table 3: Geometric means (in parentheses: medians) of the relative mean squared errors (RMSEs) r and r^* for the true and estimated CPDAGs, respectively, over 100 repetitions for scenarios with different numbers of nodes (p), $d = 1$ expected neighbours per node, and sample size $n = 100$.

The additional results for the sparse graphs are shown in Figure 7 and Table 3. Optimal IDA outperformed local IDA regardless of the number of nodes p when the CPDAG was known. When the CPDAG was not known, the median RMSE was about 1 for all p . The geometric mean varied around 1 with no obvious pattern, which may be due to the small number of replications. The results suggest that in the high-dimensional setting, the primary difficulty is learning the graph, limiting the advantage of optimal IDA over local IDA.

In summary, based on the simulation results, we recommend using optimal IDA when there is high confidence in the estimated graph. The advantage over local IDA will be most pronounced when the number of nodes is at least 20, or better 50 or more.

5. The O-Set and Non-Graphical Variable Selection

We now assume that neither the causal DAG \mathcal{D} nor a CPDAG or maxPDAG is known to us, therefore we wish to select a valid adjustment set in a non-graphical manner. We restrict our discussion to the case where we have a univariate treatment X and outcome Y of interest.

In multiple regression analyses, it is common to apply variable selection procedures, e.g. backward selection, to find a set of relevant predictors for an outcome Y . In high-dimensional settings, regularisation methods that combine selection and estimation, such as the Lasso or the Elastic Net, are commonly used (Tibshirani, 1996; Zou and Hastie, 2005). While variable selection for prediction is in general a different task than finding an efficient or optimal adjustment set for causal effect estimation, we will discuss next under what assumptions and modifications these tasks coincide. For a general overview of the relation between variable and confounder selection see Witte and Didelez (2019). A basic assumption for the validity of a selected adjustment set is that the set \mathbf{Z} from which we select the variables must itself be a valid adjustment set, as defined in Section 2. A number of selection procedures can then be used to determine different types of valid adjustment sets as subsets of \mathbf{Z} , e.g. a minimal valid adjustment set (Witte and Didelez, 2019; de Luna et al., 2011).

Consider first Algorithm 3, which shows the template for backward regression selection (see e.g. Kleinbaum and Kupper, 1978; Montgomery et al., 2012) with the above basic assumption added at the outset. Under the linear model assumptions with Gaussian errors, $Y \perp\!\!\!\perp Z_i \mid (\mathbf{Z}'_{-i}, X)$ can be tested by comparing the models with regressors $\mathbf{Z}'_{-i} \cup \{X\}$ versus $\mathbf{Z}' \cup \{X\}$, using a t-test with null hypothesis $\beta_{y z_i . x \mathbf{z}'_{-i}} = 0$. ‘Pval’ in line 6 is a function that outputs the p-value of a test for the null hypothesis specified in the argument. The maximum p-value is compared in line 9 to a threshold α . For a given α , Algorithm 3 implements the classical ‘p-value method’ (see e.g. Greenland and Pearce, 2015). Denote by $F_{\chi^2_1}(\cdot)$ the distribution function of the χ^2 distribution with one degree of freedom. For a given sample size n , Algorithm 3 with $\alpha = 1 - F_{\chi^2_1}(2)$ or $\alpha = 1 - F_{\chi^2_1}(\log(n))$ is equivalent to backward selection using the AIC or BIC, respectively (e.g. Murtaugh, 2014; Derryberry et al., 2018; although the motivation for using them stems from frameworks other than independence testing, see Akaike, 1974 and Schwarz, 1978). Algorithm 3 can easily be adapted to work with a measure of conditional independence other than the p-value of the

t-test. For example, two non-parametric implementations of Algorithm 3 were proposed by Li et al. (2005).

If the true independence relations are known, Algorithm 3 can be condensed to its oracle version, Algorithm 4. Comparing p-values is then redundant, and every Z_i needs to be visited only once, as it follows from general properties of conditional independence that the ordering Z_1, Z_2, \dots, Z_p does not matter, provided the joint probability of all variables is strictly positive. Essentially, Algorithm 4 eliminates variables until only the ‘direct predictors’ of Y are left, i.e. those variables with non-zero coefficients in the oracle regression of Y on X and \mathbf{Z} .

Algorithm 4 is the non-graphical version of the pruning algorithm introduced in HPM19 which uses d-separation relationships to prune a valid adjustment set to a subset such that the resultant effect estimator has a smaller asymptotic variance. Assume now that an underlying graph exists. The following Proposition 14 formalises how the \mathbf{O} -set can be viewed as the target set of backward variable selection algorithms and follows from Proposition 3.6 of HPM19 and Theorem 1 in RS20.

Algorithm 3 Backward regression selection.

Require: i.i.d. observations for variables X , Y and \mathbf{Z} , such that \mathbf{Z} is a valid adjustment set relative to (X, Y)

```

1:  $\mathbf{Z}' \leftarrow \mathbf{Z}$ 
2:  $P_{\max} \leftarrow 1$ 
3: while  $P_{\max} > \alpha$  do
4:    $P_{\text{list}} \leftarrow$  empty list of length  $|\mathbf{Z}'|$ 
5:   for all  $i$  in 1 to  $|\mathbf{Z}'|$  do
6:      $P_{\text{list}}[i] \leftarrow P_{\text{val}}(Y \perp\!\!\!\perp Z_i \mid (X, \mathbf{Z}'_{-i}))$ 
7:   end for
8:    $P_{\max} \leftarrow \max(P_{\text{list}})$ 
9:   if  $P_{\max} > \alpha$  then
10:     $\mathbf{Z}' \leftarrow \mathbf{Z}'_{-\text{argmax}(P_{\text{list}})}$ 
11:   end if
12: end while
13: return  $\mathbf{Z}'$ 

```

Algorithm 4 Oracle backward regression selection.

Require: independence relations between variables X , Y and \mathbf{Z} , such that \mathbf{Z} is a valid adjustment set relative to (X, Y)

```

1:  $\mathbf{Z}' \leftarrow \mathbf{Z}$ 
2: for all  $i$  in 1 to  $|\mathbf{Z}'|$  do
3:   if  $Y \perp\!\!\!\perp Z_i \mid (X, \mathbf{Z}'_{-i})$  then
4:      $\mathbf{Z}' \leftarrow \mathbf{Z}'_{-i}$ 
5:   end if
6: end for
7: return  $\mathbf{Z}'$ 

```

Proposition 14 *Let X and Y be nodes in a causal DAG, CPDAG or maxPDAG \mathcal{G} with node set \mathbf{V} and let \mathbf{V} follow a causal model with a joint density faithful to \mathcal{G} . Let \mathbf{Z} be a valid adjustment set relative to (X, Y) in \mathcal{G} and let \mathbf{Z}' be the output of Algorithm 4 when applied to \mathbf{Z} .*

- (a) *\mathbf{Z}' is a valid adjustment set and does not depend on the order in which the variables in \mathbf{Z} are considered in Algorithm 4.*
- (b) *If $\mathbf{O}(X, Y, \mathcal{G}) \subseteq \mathbf{Z}$, then $\mathbf{Z}' = \mathbf{O}(X, Y, \mathcal{G})$.*
- (c) *If \mathbf{V} follows a causal linear model, then $a.\text{var}(\hat{\beta}_{yx.\mathbf{z}'}) \leq a.\text{var}(\hat{\beta}_{yx.\mathbf{z}})$.*
- (d) *For every pair of values $x, x' \in \mathcal{X}$, $a.\text{var}(\hat{\Delta}_{yxx'.\mathbf{z}'}) \leq a.\text{var}(\hat{\Delta}_{yxx'.\mathbf{z}})$.*

As mentioned earlier, Lasso estimation can also be regarded as variable selection, even though its original motivation and common usage mostly concerns prediction. Under specific assumptions, the Lasso asymptotically selects all and only all the ‘direct predictors’ of Y with probability 1 (Zhao and Yu, 2006; Lounici, 2008). Thus, although Lasso uses a different principle than backward selection, it follows from Proposition 14 that when the starting set \mathbf{Z} is a valid adjustment set and a superset of the \mathbf{O} -set, the \mathbf{O} -set can also be viewed as the target set of the Lasso.

We emphasise again that the \mathbf{O} -set cannot be determined in a purely data-driven way. Neither the assumption that \mathbf{Z} is valid nor $\mathbf{O}(X, Y, \mathcal{G}) \subseteq \mathbf{Z}$ can be verified empirically. Hence, prior causal knowledge is essential before any variable selection algorithm can be applied (Witte and Didelez, 2019). In contrast to (semi-)local or optimal IDA, however, selection of an adjustment set based on Algorithm 4 allows some latent structures, as long as the assumption that \mathbf{Z} is a valid adjustment set continues to hold. This may be of advantage when only a subset of the variables have been measured.

The guarantees of Proposition 14 for Algorithm 4 do not translate to the finite sample version Algorithm 3. Regression selection in finite samples is known to have several weaknesses (see e.g. Harrell, 2010). Some issues are that the output may only be a local optimum, and that valid post-selection inference is difficult (Leeb and Pötscher, 2008). There is, however, a growing literature on post-selection inference both in the context of OLS-based approaches (e.g. Berk et al., 2013; Rinaldo et al., 2019) and of Lasso-based approaches (e.g. Lockhart et al., 2014; Lee et al., 2016). For causal effect estimation in a non-graphical context, post-selection inference has been considered by Belloni et al. (2014), Duker and Vansteelandt (2020a), Duker and Vansteelandt (2020b), Chernozhukov et al. (2018) and others.

6. Conclusions

In this paper, we provided insight into the construction and properties of the \mathbf{O} -set introduced by HPM19. We showed that the \mathbf{O} -set equals the set of parents of \mathbf{Y} in the latent projection over the forbidden nodes (Proposition 10). This lends formal support to the intuition that adjusting for all direct causes of \mathbf{Y} minimises the residual variance and hence improves precision when estimating the causal effect of \mathbf{X} on \mathbf{Y} .

The forbidden projection is a useful tool in its own right when the aim is to estimate a causal effect via adjustment. It displays all variables of interest, while the forbidden variables, which must not be adjusted for, are marginalised out. The forbidden projection thus reduces the complexity of the causal graph while preserving all information relevant for choosing an adjustment set (see Propositions 8 and 7).

We further proposed a new modification of the IDA algorithm, called optimal IDA, which outputs multisets of estimates of possible causal effects by adjusting for the possible \mathbf{O} -sets. We showed that this increases estimation precision also in cases where the causal structure is a-priori unknown and needs to be estimated. Moreover, this extends the applicability of optimal adjustment to non-amenable CPDAGs/maxPDAGs. Optimal IDA has been implemented in the R package `pcaIlg`. While causal search methods in general have some well-known shortcomings, IDA has proved to be a valuable tool for instance for screening purposes in large datasets (Le et al., 2013; Engelmann et al., 2015; Luo et al., 2018). The ‘optimal’ version can further improve its performance.

Finally, we detailed the prerequisites and assumptions under which non-graphical algorithms for backward variable selection can be viewed as aiming at selecting the \mathbf{O} -set. Essentially, we need to assume that the set of variables to select from consists of all nodes in the forbidden projection, or a suitable subset thereof. The algorithm then determines the parents / direct causes of Y based on detected conditional independencies. If the input contains forbidden nodes, however, or lacks certain confounders, the algorithm might select an invalid adjustment set. To avoid the latter, sufficient prior knowledge on the set of variables corresponding to forbidden nodes is therefore a key prerequisite when automated variable selection is to be used for causal inference. While this prerequisite may not require full knowledge of the underlying causal DAG, it is important to recognise that such prior knowledge cannot be established in a purely data-driven way (Witte and Didelez, 2019).

