

# Efficient Least Squares for Estimating Total Effects under Linearity and Causal Sufficiency

**F. Richard Guo**

RICGUO@UW.EDU

**Emilija Perković**

PERKOVIC@UW.EDU

*Department of Statistics*

*University of Washington*

*Seattle, WA 98195-4322, USA*

**Editor:** Joris Mooij

## Abstract

Recursive linear structural equation models are widely used to postulate causal mechanisms underlying observational data. In these models, each variable equals a linear combination of a subset of the remaining variables plus an error term. When there is no unobserved confounding or selection bias, the error terms are assumed to be independent. We consider estimating a total causal effect in this setting. The causal structure is assumed to be known only up to a maximally oriented partially directed acyclic graph (MPDAG), a general class of graphs that can represent a Markov equivalence class of directed acyclic graphs (DAGs) with added background knowledge. We propose a simple estimator based on recursive least squares, which can consistently estimate any identified total causal effect, under point or joint intervention. We show that this estimator is the most efficient among all regular estimators that are based on the sample covariance, which includes covariate adjustment and the estimators employed by the joint-IDA algorithm. Notably, our result holds without assuming Gaussian errors.

**Keywords:** structural equation model, causal inference, semiparametric efficiency, partially directed acyclic graph, observational study

## 1. Introduction

A linear structural equation model (SEM) specifies a causal mechanism underlying a set of variables (Bollen, 1989). Each variable equals a linear combination of a subset of the remaining variables plus an error term. A SEM is *associated* with a mixed graph, also known as a path diagram (Wright, 1921, 1934), which consists of both directed edges and bi-directed edges. A directed edge  $i \rightarrow j$  represents that variable  $i$  appears as a covariate in the structural equation defining variable  $j$ . The equation for variable  $j$  takes the form

$$X_j = \sum_{i:i \rightarrow j} \gamma_{ij} X_i + \epsilon_j, \quad (1)$$

where  $\epsilon_j$  is an error term. Often, the errors are assumed to follow a multivariate normal distribution, but it need not be the case. A bi-directed edge  $i \leftrightarrow j$  indicates that errors  $\epsilon_i$  and  $\epsilon_j$  are dependent, which is assumed when there exists an unobserved (i.e., latent) confounder between  $i$  and  $j$ . The mixed graph is usually assumed to be *acyclic*, i.e., the graph does not contain cycles made of directed edges.

We focus on the setting when there is no unobserved confounder or selection bias, a condition also known as *causal sufficiency*; see Spirtes et al. (2000, Chap. 3) and Pearl (2009, Chap. 6). In this setting, all the error terms are assumed to be *mutually independent* and the mixed graph associated with the linear SEM is a *directed acyclic graph* (DAG), often called a *causal DAG*. Aside from being a statistical model for observational data, the linear SEM is also a causal model in the sense that it specifies the behavior of the system under interventions (see Section 3.2). Therefore, the total causal effect of one treatment variable (point intervention) or several treatment variables (joint intervention) on some outcome variables can be defined.

The underlying causal DAG is usually unknown. In fact, linear SEMs associated with different DAGs may define the same *observed* distribution (Drton et al., 2011). Without further assumptions on the error distributions, the underlying DAG can only be learned from observational data up to its Markov equivalence class, which can be uniquely represented by a completed partially directed acyclic graph (CPDAG) (Meek, 1995; Andersson et al., 1997). Additional background knowledge, such as knowledge of certain causal relationships (Meek, 1995; Fang and He, 2020) or partial orderings (Scheines et al., 1998), restrictions on the error distributions (Shimizu et al., 2006, 2011; Shimizu, 2014; Hoyer et al., 2008; Peters and Bühlmann, 2014), and other assumptions (Hauser and Bühlmann, 2012; Wang et al., 2017; Rothenhäusler et al., 2018; Eigenmann et al., 2017) can be used to further refine the Markov equivalence class of DAGs, resulting in representing the causal structure as a maximally oriented PDAG (MPDAG), which is a rather general class of graphs containing directed and undirected edges that subsumes DAGs and CPDAGs (Meek, 1995). A given total causal effect is *identified* given a graph, if it can be expressed as a functional of the observed distribution, which is the same for every DAG in the equivalence class. Recently, a necessary and sufficient graphical criterion for identification given an MPDAG has been shown by Perković (2020). In general, there may be more than one identifying functional.

Naturally, the next step is to develop estimators for an identified total effect with desirable properties. When the effect is unidentified, the reader is referred to IDA-type (Maathuis et al., 2009; Nandy et al., 2017) or enumerative (Guo and Perkovic, 2021) approaches, which are beyond the scope of this paper. Among others, we consider the following desiderata.

**Completeness.** Can the estimator consistently estimate every identified effect, under either point or joint interventions?

**Efficiency.** Does the estimator achieve the smallest asymptotic (co)-variance compared to a reasonably large class of estimators?

To the best of our knowledge, no estimator proposed in the literature fulfills both desiderata. Indeed, the commonly used covariate adjustment estimators (Pearl, 1993; Shpitser et al., 2010; Maathuis and Colombo, 2015; Perković et al., 2015) do not exist for certain total effects under joint interventions (Nandy et al., 2017; Perković et al., 2018; Perković, 2020). Furthermore, when they exist, even with an optimal adjustment set chosen to maximize efficiency (Henckel et al., 2022; Rotnitzky and Smucler, 2020; Witte et al., 2020), we will show in Section 7 that covariate adjustment can compare less favorably against a larger class of estimators.