Much research on variable selection in causal graphs has focussed on finding small or minimal adjustment sets (de Luna et al., 2011; Textor and Liškiewicz, 2011; Knüppel and Stang, 2010). Small adjustment sets are useful during study planning, for instance when data collection is expensive and costs are to be minimised. Moreover, they entail desirable statistical properties e.g. for matching estimators, because suitable matches are more easily found when matching on a few variables only. In general, the \mathbf{O} -set is not minimal, but instead entails optimality of causal effect estimation by regression adjustment in linear causal models and non-parametric settings. Simulation results further indicate that the optimality of the \mathbf{O} -set extends to other parametric settings and estimation methods, e.g. estimation of the marginal odds ratio via standardised logistic regression (Witte and Didelez, 2019). Combining the benefits of small and optimal adjustment sets, RS20 show that the optimal minimal set, i.e. the set among all minimal adjustment sets yielding the most precise estimation in the class of regular asymptotically linear estimators, must be a subset of the \mathbf{O} -set, underlining its relevance and importance.

We note that adjustment is only one of several possible ways of identifying causal effects. While adjusting for the \mathbf{O} -set is asymptotically more efficient than adjusting for any other valid adjustment set, it is possible that an even smaller asymptotic variance can be obtained by using an alternative identification strategy, e.g. the front-door strategy (Pearl, 2009). This is further investigated in RS20 and in Guo and Perković (2020).

Finally, we would like to discuss some avenues for future research. First, given the results by RS20, a natural question is whether a non-parametric version of optimal IDA is feasible. Those aspects of IDA that relate to finding different possible valid adjustment sets are obviously not limited to the causal linear model, and estimators such as in RS20 could also be employed for any given X , Y and adjustment set. The simplifications for graph search algorithms and IDA under linearity, however, are considerable. For instance, greedy equivalence search with a Gaussian score has been shown to be consistent (Chickering, 2002), and has the advantage of always returning a CPDAG. Non-parametric graph search algorithms exist, but often come with large computational burdens and/or a low power to detect edges (Shah and Peters, 2020; Ramsey, 2014). Further, under a causal linear model, the causal effect of X on Y is a single value, see Section 2; and the marginal and conditional causal effects for different valid adjustment sets are all identical. For the non-parametric case, instead, the causal effect of X on Y is an unspecified function, and issues of non-collapsibility might also come into play. Solving these conceptual hurdles for non-parametric optimal IDA remains an open question.

Second, we assumed throughout this paper that all variables are observed. HPM19 have shown that in the presence of hidden variables, an asymptotically optimal set may not exist. Smucler et al. (2020), however, gave a sufficient condition under which an optimal adjustment set exists when the underlying DAG includes hidden variables, and showed that an optimal minimal adjustment set always exists. A necessary and sufficient condition for the existence of an optimal adjustment set, however, has not yet been formulated.

Acknowledgments

We thank the reviewers and the editor for their valuable comments and suggestions. We gratefully acknowledge financial support by the German Research Foundation (DFG—Project DI 2372/1-1).

Appendix A. Terminology

The following terminology is used throughout this paper. It is consistent with, and extends HPM19 where needed.

Graphs. A graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ consists of a node set \mathbf{V} and a set of edges \mathbf{E} . We consider three types of edges: *directed* (\rightarrow), *bi-directed* (\leftrightarrow) and *undirected* ($-$). There can be more than one edge between a given pair of nodes. We only consider loop-free graphs, i.e. an edge between a node and itself is not allowed. A loop-free graph where there is at most one edge between a given pair of nodes is called a *simple graph*. Two nodes joined by at least one edge are called *endpoints of the edge* and *adjacent*. A directed edge $A \rightarrow B$ is said to be *out of* A and *into* B . A graph $\mathcal{G}' = (\mathbf{V}', \mathbf{E}')$ is the *induced subgraph* of $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ with respect to \mathbf{V}' if $\mathbf{V}' \subseteq \mathbf{V}$ and \mathbf{E}' includes all edges in \mathbf{E} that are between nodes in \mathbf{V}' .

Paths. A *path* is a sequence of nodes and edges $(V_0, e_1, V_1, \dots, e_K, V_K)$, $K \geq 1$, such that every node occurs only once and for $k = 1, \dots, K$, e_k has endpoints V_{k-1} and V_k . In a simple graph, the path $(V_0, e_1, V_1, \dots, e_K, V_K)$ can unambiguously be identified by the sequence of nodes (V_0, V_1, \dots, V_K) alone. V_0 and V_K are called *endpoints of the path* $(V_0, e_1, V_1, \dots, e_K, V_K)$ and the path is said to be *between* V_0 and V_K or *from* V_0 to V_K , irrespective of the direction of the edges. For sets of nodes \mathbf{A} and \mathbf{B} , a path is said to be *between* \mathbf{A} and \mathbf{B} or *from* \mathbf{A} to \mathbf{B} if its first node is in \mathbf{A} and the last node is in \mathbf{B} . A path from \mathbf{A} to \mathbf{B} is *proper* if only its first node V_0 is in \mathbf{A} . Let $p = (V_0, e_1, V_1, \dots, e_K, V_K)$ and $k = 1, \dots, K$. Then an edge $V_k \leftarrow V_{k+1}$ on p is said to point towards V_0, \dots, V_k , while an edge $V_k \rightarrow V_{k+1}$ on p is said to point towards V_{k+1}, \dots, V_K . A path is *directed* from V_0 to V_K if all edges in the sequence are directed and point towards V_K . A path p is *possibly directed* from V_0 to V_K if all edges on p are either directed or undirected and there are no i, j , $1 \leq i < j \leq K$, such that $V_i \leftarrow V_j$ (cf. Perković et al. (2017); this definition of a possibly directed path is non-standard as V_i and V_j are not necessarily adjacent nodes on the path, which is required for maxPDAGs later). We define the concatenation of two paths $p = (V_0, e_1, V_1, \dots, e_K, V_K)$ and $q = (V_K, e_{K+1}, V_{K+1}, \dots, e_{K+L}, V_{K+L})$ as $p \oplus q = (V_0, e_1, V_1, \dots, e_{K+L}, V_{K+L})$, where we require that the nodes V_0, \dots, V_{K+L} are distinct.

Ancestry. If there is a directed path from A to B , or if $A = B$, then A is an *ancestor* of B and B is a *descendant* of A . If there is a possibly directed path from A to B , or if $A = B$, then A is a *possible ancestor* of B and B is a *possible descendant* of A . If there is an edge $A \rightarrow B$, then A is a *parent* of B and B is a *child* of A . If there is an edge $A - B$, A and B are *siblings*. Note that in our terminology, a node is a (possible) ancestor and (possible) descendant of itself, but not a parent/child/sibling of itself. For a node V in a simple graph \mathcal{G} , we denote the set of all ancestors, possible ancestors, descendants, possible descendants, parents, children and siblings of V in \mathcal{G} as $\text{an}(V, \mathcal{G})$, $\text{possan}(V, \mathcal{G})$, $\text{de}(V, \mathcal{G})$, $\text{possde}(V, \mathcal{G})$, $\text{pa}(V, \mathcal{G})$, $\text{ch}(V, \mathcal{G})$, $\text{sib}(V, \mathcal{G})$, respectively. For a set of nodes \mathbf{W} , the set $\text{an}(\mathbf{W}, \mathcal{G})$ is defined as $\bigcup_{W \in \mathbf{W}} \text{an}(W, \mathcal{G})$, with analogous definitions for $\text{possan}(\mathbf{W}, \mathcal{G})$, $\text{de}(\mathbf{W}, \mathcal{G})$, $\text{possde}(\mathbf{W}, \mathcal{G})$, $\text{pa}(\mathbf{W}, \mathcal{G})$, $\text{ch}(\mathbf{W}, \mathcal{G})$ and $\text{sib}(\mathbf{W}, \mathcal{G})$.

Colliders, definite-status paths and v-structures. A non-endpoint node V is a *collider* on a path p if both edges adjoining V on p have arrowheads at V , i.e. $\rightarrow V \leftarrow$, $\leftrightarrow V \leftarrow$, $\rightarrow V \leftrightarrow$, $\leftrightarrow V \leftrightarrow$. A non-endpoint node V is a *non-collider* on a path p if at least one of the edges adjoining V on p is out of V , i.e. $\rightarrow V \rightarrow$, $-V \rightarrow$, $\leftrightarrow V \rightarrow$, $\leftarrow V \rightarrow$, $\leftarrow V \leftrightarrow$, $\leftarrow V -$, $\leftarrow V \leftarrow$, or if both edges adjoining V on p are undirected edges and the

two nodes adjacent to V on p are not adjacent to each other. A *definite-status path* is a path on which every non-endpoint is either a collider or a non-collider. In a DAG or an ADMG, all paths are of definite status. Three nodes A , B and C form a *v-structure* in a graph \mathcal{G} if $A \rightarrow B \leftarrow C$ is the induced subgraph \mathcal{G}' on $\{A, B, C\}$.

ADMGs, DAGs and PDAGs. A directed path from A to B , together with an edge $A \leftarrow B$ forms a *directed cycle*. A graph with only directed and bi-directed edges and without directed cycles is called an *acyclic directed mixed graph* (ADMG). A simple graph with only directed edges and without directed cycles is called a *directed acyclic graph* (DAG). A simple graph with only directed and undirected edges containing no directed cycles is called a *partially directed acyclic graph* (PDAG).

Blocking and separation. (Richardson, 2003; Maathuis and Colombo, 2015; Pearl, 2009) A definite-status path p in an ADMG or PDAG \mathcal{G} is *blocked* by a node set \mathbf{C} if (i) p contains a non-collider in \mathbf{C} or (ii) p contains a collider that is not in $\text{an}(\mathbf{C}, \mathcal{G})$. Otherwise the path p is *open* given \mathbf{C} . Node sets \mathbf{A} and \mathbf{B} are said to be *m-separated* given a set \mathbf{C} if every path between an $A \in \mathbf{A}$ and a $B \in \mathbf{B}$ is blocked by \mathbf{C} . We then write $\mathbf{A} \perp_{\mathcal{G}} \mathbf{B} \mid \mathbf{C}$. In DAGs, m-separation is called d-separation.

Markov equivalence and CPDAGs. (Andersson et al., 1997) The *(Markov) equivalence class* of a DAG \mathcal{D} is the set of DAGs that imply the same d-separation relationships as \mathcal{D} . These are all DAGs with the same adjacencies and v-structures Verma and Pearl (1990). Markov equivalence classes can be represented as *completed partially directed acyclic graphs* (CPDAGs), which are simple graphs with directed or undirected edges, without directed cycles and with certain restrictions regarding the patterns of edges that can occur. The equivalence class represented by a CPDAG \mathcal{G} is denoted by $[\mathcal{G}]$. A directed edge $A \rightarrow B$ in \mathcal{G} means that this edge is present in all DAGs in the equivalence class $[\mathcal{G}]$. An undirected edge $A - B$ in \mathcal{G} means that A and B are adjacent in every DAG in $[\mathcal{G}]$ and there is at least one DAG in $[\mathcal{G}]$ with $A \rightarrow B$ and at least one with $A \leftarrow B$.

Meek’s rules and maxPDAGs. (Perković et al., 2017) Certain subsets of equivalence classes of DAGs can be represented by maximally oriented PDAGs (maxPDAGs), which are PDAGs with edge orientations that are closed under the orientation rules in Figure 8 (*Meek’s rules*, Meek (1995)). The set of DAGs represented by a maxPDAG \mathcal{G} is denoted by $[\mathcal{G}]$. The edges in maxPDAGs have the same interpretation as in CPDAGs. DAGs and CPDAGs are special cases of maxPDAGs.



Figure 8: Meek’s orientation rules. Let \mathcal{G} be a simple graph with only directed and undirected edges and without directed cycles. If the graph on the left is an induced subgraph of \mathcal{G} , then orient the undirected edges in \mathcal{G} according to the graph on the right (Meek, 1995). The rules prevent directed cycles and new v-structures.

Partial topological ordering. Let \mathcal{D} be a DAG with node set \mathbf{V} and let $\mathbf{V}_1, \dots, \mathbf{V}_p$ be a partition of \mathbf{V} . Then $\mathbf{V}_1 < \dots < \mathbf{V}_p$ is a *partial topological ordering* of \mathbf{V} if for every $i > j$, there are no directed edges from \mathbf{V}_i to \mathbf{V}_j .

Independence and faithfulness. For sets of random variables \mathbf{X} , \mathbf{Y} and \mathbf{Z} , if \mathbf{X} and \mathbf{Y} are conditionally independent given \mathbf{Z} , we write $\mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z}$. A joint density $f(\mathbf{v})$ over a set of random variables \mathbf{V} is *Markov* with respect to a DAG \mathcal{D} with node set \mathbf{V} if for disjoint $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \subseteq \mathbf{V}$, $\mathbf{X} \perp_{\mathcal{D}} \mathbf{Y} \mid \mathbf{Z} \Rightarrow \mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z}$; the density $f(\mathbf{v})$ is *faithful* to \mathcal{D} if also $\mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z} \Rightarrow \mathbf{X} \perp_{\mathcal{D}} \mathbf{Y} \mid \mathbf{Z}$.

Causal DAGs, CPDAGs, maxPDAGs and ADMGs. Intuitively, a *causal DAG* is a DAG where an edge $A \rightarrow B$ means that A is a direct cause of B (relative to the variables included). This can be formalised using the intervention operator, denoted by $do(\cdot)$ in Pearl (2009). For random variables \mathbf{V} and $\mathbf{X} \subseteq \mathbf{V}$, the *post-intervention density* $f(\mathbf{v} \mid do(\mathbf{x}'))$ is the joint density of \mathbf{V} in a (hypothetical) experiment that fixes \mathbf{X} to \mathbf{x}' for everyone in the population by an external intervention. A joint density $f(\mathbf{v})$ is *compatible with a causal DAG* $\mathcal{D} = (\mathbf{V}, \mathbf{E})$ if for all $\mathbf{X} \subseteq \mathbf{V}$, the post-intervention density $f(\mathbf{v} \mid do(\mathbf{x}'))$ can be written as

$$f(\mathbf{v} \mid do(\mathbf{x}')) = \mathbf{1}(\mathbf{x} = \mathbf{x}') \prod_{V \in \mathbf{V} \setminus \mathbf{X}} f(v \mid \text{pa}(V, \mathcal{D})),$$

where $\mathbf{1}(\mathbf{x} = \mathbf{x}')$ is the indicator function that is 1 if $\mathbf{x} = \mathbf{x}'$ and 0 otherwise. This is known as the truncated factorisation formula (Spirtes et al., 2000; Pearl, 2009). A CPDAG or maxPDAG \mathcal{G} is called a *causal CPDAG* or *causal maxPDAG* if $[\mathcal{G}]$ contains a causal DAG. A *causal ADMG* is an ADMG that has been obtained by subjecting a causal DAG to a latent projection, see Definition 4.

(Possibly) causal nodes and forbidden nodes. See Section 2.

Valid adjustment sets and amenability. Let \mathbf{X} , \mathbf{Y} and \mathbf{Z} be disjoint sets of random variables, where \mathbf{Z} is possibly empty. Then \mathbf{Z} is a *valid adjustment set* relative to (\mathbf{X}, \mathbf{Y}) if we have

$$f(\mathbf{y} \mid do(\mathbf{x})) = \begin{cases} f(\mathbf{y} \mid \mathbf{x}) & \text{if } \mathbf{Z} = \emptyset, \\ \int_{\mathbf{z}} f(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) f(\mathbf{z}) d\mathbf{z} & \text{otherwise.} \end{cases} \quad (1)$$

Relative to a causal DAG, CPDAG, maxPDAG or ADMG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, a valid adjustment set is defined as follows: Let \mathbf{X} , \mathbf{Y} and \mathbf{Z} be disjoint subsets of \mathbf{V} , where \mathbf{Z} is possibly empty. Then \mathbf{Z} is a *valid adjustment set* relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} if equation (1) holds for every joint density $f(\mathbf{v})$ compatible with \mathcal{G} (Perković et al., 2018). Further, \mathcal{G} is said to be *amenable* for adjustment relative to (\mathbf{X}, \mathbf{Y}) if every proper possibly causal path from \mathbf{X} to \mathbf{Y} starts with a directed edge out of \mathbf{X} (Perković et al., 2018).

Generalised adjustment criterion. (Perković et al., 2017, 2018; Shpitser et al., 2010) Let \mathbf{X} , \mathbf{Y} and \mathbf{Z} be disjoint node sets in a causal DAG, CPDAG, maxPDAG or ADMG \mathcal{G} . Then \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} if and only if the following three conditions hold:

- (a) \mathcal{G} is amenable relative to (\mathbf{X}, \mathbf{Y}) ,
- (b) $\mathbf{Z} \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \emptyset$,
- (c) all proper non-causal definite-status paths from \mathbf{X} to \mathbf{Y} are blocked by \mathbf{Z} .

Causal linear model. Let \mathcal{D} be a causal DAG with node set $\mathbf{V} = (V_1, \dots, V_p)$. Then \mathbf{V} is said to follow a *causal linear model* compatible with \mathcal{D} if the distribution of each $V_i \in \mathbf{V}$ can be described by an equation of the form

$$V_i = \sum_{V_j \in \text{pa}(V_i, \mathcal{D})} \alpha_{ij} V_j + \epsilon_{v_i}$$

with $\alpha_{ij} \in \mathbb{R}$ and ϵ_{v_i} a random variable with mean 0 and finite variance such that $\epsilon_{v_1}, \dots, \epsilon_{v_p}$ are jointly independent. For a causal CPDAG or maxPDAG \mathcal{G} , \mathbf{V} is said to follow a causal linear model compatible with \mathcal{G} if \mathbf{V} follows a causal linear model compatible with a DAG in $[\mathcal{G}]$.

Partial variance notation. Consider a random variable S and a random vector \mathbf{T} . We denote the covariance matrix of \mathbf{T} by $\Sigma_{\mathbf{tt}}$ and the row vector of covariances between S and \mathbf{T} by $\Sigma_{\mathbf{st}}$. The partial variance of S given \mathbf{T} is defined as $\sigma_{ss, \mathbf{t}} = \text{Var}(S) - \Sigma_{\mathbf{st}} \Sigma_{\mathbf{tt}}^{-1} \Sigma_{\mathbf{st}}^{-1}$.

Asymptotic variance. Consider a sequence of estimators $(\hat{\beta}_n)_{n \in \mathbb{N}}$ such that $\sqrt{n}(\hat{\beta}_n - \beta)$ converges in distribution to $\mathcal{N}(0, v)$. We call v the asymptotic variance of $\hat{\beta}$ and write $a.\text{var}(\hat{\beta}) = v$.

Appendix B. Proofs for Section 3

In this appendix, we prove our claims about the forbidden projection made in Section 3.

Proposition 6 *Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a causal DAG \mathcal{D} such that $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$. Then a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} exists if and only if there is no bi-directed edge between any $X \in \mathbf{X}$ and $Y \in \mathbf{Y}$ in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.*

Proof We show that a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} cannot exist if and only if there is a bi-directed edge between a $X \in \mathbf{X}$ and a $Y \in \mathbf{Y}$ in the forbidden projection $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.

Assume first that there is a bi-directed edge in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ between some $X \in \mathbf{X}$ and a $Y \in \mathbf{Y}$. Then according to Definition 4 of the latent projection there is a path in \mathcal{D} between X and Y on which all nodes are non-colliders and contained in the forbidden set. This constitutes a non-causal path that cannot be blocked by any sets of nodes that are not forbidden. Hence no valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) and \mathcal{D} exists.

Assume now that there is no valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} . Then Lemma 15 implies $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D}), \mathcal{D}) \neq \emptyset$. Let $X^* \in \mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D}), \mathcal{D})$. Then there must exist a node $C^* \in \text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$ and a node $Y^* \in \mathbf{Y}$ such that there is a path of the form $X^* \leftarrow \dots \leftarrow C^* \rightarrow \dots \rightarrow Y^*$ where all non-endpoints are non-colliders on the path and in the forbidden set. It follows from Definition 4 of the latent projection that $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ contains a bi-directed edge $X^* \leftrightarrow Y^*$. ■

Lemma 15 (Corollary 27 in Perković et al., 2018) *Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a causal DAG \mathcal{D} such that $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$. Then a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} exists if and only if $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D}), \mathcal{D}) = \emptyset$.*

Proposition 7 *Let \mathbf{X} and $\{Y\}$ be disjoint node sets in a causal DAG \mathcal{D} such that $Y \in \text{de}(\mathbf{X}, \mathcal{D})$. Then $\mathcal{D}^{\mathbf{X}Y}$ is a causal DAG if and only if there exists a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} .*

Proof First assume that $\mathcal{D}^{\mathbf{X}Y}$ is a causal DAG. Then by Proposition 6, a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} exists.

Now assume that a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} exists. We show that (1) $\mathcal{D}^{\mathbf{X}Y}$ is a DAG syntactically, i.e. a directed graph without cycles, (2) semantically, applying the d -separation criterion to sets of nodes in $\mathcal{D}^{\mathbf{X}Y}$ yields the same separations as applying the d -separation criterion to the same node sets in \mathcal{D} , and (3) $\mathcal{D}^{\mathbf{X}Y}$ is a causal DAG for $(\mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{D})) \cup \mathbf{X} \cup \{Y\}$.

(1) As we assume that a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} exists, it follows from Proposition 6 together with Lemma 16 that $\mathcal{D}^{\mathbf{X}Y}$ does not contain bi-directed edges. Acyclicity of latent projections is guaranteed by property 1 of Definition 4 of the latent projection: every directed edge in $\mathcal{D}(\mathbf{W})$ corresponds to a directed path in \mathcal{D} , hence if $\mathcal{D}^{\mathbf{X}Y}$ had a directed cycle then so would \mathcal{D} . It follows that $\mathcal{D}^{\mathbf{X}Y}$ is acyclic.

(2) The m -separations in a latent projection $\mathcal{D}(\mathbf{W})$ correspond to the d -separations between nodes in \mathbf{W} in the original DAG \mathcal{D} (Richardson et al., 2017). In our case, $\mathcal{D}(\mathbf{W}) = \mathcal{D}^{\mathbf{X}Y}$ is itself a DAG syntactically, and for DAGs m -separation and d -separation are equivalent.

(3) Since \mathcal{D} is a causal DAG for the random variables \mathbf{V} , the truncated factorisation derived from \mathcal{D} holds for all interventions $do(\mathbf{T} = \mathbf{t}')$ with $\mathbf{T} \subseteq \mathbf{V}$:

$$f(\mathbf{v} \mid do(\mathbf{t}')) = \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{V \in \mathbf{V} \setminus \mathbf{T}} f(v \mid \text{pa}(V, \mathcal{D})). \quad (2)$$

We need to show that the truncated factorisation implied by $\mathcal{D}^{\mathbf{X}Y}$ holds for the joint marginal distribution of $(\mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{D})) \cup \mathbf{X} \cup \{Y\}$. We distinguish two cases. In the first case, $Y \notin \text{de}(\mathbf{X}, \mathcal{D})$. This case is trivial, as the forbidden set is then empty and $\mathcal{D} = \mathcal{D}^{\mathbf{X}Y}$. For the second case, $Y \in \text{de}(\mathbf{X}, \mathcal{D})$ we define the following sets: $\mathbf{A} = (\mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{D})) \cup \mathbf{X}$ is the node set of $\mathcal{D}^{\mathbf{X}Y}$ without Y , $\mathbf{C} = \text{cn}(\mathbf{X}, Y, \mathcal{D}) \setminus (\mathbf{X} \cup \{Y\})$ is the set of forbidden nodes that are ancestors of Y , excluding \mathbf{X} and Y , and $\mathbf{M} = \text{forb}(\mathbf{X}, Y, \mathcal{D}) \setminus (\mathbf{X} \cup \{Y\} \cup \mathbf{C})$ is the forbidden set excluding \mathbf{X} , Y and \mathbf{C} , so that $\mathbf{C} \cup \mathbf{M} = \text{forb}(\mathbf{X}, Y, \mathcal{D}) \setminus (\mathbf{X} \cup \{Y\})$ is the set over which we marginalise when subjecting \mathcal{D} to the forbidden projection. Then the following partial topological ordering holds: $\mathbf{A} < \mathbf{C} < Y < \mathbf{M}$. (Note that Y cannot have descendants in \mathbf{X} , as otherwise no valid adjustment set would exist by Lemma 15.)

We can now rewrite equation (2) as

$$f(\mathbf{v} \mid do(\mathbf{t}')) = \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D})) \prod_{C \in \mathbf{C} \setminus \mathbf{T}} f(c \mid \text{pa}(C, \mathcal{D})) \\ f(y \mid \text{pa}(Y, \mathcal{D}))^{1(Y \notin \mathbf{T})} \prod_{M \in \mathbf{M} \setminus \mathbf{T}} f(m \mid \text{pa}(M, \mathcal{D})).$$

Consider now interventions only in nodes $\mathbf{T} \subseteq (\mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{D})) \cup \mathbf{X} \cup \{Y\}$, then $\mathbf{C} \setminus \mathbf{T} = \mathbf{C}$ and $\mathbf{M} \setminus \mathbf{T} = \mathbf{M}$. Upon marginalising the above intervention distribution over

M the last term in the product vanishes but the remaining terms do not change:

$$f(\mathbf{a}, y, \mathbf{c} \mid do(\mathbf{t}')) = \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D})) \prod_{C \in \mathbf{C}} f(c \mid \text{pa}(C, \mathcal{D})) f(y \mid \text{pa}(Y, \mathcal{D}))^{\mathbf{1}(Y \notin \mathbf{T})}.$$

Further marginalising over \mathbf{C} , the partial topological order guarantees that the variables in \mathbf{A} do not have parents in \mathbf{C} . This yields

$$f(\mathbf{a}, y \mid do(\mathbf{t}')) = \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D})) \int_{\mathbf{c}} \prod_{C \in \mathbf{C}} f(c \mid \text{pa}(C, \mathcal{D})) f(y \mid \text{pa}(Y, \mathcal{D}))^{\mathbf{1}(Y \notin \mathbf{T})} d\mathbf{c}.$$

A variable is conditionally independent of its non-descendants given its parents. All variables in $\mathbf{A} \cup \mathbf{C}$ are non-descendants of Y , hence $Y \perp\!\!\!\perp \mathbf{A} \cup \mathbf{C} \mid \text{pa}(Y, \mathcal{D})$ and $f(y \mid \text{pa}(Y, \mathcal{D})) = f(y \mid \text{pa}(Y, \mathcal{D}) \cup \mathbf{a} \cup \mathbf{c}) = f(y \mid \mathbf{a} \cup \mathbf{c})$. The second equality holds because the parents of Y , if there are any, form a subset of $\mathbf{A} \cup \mathbf{C}$. Similarly, all variables in \mathbf{A} are non-descendants of all variables in \mathbf{C} , hence $f(c \mid \text{pa}(C, \mathcal{D})) = f(c \mid \text{pa}(C, \mathcal{D}) \cup \mathbf{a})$. Further, all parents of variables in \mathbf{C} are in $\mathbf{A} \cup \mathbf{C}$, hence $\prod_{C \in \mathbf{C}} f(c \mid \text{pa}(C, \mathcal{D}) \cup \mathbf{a}) = f(\mathbf{c} \mid \mathbf{a})$. We obtain

$$\begin{aligned} f(\mathbf{a}, y \mid do(\mathbf{t}')) &= \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D})) \int_{\mathbf{c}} f(\mathbf{c} \mid \mathbf{a}) f(y \mid \mathbf{a} \cup \mathbf{c})^{\mathbf{1}(Y \notin \mathbf{T})} d\mathbf{c} \\ &= \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D})) \int_{\mathbf{c}} f(\mathbf{c}, y \mid \mathbf{a})^{\mathbf{1}(Y \notin \mathbf{T})} f(\mathbf{c} \mid \mathbf{a})^{\mathbf{1}(Y \in \mathbf{T})} d\mathbf{c} \\ &= \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D})) f(y \mid \mathbf{a})^{\mathbf{1}(Y \notin \mathbf{T})}. \end{aligned}$$

Two things remain to be shown. First, for every $A \in \mathbf{A}$, $\text{pa}(A, \mathcal{D}) = \text{pa}(A, \mathcal{D}^{\mathbf{X}Y})$ because A does not have parents in the node set over which we marginalised in the projection. Second, $f(y \mid \mathbf{a}) = f(y \mid \text{pa}(Y, \mathcal{D}^{\mathbf{X}Y}))$, which follows from the fact that all conditional independencies between variables in $\mathcal{D}^{\mathbf{X}Y}$ can be read off $\mathcal{D}^{\mathbf{X}Y}$ using the d -separation criterion, as we showed in part (2) of this proof. Hence, we have

$$\begin{aligned} f(\mathbf{a}, y \mid do(\mathbf{t}')) &= \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{A \in \mathbf{A} \setminus \mathbf{T}} f(a \mid \text{pa}(A, \mathcal{D}^{\mathbf{X}Y})) f(y \mid \text{pa}(Y, \mathcal{D}^{\mathbf{X}Y}))^{\mathbf{1}(Y \notin \mathbf{T})} \\ &= \mathbf{1}(\mathbf{t} = \mathbf{t}') \prod_{V \in (\mathbf{A} \cup \{Y\}) \setminus \mathbf{T}} f(v \mid \text{pa}(V, \mathcal{D}^{\mathbf{X}Y})), \end{aligned}$$

which is exactly the truncated factorisation formula implied by $\mathcal{D}^{\mathbf{X}Y}$. Hence, $\mathcal{D}^{\mathbf{X}Y}$ is a causal DAG for the random variables $\mathbf{A} \cup \{Y\} = (\mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{D})) \cup \mathbf{X} \cup \{Y\}$. \blacksquare

Lemma 16 *Let \mathcal{D} be a DAG with node set \mathbf{V} and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} exists. Then any edge that is in $\mathcal{D}^{\mathbf{X}Y}$ but not in \mathcal{D} is a directed edge into Y .*

Proof We only consider the case where $Y \in \text{de}(\mathbf{X}, \mathcal{D})$, as otherwise $\mathcal{D} = \mathcal{D}^{\mathbf{X}\mathbf{Y}}$ and our statement follows trivially. Define $\mathbf{F} = \text{forb}(\mathbf{X}, Y, \mathcal{D}) \setminus (\mathbf{X} \cup \{Y\})$.

By Definition 5, an edge present in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ but not in \mathcal{D} only occurs if \mathcal{D} contains a node $W_j \in \mathbf{V} \setminus \mathbf{F}$ that has an ancestor in \mathbf{F} . We show that the only node in $\mathbf{V} \setminus \mathbf{F}$ that can have an ancestor in \mathbf{F} is Y .

Consider first a $W \in \mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{D})$. W does not have ancestors in \mathbf{F} , as otherwise W would be a forbidden node itself. Consider next a node $X \in \mathbf{X}$. X does not have ancestors in \mathbf{F} either, as every node in \mathbf{F} is a descendant of $\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$, but we assume that a valid adjustment set exists relative to (\mathbf{X}, Y) in \mathcal{D} , implying $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D}), \mathcal{D}) = \emptyset$ by Lemma 15. Hence, Y is the only node in $\mathbf{V} \setminus \mathbf{F}$ that can have an ancestor in \mathbf{F} . ■

Proposition 8 *Let \mathbf{X}, \mathbf{Y} and \mathbf{Z} be disjoint node sets in a causal DAG \mathcal{D} . Then \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} if and only if \mathbf{Z} is also a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.*

Proof Throughout, let $\mathbf{F} = \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \setminus (\mathbf{X} \cup \mathbf{Y})$.

We first suppose that \mathbf{Z} is a valid adjustment set in \mathcal{D} and show that this implies that it is also a valid adjustment set in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. Hence, $\mathbf{Z} \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \emptyset$ and $\mathbf{Z} \cap \mathbf{Y} = \emptyset$, so that every node in \mathbf{Z} is also a node of $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. Further, $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) \subseteq \mathbf{X} \cup \mathbf{Y}$ and hence $\mathbf{Z} \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) = \emptyset$. Amenability trivially holds in both \mathcal{D} and $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ by assumption.

It remains to show that every proper non-causal path from \mathbf{X} to \mathbf{Y} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ is blocked by \mathbf{Z} , which we do by contradiction. So suppose that \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} and that there exists a proper non-causal path $p = (V_0, e_1, V_1, \dots, e_K, V_K)$ from \mathbf{X} to \mathbf{Y} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ that is open given \mathbf{Z} . We denote $\mathbf{Y}^F = \mathbf{Y} \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) = \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) \setminus \mathbf{X}$ and note that $\text{de}(\mathbf{Y}^F, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) \subseteq \mathbf{Y}^F$.

Let $V_L \in \mathbf{Y}$ be the first node on p that is in \mathbf{Y} and consider the path segment $p' = (V_0, e_1, V_1, \dots, e_L, V_L)$. Suppose that $L < K$ and that $V_L \in \mathbf{Y}^F$. If p' is causal, then p must either be causal or contain a collider in \mathbf{Y}^F , contradicting our assumption that it is open given \mathbf{Z} . If $V_L \in \mathbf{Y} \setminus \mathbf{Y}^F$ then p' cannot be causal. Hence, we can suppose that $L = K$ or replace p with p' without loss of generality.

Consider now the case that $V_K \in \mathbf{Y} \setminus \mathbf{Y}^F$. This implies that all nodes on p except V_0 are not in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}^{\mathbf{X}\mathbf{Y}})$. Since $\text{de}(\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) \setminus \mathbf{X}, \mathcal{D}) \subseteq \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$ and by definition of latent projections, this implies that p is also a path in \mathcal{D} . As for any node V in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$, $\text{de}(V, \mathcal{D}) \subseteq \text{de}(V, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) \cup \mathbf{F}$, and $\mathbf{Z} \cap \mathbf{F} = \emptyset$, it follows that p is also open given \mathbf{Z} in \mathcal{D} .

Consider now the case that $V_K \in \mathbf{Y}^F$. The path p cannot be a one-edge path, as the two possible such paths would require the existence of paths in \mathcal{D} implying that no valid adjustment sets relative to (\mathbf{X}, \mathbf{Y}) exist in \mathcal{D} . By the fact that $\text{de}(\mathbf{Y}^F, \mathcal{D}) \subseteq \mathbf{Y}^F$, the last edge of p must be of the form $p'' = V_{K-1} \rightarrow V_K$. By the same argument and the definition of the forbidden projection, the segment $p' = (V_0, e_1, V_1, \dots, V_{K-1})$ is also a path in \mathcal{D} , which by definition of p and p'' must be non-causal. The path p'' corresponds to a causal path q'' in \mathcal{D} , such that all nodes except for V_{K-1} on q'' are forbidden.

The path $q = p' \oplus q''$ is a proper non-causal path from \mathbf{X} to \mathbf{Y} in \mathcal{D} ; we now show that it is open given \mathbf{Z} . Since $\text{de}(V, \mathcal{D}) \subseteq \text{de}(V, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) \cup \mathbf{F}$, for any node $V \notin \mathbf{F}$ in \mathcal{D} , it follows from the fact that p' is open given \mathbf{Z} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ that it is also open given \mathbf{Z} in \mathcal{D} . Since $\mathbf{F} \cap \mathbf{Z} = \emptyset$,

q'' is also open given \mathbf{Z} . The node V_{K-1} is a non-collider on p and hence by the assumption that p is open given \mathbf{Z} it follows that $V_{K-1} \notin \mathbf{Z}$. Since V_{K-1} is also a non-collider on q it follows that q is open given \mathbf{Z} in \mathcal{D} .

We now turn to the second part of the proof showing that if a set \mathbf{Z} containing no nodes in \mathbf{F} is not a valid adjustment relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} then it is also not a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.

Suppose that $\mathbf{Z} \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \neq \emptyset$. Since $\mathbf{Z} \cap \mathbf{F} = \emptyset$ it follows that $\mathbf{Z} \cap (\mathbf{X} \cup \mathbf{Y}) \neq \emptyset$; but this clearly implies that \mathbf{Z} cannot be a valid adjustment set in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$.

Suppose now that $\mathbf{Z} \cap (\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \cup \mathbf{Y}) = \emptyset$ and that \mathbf{Z} is not a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} . This implies the existence of a proper non-causal path $p = (V_0, e_1, V_1, \dots, e_K, V_K)$ from \mathbf{X} to \mathbf{Y} in \mathcal{D} that is open given \mathbf{Z} .

Consider first the case that p contains no nodes in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Then p also exists in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ and by the fact that $\text{de}(V, \mathcal{D}^{\mathbf{X}\mathbf{Y}}) = \text{de}(V, \mathcal{D}) \setminus \mathbf{F}$ it follows that p is also open given \mathbf{Z} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. Suppose now that p contains at least one node in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Since p is open given \mathbf{Z} , it cannot contain a collider in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. If we suppose that all nodes on p are in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$, then its existence implies that no valid adjustment exists in \mathcal{D} , while the corresponding edge in the forbidden projection would imply the same for $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. Hence, we can suppose that p contains at least one node not in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$ and let V_L be the last such node. Let $p' = (V_0, e_1, V_1, \dots, e_L, V_L)$ and $p'' = (V_L, e_{L+1}, V_{L+1}, \dots, e_K, V_K)$. By construction, $V_{L+1} \in \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$ and since $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \setminus \mathbf{X} \subseteq \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$, it follows that p'' is causal. Thus the forbidden projection will map p'' to the path $q'' = V_L \rightarrow V_K$. This also implies that p' is non-causal.

Suppose first that $V_0 \in \mathbf{X} \cap \text{de}(\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \setminus \mathbf{X}, \mathcal{D})$. Then no valid adjustment set exists in \mathcal{D} . Further, there must be a bi-directed edge from \mathbf{X} to \mathbf{Y} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ and hence that no valid adjustment set exists in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ either. We can hence suppose that $V_0 \notin \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \setminus \mathbf{X}$. This implies that all nodes on p' except V_0 are not in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Since this implies that no node on p' is in $\text{de}(\mathbf{F}, \mathcal{D})$ it follows that p' is also a path in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. The path $q = p' \oplus p''$ is a proper non-causal path from \mathbf{X} to \mathbf{Y} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. By the usual argument p' is also open given \mathbf{Z} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ and trivially, this is also true for q'' . Further, $V_L \notin \mathbf{Z}$ is a non-collider on q'' and hence, q is open given \mathbf{Z} in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. \blacksquare

Proposition 10 *Let \mathbf{X} and \mathbf{Y} be disjoint subsets of the node set \mathbf{V} of a DAG \mathcal{D} such that $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$. Then $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D})$.*

Proof We first show that $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \subseteq \mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Let $Z \in \mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. By Definition 1, $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \cap \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \emptyset$ and hence Z is a node in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$. Furthermore, since $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \subseteq \text{pa}(\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D}), \mathcal{D})$, and $\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \subseteq \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$, there is a node $Y \in \mathbf{Y}$ such that \mathcal{D} contains a directed path $Z \rightarrow \dots \rightarrow Y$ on which all non-endpoint nodes are in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Due to property 1 of Definition 4, this corresponds to an edge $Z \rightarrow Y$ in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$, hence $Z \in \mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D})$.

Next, we show that $\mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \subseteq \mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Let $Z^* \in \mathbf{O}^*(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. By Definition 5, this implies that $Z^* \in \mathbf{V} \setminus (\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \cup \mathbf{X} \cup \mathbf{Y})$. Moreover, by Definition 9, there is an edge $Z^* \rightarrow Y^*$ in $\mathcal{D}^{\mathbf{X}\mathbf{Y}}$ with $Y^* \in \mathbf{Y}$. In \mathcal{D} , this corresponds to a directed path $Z^* \rightarrow \dots \rightarrow Y^*$ on which all non-endpoint nodes are in $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$, and

$Z^* \notin \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. Denote the path by p . There are two cases: In the first case, p has no non-endpoint nodes, i.e. \mathcal{D} contains the edge $Z^* \rightarrow Y^*$. Since we assume $\mathbf{Y} \subseteq \text{de}(\mathbf{X}, \mathcal{D})$, Y^* must be in $\text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$, hence $Z^* \in \mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. In the second case, p has at least one non-endpoint node. This means that $Z^* \in \text{pa}(W, \mathcal{D})$, where $W \in \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) \setminus (\mathbf{X} \cup \mathbf{Y})$ and $W \in \text{an}(Y^*, \mathcal{D})$. Since in a DAG, all forbidden nodes are descendants of \mathbf{X} , we also have $W \in \text{de}(\mathbf{X}, \mathcal{D})$, and hence $W \in \text{cn}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. It follows that $Z^* \in \mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$. \blacksquare

Appendix C. Generalisation of the Forbidden Projection and the \mathbf{O}^* -Set to Amenable MaxPDAGs

In this appendix, we generalise the forbidden projection (Definition 5) and the \mathbf{O}^* -set (Definition 9) to amenable maxPDAGs and show that Propositions similar to 6, 8, 7 and 10 still hold for the more general definitions.

The latent projection in general (Definition 4) cannot be generalised to (amenable) maxPDAGs as marginalising does not generally result in an ADMG. As an example, consider the maxPDAG $W_1 - L \rightarrow W_2$ with latent node L . It is not clear how the projection should be constructed in this case: $W_1 \rightarrow W_2$ would give the wrong impression that W_1 is an ancestor of W_2 (instead of a *possible* ancestor), while $W_1 - W_2$ would imply that W_2 is a possible ancestor of W_1 . As we will show in the following propositions, however, the latent projection can be meaningfully generalised to amenable maxPDAGs when projecting over the special case of a forbidden set.

Definition 17 (Forbidden projection for amenable maxPDAGs) *Let \mathcal{G} be a maxPDAG with node set \mathbf{V} , and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that \mathcal{G} is amenable relative to (\mathbf{X}, Y) . Define $\mathbf{F} = \text{forb}(\mathbf{X}, Y, \mathcal{G}) \setminus (\mathbf{X} \cup \{Y\})$. The forbidden projection $\mathcal{G}^{\mathbf{X}Y}$ of \mathcal{G} is a graph with node set $\mathbf{V} \setminus \mathbf{F}$ and edges as follows: For distinct nodes $W_i, W_j \in \mathbf{V} \setminus \mathbf{F}$,*

1. $\mathcal{G}^{\mathbf{X}Y}$ contains a directed edge $W_i \rightarrow W_j$ if and only if \mathcal{G} contains a directed path $W_i \rightarrow \dots \rightarrow W_j$ on which all non-endpoint nodes are in \mathbf{F} ,
2. $\mathcal{G}^{\mathbf{X}Y}$ contains a bi-directed edge $W_i \leftrightarrow W_j$ if and only if \mathcal{G} contains a path, with at least one non-endpoint node, of the form $W_i \leftarrow \dots \rightarrow W_j$ on which all non-endpoints are non-colliders and in \mathbf{F} ,
3. $\mathcal{G}^{\mathbf{X}Y}$ contains an undirected edge $W_i - W_j$ if and only if \mathcal{G} contains $W_i - W_j$.

Note that we restrict the definition to singleton Y . This is because with a set \mathbf{Y} , we run into similar construction/interpretation problems as described above. Consider, for example, an amenable maxPDAG with node set $\{X, F, Y_1, Y_2\}$ and edges $X \rightarrow Y_1 - F \rightarrow Y_2$ as well as $X \rightarrow F$. None of $Y_1 \rightarrow Y_2$, $Y_1 \leftrightarrow Y_2$ or $Y_1 - Y_2$ are correct representations of the marginal distribution.

Before generalising the \mathbf{O}^* -set, we now describe the properties of the forbidden projection for maxPDAGs. A key property is that if a valid adjustment set exists, the forbidden projection of a maxPDAG is itself a maxPDAG (Proposition 22). This is analogous to Proposition 7 for DAGs. Proposition 19, in analogy to Proposition 6, states that if \mathbf{X} has

a causal effect on Y , then the forbidden projection can be used to check whether a valid adjustment set exists. In Proposition 25, we show that a set \mathbf{Z} is a valid adjustment set in the forbidden projection if and only if it is a valid adjustment set in the original graph, which is analogous to Proposition 8.

We begin with a lemma that will allow us to use $\text{possde}(\mathbf{X}, \mathcal{G})$ and $\text{de}(\mathbf{X}, \mathcal{G})$ interchangeably when \mathcal{G} is an amenable maxPDAG.

Lemma 18 *Let $p = (V_1, V_2, \dots, V_K)$ be a possibly directed path in a maxPDAG \mathcal{G} such that no node on p shares an undirected edge with V_1 . Then a subsequence of p forms a directed path from V_1 to V_K in \mathcal{G} .*

Proof We show this by induction. By assumption, (V_1, V_2) is a subsequence of p and forms a directed path from V_1 to V_2 . Now assume that a subsequence of p forms a directed path from V_1 to V_{k-1} , for $2 < k \leq K$. Denote this subsequence by $(V_1 = W_1, W_2, \dots, W_Q = V_{k-1})$. Clearly, if $V_{k-1} \rightarrow V_k$ then $(V_1 = W_1, W_2, \dots, W_Q = V_{k-1}, V_k)$ is a subsequence of p and forms a directed path from V_1 to V_k in \mathcal{G} , which is what we wanted to show. If, on the other hand, if $V_{k-1} - V_k$ then there are four cases, three of which lead to a contradiction:

(1) The induced subgraph of \mathcal{G} on $\{W_{Q-1}, W_Q = W_{k-1}, V_k\}$ is Graph 1 in Figure 9. Then \mathcal{G} is not closed under Meek's Rule 1 (see Figure 8), which is a contradiction.

(2) The induced subgraph of \mathcal{G} on $\{W_{Q-1}, W_Q = W_{k-1}, V_k\}$ is Graph 2 in Figure 9. Then \mathcal{G} is not closed under Meek's Rule 2, which is a contradiction.

(3) The induced subgraph of \mathcal{G} on $\{W_{Q-1}, W_Q = W_{k-1}, V_k\}$ is Graph 3 in Figure 9. This implies the induced subgraph of \mathcal{G} on $\{W_{Q-2}, W_{Q-1}, V_k\}$ would also be graph 3 (with the same reasons as above excluding graphs 1 and 2). Repeating the argument for W_{Q-3}, W_{Q-4}, \dots implies an undirected edge between $W_1 = V_1$ and V_k , which is a contradiction.

(4) The induced subgraph of \mathcal{G} on $\{W_{Q-1}, W_Q = W_{k-1}, V_k\}$ is Graph 4 in Figure 9. Then $(V_1 = W_1, W_2, \dots, W_{Q-1}, V_k)$ is a subsequence of p and forms a directed path from V_1 to V_k , which is what we wanted to show. \blacksquare




Figure 9: Graphs for the proof of Lemma 18.

Proposition 19 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} , and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that $Y \in \text{possde}(\mathbf{X}, \mathcal{G})$ and \mathcal{G} is amenable relative to (\mathbf{X}, Y) . Then a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} exists if and only if there is no bi-directed edge between any nodes in the forbidden projection $\mathcal{G}^{\mathbf{X}Y}$.*

Proof By Lemma 18, $Y \in \text{possde}(\mathbf{X}, \mathcal{G})$ implies $Y \in \text{de}(\mathbf{X}, \mathcal{G})$. The proof is now analogous to the proof of Proposition 6, with Lemma 15 replaced by Lemma 20. \blacksquare

Lemma 20 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} , and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that $Y \in \text{de}(\mathbf{X}, \mathcal{G})$ and \mathcal{G} is amenable relative to (\mathbf{X}, Y) . Then a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} exists if and only if $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}), \mathcal{G}) = \emptyset$.*

Proof We show that *no* valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} exists if and only if $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}), \mathcal{G}) \neq \emptyset$. The proof is similar to the proof of Corollary 27 in Perković et al. (2018).

Assume first that no valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} exists, then by Lemma 21, there is a proper non-causal definite-status path from some $X \in \mathbf{X}$ to Y that is open given $\text{adjust}(\mathbf{X}, Y, \mathcal{G}) = \text{possan}(\mathbf{X} \cup \{Y\}, \mathcal{G}) \setminus (\mathbf{X} \cup \text{forb}(\mathbf{X}, Y, \mathcal{G}))$. Denote one such path by p . Assume for contradiction that p contains a collider and denote the collider by C . As p is open given $\text{adjust}(\mathbf{X}, Y, \mathcal{G})$, a descendant of C is in $\text{adjust}(\mathbf{X}, Y, \mathcal{G})$. This implies that $\text{an}(C, \mathcal{G}) \cap \text{forb}(\mathbf{X}, Y, \mathcal{G}) = \emptyset$, as otherwise all descendants of C would be in $\text{forb}(\mathbf{X}, Y, \mathcal{G})$ and could not be in $\text{adjust}(\mathbf{X}, Y, \mathcal{G})$. As $Y \in \text{forb}(\mathbf{X}, Y, \mathcal{G})$, it follows that at least one of the nodes adjacent to C on p must be a non-endpoint non-collider on p . Denote one such node by B . As $B \in \text{pa}(C, \mathcal{G})$, $B \notin \text{forb}(\mathbf{X}, Y, \mathcal{G})$ and $B \in \text{adjust}(\mathbf{X}, Y, \mathcal{G})$. But then p is not open given $\text{adjust}(\mathbf{X}, Y, \mathcal{G})$, which is a contradiction. Hence, p does not contain a collider. As p is non-causal, p cannot be directed towards Y , and as we assume that $Y \in \text{de}(\mathbf{X}, \mathcal{G})$, p cannot be directed towards X . Hence, p is a path of the form $X \leftarrow \dots \leftarrow A \rightarrow \dots \rightarrow Y$, where every non-endpoint is a non-collider not in $\text{adjust}(\mathbf{X}, Y, \mathcal{G})$. It follows that A is in $\text{forb}(\mathbf{X}, Y, \mathcal{G})$ and thus is a descendant of \mathbf{X} , which implies that $X \in \text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}), \mathcal{G})$ and hence $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}), \mathcal{G}) \neq \emptyset$.

Assume now that $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}), \mathcal{G}) \neq \emptyset$. Pick a node from $\mathbf{X} \cap \text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}), \mathcal{G})$ and denote it by X^* . Then there must exist a node $C^* \in \text{cn}(\mathbf{X}, Y, \mathcal{G})$ such that there is a path of the form $X^* \leftarrow \dots \leftarrow C^* \rightarrow \dots \rightarrow Y$ where that all non-endpoint non-colliders on the path are in the forbidden set. This path cannot be blocked by any set of non-forbidden nodes. \blacksquare

Lemma 21 (Theorem 5.6 in Perković et al., 2017) *Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a causal maxPDAG \mathcal{G} such that \mathcal{G} is amenable relative to (\mathbf{X}, \mathbf{Y}) , and let $\text{adjust}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \text{possan}(\mathbf{X} \cup \mathbf{Y}, \mathcal{G}) \setminus (\mathbf{X} \cup \mathbf{Y} \cup \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}))$. Then a valid adjustment set relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{D} exists if and only if all proper non-causal definite-status paths from \mathbf{X} to \mathbf{Y} are blocked by $\text{adjust}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$ in \mathcal{G} .*

Proposition 22 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} , and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that \mathcal{G} is amenable relative to (\mathbf{X}, Y) and there exists a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} . Denote the set of DAGs represented by \mathcal{G} by $[\mathcal{G}] = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_M\}$. Then the forbidden projection $\mathcal{G}^{\mathbf{X}Y}$ is the causal maxPDAG representing the DAGs in $\{\mathcal{D}_1^{\mathbf{X}Y}, \mathcal{D}_2^{\mathbf{X}Y}, \dots, \mathcal{D}_M^{\mathbf{X}Y}\}$.*

Proof We only consider the case that $Y \in \text{possde}(\mathbf{X}, \mathcal{G})$, as otherwise the proposition follows trivially from the fact that $\mathcal{G}^{\mathbf{X}Y} = \mathcal{G}$. By Lemma 18, $Y \in \text{possde}(\mathbf{X}, \mathcal{G})$ implies $Y \in \text{de}(\mathbf{X}, \mathcal{G})$. We know from Propositions 6 and 19 that none of $\mathcal{G}^{\mathbf{X}Y}, \mathcal{D}_1^{\mathbf{X}Y}, \mathcal{D}_2^{\mathbf{X}Y}, \dots, \mathcal{D}_M^{\mathbf{X}Y}$ contain any bi-directed edges. Consider edges present in the latent projections but not in the original graphs: For the maxPDAG \mathcal{G} , denote the set of edges in $\mathcal{G}^{\mathbf{X}Y}$ but not in \mathcal{G} by $\mathbf{e}(\mathcal{G})$, and define analogous sets $\mathbf{e}(\mathcal{D}_1), \mathbf{e}(\mathcal{D}_2), \dots, \mathbf{e}(\mathcal{D}_M)$ for the DAGs $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_M$. None of $\mathbf{e}(\mathcal{G}), \mathbf{e}(\mathcal{D}_1), \mathbf{e}(\mathcal{D}_2), \dots, \mathbf{e}(\mathcal{D}_M)$ contain any undirected edges. Further, any directed edges in any of $\mathbf{e}(\mathcal{G}), \mathbf{e}(\mathcal{D}_1), \mathbf{e}(\mathcal{D}_2), \dots, \mathbf{e}(\mathcal{D}_M)$ are into Y . This is because for every $\mathcal{G}' \in \{\mathcal{G}, \mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_M\}$, an edge in $\mathbf{e}(\mathcal{G}')$ corresponds to a directed path with at least one forbidden non-endpoint node in \mathcal{G}' . If an edge in $\mathbf{e}(\mathcal{G}')$ was into a node $V \in \mathbf{V} \setminus (\mathbf{X} \cup \{Y\})$, then V would itself be forbidden, which is a contradiction. If an edge in $\mathbf{e}(\mathcal{G}')$ was into a node $X \in \mathbf{X}$, then X would be in $\text{de}(\text{cn}(\mathbf{X}, Y, \mathcal{G}'), \mathcal{G}')$, which by Lemma 20 contradicts our assumption that a valid adjustment set exists relative to (\mathbf{X}, Y) in \mathcal{G}' . Hence, all edges in all of $\mathbf{e}(\mathcal{G}), \mathbf{e}(\mathcal{D}_1), \mathbf{e}(\mathcal{D}_2), \dots, \mathbf{e}(\mathcal{D}_M)$ are into Y . In fact, by Lemma 23 below, $\mathbf{e}(\mathcal{G}) = \mathbf{e}(\mathcal{D}_1) = \mathbf{e}(\mathcal{D}_2) = \dots = \mathbf{e}(\mathcal{D}_M) = \mathbf{e}$. The graphs $\mathcal{G}^{\mathbf{X}Y}, \mathcal{D}_1^{\mathbf{X}Y}, \mathcal{D}_2^{\mathbf{X}Y}, \dots, \mathcal{D}_M^{\mathbf{X}Y}$ can thus be constructed by copying the induced subgraphs of $\mathcal{G}, \mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_M$ with respect to $(\mathbf{V} \setminus \text{forb}(\mathbf{X}, Y, \mathcal{G})) \cup (\mathbf{X}, Y)$ and adding the edges in \mathbf{e} . Hence, $\mathcal{G}^{\mathbf{X}Y}$ represents the DAGs in $\{\mathcal{D}_1^{\mathbf{X}Y}, \mathcal{D}_2^{\mathbf{X}Y}, \dots, \mathcal{D}_M^{\mathbf{X}Y}\}$ in the sense that every directed edge in $\mathcal{G}^{\mathbf{X}Y}$ is also present in all DAGs in $\{\mathcal{D}_1^{\mathbf{X}Y}, \mathcal{D}_2^{\mathbf{X}Y}, \dots, \mathcal{D}_M^{\mathbf{X}Y}\}$, and for every undirected edge $V_i - V_j$ in $\mathcal{G}^{\mathbf{X}Y}$, there is at least one DAG in $\{\mathcal{D}_1^{\mathbf{X}Y}, \mathcal{D}_2^{\mathbf{X}Y}, \dots, \mathcal{D}_M^{\mathbf{X}Y}\}$ with $V_i \rightarrow V_j$ and at least one with $V_i \leftarrow V_j$.

In order show that $\mathcal{G}^{\mathbf{X}Y}$ has all the characteristics of a maxPDAG, we show that $\mathcal{G}^{\mathbf{X}Y}$ is closed under Meek's rules. Referring to Figure 8, we argue that the graphs on the left-hand sides of Rules 1 – 4 cannot be induced subgraphs of $\mathcal{G}^{\mathbf{X}Y}$. Assume for contradiction that the left-hand graph of Rule 1, $\rightarrow -$, was an induced subgraph of $\mathcal{G}^{\mathbf{X}Y}$. As this graph is not an induced subgraph of \mathcal{G} by assumption, and all of $\mathbf{e}(\mathcal{G}), \mathbf{e}(\mathcal{D}_1), \mathbf{e}(\mathcal{D}_2), \dots, \mathbf{e}(\mathcal{D}_M)$ consist of only directed edges into Y , we can conclude that the directed edge in $\rightarrow -$ is into Y , i.e. $\rightarrow Y -$. Hence, Y shares an undirected edge with some node V in \mathcal{G} , but this means that V is a forbidden node in some $\mathcal{D} \in [\mathcal{G}]$, which is not allowed according to Lemma 24. By similar arguments, none of the graphs on the left-hand sides of Rules 1 – 4 in Figure 8 is an induced subgraph of $\mathcal{G}^{\mathbf{X}Y}$. \blacksquare

Lemma 23 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} , and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that \mathcal{G} is amenable relative to (\mathbf{X}, Y) and there exists a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} . Define $\mathbf{F} = \text{forb}(\mathbf{X}, Y, \mathcal{G}) \setminus (\mathbf{X} \cup \{Y\})$ and pick a node $V_1 \in \mathbf{V} \setminus \mathbf{F}$. Then the following two statements are equivalent:*

- (i) *A DAG $\mathcal{D} \in [\mathcal{G}]$ contains a directed path $p = (V_1, V_2, \dots, V_K = Y)$ such that all non-endpoint nodes on p are in \mathbf{F} .*
- (ii) *The maxPDAG \mathcal{G} contains a directed path $q = (V_1 = W_1, W_2, \dots, W_Q = Y)$ such that all non-endpoint nodes on q are in \mathbf{F} .*

Proof Statement (ii) implies that the directed path p is present in all DAGs in $[\mathcal{G}]$ by the defining properties of a maxPDAG. Hence, we only show that (i) implies (ii). Again by the properties of maxPDAGs, the sequence of nodes $(V_1, V_2, \dots, V_K = Y)$ forms a possibly

directed path from V_1 to Y in \mathcal{G} . We first show that no node in $\{V_2, \dots, V_K = Y\}$ shares an undirected edge with V_1 . Suppose, for contradiction, that node $V_k, 2 \leq k \leq K$ shares an undirected edge with V_1 and distinguish two cases: (1) $V_1 \in \mathbf{X}$, (2) $V_1 \in \mathbf{V} \setminus \mathbf{X}$. The first case contradicts our assumption that \mathcal{G} is amenable relative to (\mathbf{X}, Y) . The second case implies that V_1 , as a possible descendant of V_k , is in \mathbf{F} , but we chose V_1 such that $V_1 \in \mathbf{V} \setminus \mathbf{F}$. Hence, no node in $\{V_2, \dots, V_K = Y\}$ shares an undirected edge with V_1 . We can thus apply Lemma 18 and conclude that a subsequence of $(V_1, V_2, \dots, V_K = Y)$ forms a directed path from V_1 to Y in \mathcal{G} , which implies that statement (ii) holds. ■

Lemma 24 (Lemma E.8 in HPM19) *Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a maxPDAG \mathcal{G} , such that \mathcal{G} is amenable relative to (\mathbf{X}, \mathbf{Y}) , and let \mathcal{D} be a DAG in $[\mathcal{G}]$. Then $\text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \text{forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D})$.*

Proposition 25 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} and let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that \mathcal{G} is amenable relative to (\mathbf{X}, Y) . Then a set \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} if and only if it is a valid adjustment set relative to (\mathbf{X}, Y) in the forbidden projection $\mathcal{G}^{\mathbf{X}Y}$.*

Proof Let $\mathcal{D} \in [\mathcal{G}]$ and $\mathcal{D}^{\mathbf{X}Y}$ its forbidden projection. By Proposition 8, a set \mathbf{Z} is a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{D} if and only if it is a valid adjustment set relative to (\mathbf{X}, Y) in $\mathcal{D}^{\mathbf{X}Y}$. By Proposition 22, the set $[\mathcal{G}^{\mathbf{X}Y}]$ contains exactly the forbidden projections of all DAGs in $[\mathcal{G}]$. Hence if \mathbf{Z} is a valid adjustment set in all $\mathcal{D} \in [\mathcal{G}]$, then it is a valid adjustment set in all $\mathcal{D}^{\mathbf{X}Y} \in [\mathcal{G}^{\mathbf{X}Y}]$ and vice versa. ■

To summarise, the forbidden projection for amenable causal maxPDAGs has very similar properties as the forbidden projection for causal DAGs, as long as we consider a singleton outcome node Y : Bi-directed edges in the projection indicate the lack of a valid adjustment set; if a valid set exists, the forbidden projection is a maxPDAG itself, preserving all the information relevant to causal effect identification via adjustment; in particular, all valid sets can be read off the forbidden projection as well as the original graph.

Finally, we now generalise Definition 9 of the \mathbf{O}^* -set and its optimality property in Proposition 10 to amenable maxPDAGs with singleton Y .

Definition 26 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} , let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that \mathcal{G} is amenable relative to (\mathbf{X}, Y) , and let $\mathcal{G}^{\mathbf{X}Y}$ be the corresponding forbidden projection. We define $\mathbf{O}^*(\mathbf{X}, Y, \mathcal{G})$ as:*

$$\mathbf{O}^*(\mathbf{X}, Y, \mathcal{G}) = \text{pa}(Y, \mathcal{G}^{\mathbf{X}Y}) \setminus \mathbf{X}.$$

Proposition 27 *Let \mathcal{G} be a causal maxPDAG with node set \mathbf{V} , let $\mathbf{X} \subset \mathbf{V}$ and $Y \in \mathbf{V} \setminus \mathbf{X}$ such that $Y \in \text{posse}(\mathbf{X}, \mathcal{D})$ and \mathcal{G} is amenable relative to (\mathbf{X}, Y) , let \mathbf{Z} be a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} and let $\mathbf{O}^* = \mathbf{O}^*(\mathbf{X}, Y, \mathcal{G})$. Then \mathbf{O}^* is a valid adjustment set relative to (\mathbf{X}, Y) in \mathcal{G} and if the variables \mathbf{V} follow a linear causal model compatible with \mathcal{G} , then, for every $X_i \in \mathbf{X}$, $a.\text{var}(\hat{\beta}_{y x_i \mathbf{x} \mathbf{O}^*}) \leq a.\text{var}(\hat{\beta}_{y x_i \mathbf{x} \mathbf{Z}})$.*

Proof We prove this by showing that $\mathbf{O}(\mathbf{X}, Y, \mathcal{G})$ and $\mathbf{O}^*(\mathbf{X}, Y, \mathcal{G})$ are equal and invoking Propositions 2 and 3. By Lemma 18, $Y \in \text{posdde}(\mathbf{X}, \mathcal{G})$ implies $Y \in \text{de}(\mathbf{X}, \mathcal{G})$. Then the equivalence follow directly from Lemma 28 in combination with Proposition 10: For every DAG $\mathcal{D} \in [\mathcal{G}]$, $\mathbf{O}(\mathbf{X}, Y, \mathcal{G}) = \mathbf{O}(\mathbf{X}, Y, \mathcal{D}) = \mathbf{O}^*(\mathbf{X}, Y, \mathcal{D}) = \mathbf{O}^*(\mathbf{X}, Y, \mathcal{G})$. \blacksquare

Lemma 28 (Lemma E.7 in HPM19) *Let \mathbf{X} and \mathbf{Y} be disjoint node sets in a maxPDAG \mathcal{G} such that \mathcal{G} is amenable relative to (\mathbf{X}, \mathbf{Y}) in \mathcal{G} , let $\mathbf{Y} \subseteq \text{posdde}(\mathbf{X}, \mathcal{G})$, and let \mathcal{D} be a DAG in $[\mathcal{G}]$. Then $\mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \mathbf{O}(\mathbf{X}, \mathbf{Y}, \mathcal{G})$.*

Appendix D. Proof for Section 4

Proposition 11 *Let X and Y be nodes in a causal CPDAG or maxPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, such that \mathbf{V} follows a causal linear model compatible with \mathcal{G} with Gaussian errors. Let $\widehat{\Theta}^{\mathbf{P}}$ and $\widehat{\Theta}^{\mathbf{O}}$ be the multisets returned by semi-local IDA and optimal IDA respectively, applied to X, Y and \mathcal{G} , with the subsets of $\text{sib}(X, \mathcal{G})$ considered in the same order for both. Then, for $i \in \{1 \dots, k\}$, with $k = |\widehat{\Theta}^{\mathbf{P}}| = |\widehat{\Theta}^{\mathbf{O}}|$,*

1. $\mathbb{E}[\widehat{\Theta}_i^{\mathbf{P}}] = \mathbb{E}[\widehat{\Theta}_i^{\mathbf{O}}]$ and
2. $a.\text{var}(\widehat{\Theta}_i^{\mathbf{P}}) \geq a.\text{var}(\widehat{\Theta}_i^{\mathbf{O}})$.

Proof Consider any set $\mathbf{S}_i \subseteq \text{sib}(X, \mathcal{G})$. Perković et al. (2017) showed that there exists a DAG $\mathcal{D} \in [\mathcal{G}]$, such that $\mathbf{P}_i = \text{pa}(X, \mathcal{D}) = \mathbf{S}_i \cup \text{pa}(X, \mathcal{G})$, if and only if directing the edges in the neighbourhood of X according to \mathbf{P}_i and applying Meek's orientation rules results in a valid maxPDAG \mathcal{G}'_i . If this is not the case for \mathbf{S}_i , both algorithms discard \mathbf{S}_i at their respective line 8. We can hence suppose that there exists a DAG $\mathcal{D} \in [\mathcal{G}]$, such that $\mathbf{P}_i = \text{pa}(X, \mathcal{D}) = \mathbf{S}_i \cup \text{pa}(X, \mathcal{G})$.

Suppose that $Y \in \text{posdde}(X, \mathcal{G}'_i)$. In this case $\widehat{\Theta}_i^{\mathbf{O}} = \widehat{\beta}_{yx.\mathbf{o}_i}$, where $\mathbf{O}_i = \mathbf{O}(X, Y, \mathcal{G}'_i)$. As \mathcal{G}'_i is amenable by construction, it follows from Lemma 28 in Appendix C that $\mathbf{O}_i = \mathbf{O}(X, Y, \mathcal{D})$. Further, $Y \in \text{posdde}(X, \mathcal{G}'_i)$ implies that $Y \notin \text{pa}(X, \mathcal{G}'_i)$ and thus $\widehat{\Theta}_i^{\mathbf{P}} = \widehat{\beta}_{yx.\mathbf{p}_i}$. By Proposition 2, \mathbf{O}_i is a valid adjustment set relative to (X, Y) in \mathcal{D} , and clearly the same holds for \mathbf{P}_i . Since we suppose multivariate Gaussianity, this implies that both $\widehat{\beta}_{yx.\mathbf{o}_i}$ and $\widehat{\beta}_{yx.\mathbf{p}_i}$ are consistent estimators of $\tau_{yx}(\mathcal{D})$, and $\mathbb{E}[\widehat{\beta}_{yx.\mathbf{o}_i}] = \mathbb{E}[\widehat{\beta}_{yx.\mathbf{p}_i}] = \tau_{yx}(\mathcal{D})$.

Further, by Lemmas E.4 and E.5 of the Supplement of HPM19, $\mathbf{P}_i \setminus \mathbf{O}_i$ is conditionally independent of Y given $\{X\} \cup \mathbf{P}_i$, and $\mathbf{O}_i \setminus \mathbf{P}_i$ is conditionally independent of X given \mathbf{P}_i , respectively. These two independencies allow us to invoke Lemma C.2 of HPM19 and conclude that $\sigma_{yy.x\mathbf{o}_i} \leq \sigma_{yy.x\mathbf{p}_i}$ as well as $\sigma_{xx.\mathbf{o}_i} \geq \sigma_{xx.\mathbf{p}_i}$. As we assume a multivariate Gaussian distribution, it follows that

$$a.\text{var}(\widehat{\beta}_{yx.\mathbf{o}_i}) = \frac{\sigma_{yy.x\mathbf{o}_i}}{\sigma_{xx.\mathbf{o}_i}} \leq \frac{\sigma_{yy.x\mathbf{p}_i}}{\sigma_{xx.\mathbf{p}_i}} = a.\text{var}(\widehat{\beta}_{yx.\mathbf{p}_i}).$$

Suppose now that $Y \notin \text{posdde}(X, \mathcal{G}'_i)$. Then $Y \notin \text{de}(X, \mathcal{D})$, hence $\tau_{yx}(\mathcal{D}) = 0$. As $Y \notin \text{posdde}(X, \mathcal{G}'_i)$, $\widehat{\Theta}_i^{\mathbf{O}} = 0$ and as a result $a.\text{var}(\widehat{\Theta}_i^{\mathbf{O}}) = 0$. If $Y \in \text{pa}(X, \mathcal{G}'_i)$, then $\widehat{\Theta}_i^{\mathbf{P}} = 0$ and as a result $a.\text{var}(\widehat{\Theta}_i^{\mathbf{P}}) = 0$. If $Y \notin \text{posdde}(X, \mathcal{G}'_i) \cup \text{pa}(X, \mathcal{G}'_i)$, then $\widehat{\Theta}_i^{\mathbf{P}} = \widehat{\beta}_{yx.\mathbf{p}_i}$ and by nature of parent sets $\mathbb{E}[\widehat{\beta}_{yx.\mathbf{p}_i}] = 0$. Clearly, $a.\text{var}(\widehat{\Theta}_i^{\mathbf{P}}) > 0$ in this case. \blacksquare

Appendix E. Violin Plots




Figure 10: Violin plots of the relative mean squared error (RMSE) r over 1000 repetitions for scenarios with different numbers of nodes (p), expected number of neighbours per node (d), and sample sizes (n). Optimal and semi-local IDA were applied to the true CPDAG \mathcal{G} . The dots mark the geometric means, the plus signs the medians.




Figure 11: Violin plots of the relative mean squared error (RMSE) r^* over 1000 repetitions for scenarios with different numbers of nodes (p), expected number of neighbours per node (d), and sample sizes (n). Optimal and semi-local IDA were applied to the estimated CPDAG \mathcal{G}^* . The dots mark the geometric means, the plus signs the medians.

References

- Hirotsugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.
- Steen A. Andersson, David Madigan, and Michael D. Perlman. A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25(2):505–541, 1997.
- Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650, 2014.
- Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, and Linda Zhao. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- M. Alan Brookhart, Sebastian Schneeweiss, Kenneth J. Rothman, Robert J. Glynn, Jerry Avorn, and Til Stürmer. Variable selection for propensity score models. *American Journal of Epidemiology*, 163(12):1149–1156, 2006.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov):507–554, 2002.
- A. Philip Dawid and Vanessa Didelez. Identifying the consequences of dynamic treatment strategies: A decision-theoretic overview. *Statistics Surveys*, 4:184–231, 2010.
- Xavier de Luna, Ingeborg Waernbaum, and Thomas S Richardson. Covariate selection for the nonparametric estimation of an average treatment effect. *Biometrika*, 98(4):861–875, 2011.
- DeWayne Derryberry, Ken Aho, John Edwards, and Teri Peterson. Model selection and regression t-statistics. *The American Statistician*, 72(4):379–381, 2018.
- Oliver Dukes and Stijn Vansteelandt. How to obtain valid tests and confidence intervals after propensity score variable selection? *Statistical Methods in Medical Research*, 29(3):677–694, 2020a.
- Oliver Dukes and Stijn Vansteelandt. Inference on treatment effect parameters in potentially misspecified high-dimensional models. *Biometrika*, 2020b.
- Julia C. Engelmann, Thomas Amann, Birgitta Ott-Rötzer, Margit Nützel, Yvonne Reinders, Jörg Reinders, Wolfgang E. Thasler, Theresa Kristl, Andreas Teufel, Christian G. Huber, Peter J. Oefner, Rainer Spang, and Claus Hellerbrand. Causal modeling of cancer-stromal communication identifies PAPP A as a novel stroma-secreted factor activating NF κ B signaling in hepatocellular carcinoma. *PLoS Computational Biology*, 11(5):e1004293, 2015.
- Sander Greenland and Neil Pearce. Statistical foundations for model-based adjustments. *Annual Review of Public Health*, 36:89–108, 2015.

- F. Richard Guo and Emilija Perković. Efficient least squares for estimating total effects under linearity and causal sufficiency. *arXiv preprint arXiv:2008.03481v2*, 2020.
- Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 66:315–331, 1998.
- Frank E. Harrell, Jr. *Regression modeling strategies*. Springer, 5th edition, 2010.
- Leonard Henckel, Emilija Perković, and Marloes H. Maathuis. Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *arXiv preprint arXiv:1907.02435*, 2019.
- Keisuke Hirano, Guido W. Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- Markus Kalisch, Martin Mächler, Diego Colombo, Marloes H. Maathuis, and Peter Bühlmann. Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47(11):1–26, 2012.
- Markus Kalisch, Alain Hauser, Martin Maechler, Diego Colombo, Doris Entner, Patrik Hoyer, Antti Hyttinen, Jonas Peters, Nicoletta Andri, Emilija Perkovic, Preetam Nandy, Philipp Ruetimann, Daniel Stekhoven, Manuel Schuerch, and Marco Eigenmann. *pcalg: Methods for Graphical Models and Causal Inference*, 2019. URL <https://CRAN.R-project.org/package=pcalg>. R package version 2.6-6.
- David G. Kleinbaum and Lawrence L. Kupper. *Applied regression analysis and other multivariable methods*. Duxbury Press, 1978.
- Sven Knüppel and Andreas Stang. DAG program: Identifying minimal sufficient adjustment sets. *Epidemiology*, 21(1):159, 2010.
- Manabu Kuroki and Zhihong Cai. Selection of identifiability criteria for total effects by using path diagrams. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI-04)*, pages 333–340, 2004.
- Manabu Kuroki and Masami Miyakawa. Covariate selection for estimating the causal effect of control plans by using causal diagrams. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1):209–222, 2003.
- Thuc Duy Le, Lin Liu, Anna Tsykin, Gregory J Goodall, Bing Liu, Bing-Yu Sun, and Jiuyong Li. Inferring microRNA–mRNA causal regulatory relationships from expression data. *Bioinformatics*, 29(6):765–771, 2013.
- Jason D. Lee, Dennis L. Sun, Yuekai Sun, and Jonathan E. Taylor. Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927, 2016.
- Hannes Leeb and Benedikt M. Pötscher. Can one estimate the unconditional distribution of post-model-selection estimators? *Econometric Theory*, 24(2):338–376, 2008.

- Lexin Li, R. Dennis Cook, and Christopher J. Nachtsheim. Model-free variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):285–299, 2005.
- Richard Lockhart, Jonathan Taylor, Ryan J. Tibshirani, and Robert Tibshirani. A significance test for the lasso. *Annals of Statistics*, 42(2):413–468, 2014.
- Karim Lounici. Sup-norm convergence rate and sign concentration property of Lasso and Dantzig estimators. *Electronic Journal of Statistics*, 2:90–102, 2008.
- Jared K. Lunceford and Marie Davidian. Stratification and weighting via the propensity score in estimation of causal treatment effects: A comparative study. *Statistics in Medicine*, 23(19):2937–2960, 2004.
- Jiawei Luo, Wei Huang, and Buwen Cao. A novel approach to identify the miRNA-mRNA causal regulatory modules in cancer. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 15(1):309–315, 2018.
- Marloes H. Maathuis and Diego Colombo. A generalised back-door criterion. *The Annals of Statistics*, 43(3):1060–1088, 2015.
- Marloes H. Maathuis, Markus Kalisch, and Peter Bühlmann. Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics*, 37(6A):3133–3164, 2009.
- Marloes H. Maathuis, Diego Colombo, Markus Kalisch, and Peter Bühlmann. Predicting causal effects in large-scale systems from observational data. *Nature Methods*, 7(4):247–248, 2010.
- Christopher Meek. Causal inference and causal explanation with background knowledge. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 403–410, 1995.
- Douglas C. Montgomery, Elizabeth A. Peck, and C. Geoffrey Vining. *Introduction to linear regression analysis*. John Wiley & Sons, 5th edition, 2012.
- Paul A. Murtaugh. In defense of p values. *Ecology*, 95(3):611–617, 2014.
- Preetam Nandy, Marloes H. Maathuis, and Thomas S. Richardson. Estimating the effect of joint interventions from observational data in sparse high-dimensional settings. *The Annals of Statistics*, 45(2):647–674, 2017.
- Judea Pearl. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI-01)*, pages 411–420, 2001.
- Judea Pearl. *Causality: Models, reasoning, and inference*. Cambridge University Press, 2nd edition, 2009.
- Emilija Perković. Identifying causal effects in maximally oriented partially directed acyclic graphs. In *Proceedings of the Thirty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI-20)*, page ID: 229, 2020.

- Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. A complete generalized adjustment criterion. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (UAI-15)*, pages 682–691, 2015.
- Emilija Perković, Markus Kalisch, and Maloes H. Maathuis. Interpreting and using CPDAGs with background knowledge. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence (UAI-17)*, page ID: 120, 2017.
- 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.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019. URL <https://www.R-project.org/>.
- Joseph D. Ramsey. A scalable conditional independence test for nonlinear, non-gaussian data. *arXiv preprint arXiv:1401.5031v2*, 2014.
- Thomas Richardson. Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1):145–157, 2003.
- Thomas S. Richardson, Robin J. Evans, James M. Robins, and Ilya Shpitser. Nested Markov properties for acyclic directed mixed graphs. *arXiv preprint arXiv:1701.06686*, art. arXiv:1701.06686, 2017.
- Alessandro Rinaldo, Larry Wasserman, and Max G’Sell. Bootstrapping and sample splitting for high-dimensional, assumption-lean inference. *The Annals of Statistics*, 47(6):3438–3469, 2019.
- James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9–12):1393–1512, 1986.
- James M. Robins and Sander Greenland. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(2):143–155, 1992.
- Andrea Rotnitzky and Ezequiel Smucler. Efficient adjustment sets for population average treatment effect estimation in non-parametric causal graphical models. *Journal of Machine Learning Research*, 21(188):1–86, 2020.
- Gideon Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.
- Rajen D. Shah and Jonas Peters. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3):1514–1538, 2020.
- Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7:2003–2030, 2006.

- Susan M. Shortreed and Ashkan Ertefaie. Outcome-adaptive lasso: Variable selection for causal inference. *Biometrics*, 73(4):1111–1122, 2017.
- Ilya Shpitser, Tyler VanderWeele, and James M. Robins. On the validity of covariate adjustment for estimating causal effects. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI-10)*, pages 527–536, 2010.
- Ilya Shpitser, Robin J. Evans, Thomas S. Richardson, and James M. Robins. Introduction to nested Markov models. *Behaviormetrika*, 41(1):3–39, 2014.
- Ezequiel Smucler, Andrea Rotnitzky, and James M. Robins. A unifying approach for doubly-robust l_1 regularized estimation of causal contrasts. *arXiv preprint arXiv:1904.03737v3*, 2019.
- Ezequiel Smucler, Facundo Sapienza, and Andrea Rotnitzky. Efficient adjustment sets in causal graphical models with hidden variables. *arXiv preprint arXiv:1912.00306*, 2020.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.
- Johannes Textor and Maciej Liśkiewicz. Adjustment criteria in causal diagrams: An algorithmic perspective. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI-11)*, pages 681–688, 2011.
- Jin Tian and Judea Pearl. On the identification of causal effects. Technical Report R-290-L, Department of Computer Science, University, University of California, Los Angeles, 2003.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 58(1):267–288, 1996.
- Dirk van Kampen. The Schizotypic Syndrome Questionnaire (SSQ): Psychometrics, validation and norms. *Schizophrenia Research*, 84(2–3):305–322, 2006.
- Dirk van Kampen. The SSQ model of schizophrenic prodromal unfolding revised: An analysis of its causal chains based on the language of directed graphs. *European Psychiatry*, 29(7):437–448, 2014.
- TS Verma and Judea Pearl. Equivalence and synthesis of causal models. Technical Report R-150, Department of Computer Science, University, University of California, Los Angeles, 1990.
- Janine Witte and Vanessa Didelez. Covariate selection strategies for causal inference: Classification and comparison. *Biometrical Journal*, 61(5):1270–1289, 2019.
- Peng Zhao and Bin Yu. On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2563, 2006.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.