We propose an estimator that is based on simple recursive least squares, that affirmatively fulfills both desiderata. In particular, our proposed estimator achieves the efficiency

bound among all regular estimators that only depend on the sample covariance; see Section 6 for the precise definition of the class of estimators. Remarkably, our result holds regardless of the type of error distribution in the underlying linear SEM. Our method is implemented in the R (R Core Team, 2020) package `eff2` (<https://cran.r-project.org/package=eff2>), which stands for “efficient effect” (estimate).

Admittedly, our estimator can be less efficient when compared to an even larger class of estimators, such as the class of all regular estimators considered in standard semiparametric theory. A semiparametric efficient estimator, relative to all regular estimators, can in principle be constructed by computing the efficient influence function and employing estimation strategies such as one-step correction or targeted maximum likelihood estimation (van der Laan and Rose, 2011). In fact, the semiparametric model we consider (see Eq. (21)) is a generalized, multivariate location-shift regression model with additional conditional independence constraints; see also Tsiatis (2006, §5.1) and Bickel et al. (1993, §4.3). While it is theoretically possible to construct a semiparametric efficient estimator by firstly estimating the error score and then solving the associated estimating equations (Bickel et al., 1993, §7.8), the resulting estimator tends to be too complicated and unstable for practical purposes unless the sample size is very large (Tsiatis, 2006, page 111). On the other hand, despite the potential loss of efficiency, our least squares estimator is easily computed and numerically stable. Hence, our proposal can be viewed as a deliberate trade-off between optimality and practicality.

The paper is organized as follows. In Section 2, we review related work on efficient estimation of total effects in over-identified settings. In Section 3, we introduce the preliminaries on linear structural equation models, causal graphs and the identification of total causal effects. The concept of bucket decomposition is introduced. In Section 4, we introduce a block-recursive representation for the observational data and identify the total causal effect under such a representation. We first derive the proposed least squares estimator by finding the maximum likelihood estimator (MLE) under the assumption of Gaussian errors in Section 5. We then prove the optimal efficiency of our proposed estimator under arbitrary error distributions in Section 6. Additional preliminaries, proofs and numerical results can be found in the Appendix.

## 2. Related work

The statistical performance of an estimator of a total causal effect, in *over-identified* settings, has recently received more attention; see, e.g, Kuroki and Miyakawa (2003); Henckel et al. (2022); Witte et al. (2020); Gupta et al. (2020); Rotnitzky and Smucler (2020); Smucler et al. (2020); Kuroki and Nanmo (2020). Here, “over-identified” (Koopmans and Reiersøl, 1950) refers to the fact that the total causal effect can be expressed as more than one functional of the (population) observed distribution, all of which coincide due to the additional conditional independence constraints obeyed by the observed distribution. For example, in the case where a total causal effect can be identified through covariate adjustment, usually there exists more than one valid adjustment set (Henckel et al., 2022). This is in contrast to the more traditional setting of causal inference, where the observed data distribution is nonparametric and is not expected to satisfy extra conditional independences.

Intuitively, the conditional independences in over-identified models can be exploited to maximize asymptotic efficiency; see, e.g., Sargan (1958); Hansen (1982) for early works in this direction. Under a linear SEM with independent errors, a total causal effect can be estimated via covariate adjustment as the least squares coefficient from the regression of the outcome on the treatment and adjustment variables. Henckel et al. (2022) recently showed that, under a linear SEM with independent errors, a valid adjustment set that minimizes asymptotic variance, also referred to as the *optimal adjustment set*, can be graphically characterized; see also Witte et al. (2020) for further properties of such an optimal set. This result was generalized by Rotnitzky and Smucler (2020) beyond linear SEMs: an optimal adjustment set is shown to always exist for point interventions, and a semiparametric efficient estimator is developed for this case. Note that, while valid adjustment sets (called “time-independent” adjustment sets by Rotnitzky and Smucler (2020)) exist for point interventions (Perković, 2020, Proposition 4.2), they may not exist for joint interventions (Nandy et al., 2017; Perković et al., 2018; Perković, 2020).

Less is known about how to efficiently estimate the total causal effect of a joint intervention, at least in a generic fashion. For linear SEMs with independent errors, with the knowledge of the parents of the treatment variables in the underlying causal DAG, Nandy et al. (2017) considered two estimators for the joint-IDA algorithm, one based on recursive least squares and one based on a modified Cholesky decomposition. However, the efficiency properties of these estimators were not explored. In Section 7, numerical comparisons will show that our proposed estimator significantly outperforms these estimators.

Other results on the linear SEM include explicit calculations and comparisons for typical examples with either a particular structure or only a few variables; see, e.g., Kuroki and Cai (2004); Gupta et al. (2020). Gaussian errors are also assumed in these calculations.

### 3. Linear SEMs, causal graphs and effect identification

#### 3.1 Linear SEMs under causal sufficiency

A linear SEM postulates a causal mechanism that generates data. Let  $X$  denote a vector of variables generated by a linear SEM, where  $X$  is indexed by  $V$  ( $X = X_V$ ). Let  $\mathcal{D}$  be the associated DAG on vertices  $V$ . For this  $|V|$ -dimensional random vector  $X$ , the model in Eq. (1) can be compactly rewritten as

$$X = \Gamma^\top X + \epsilon, \quad \Gamma = (\gamma_{ij}), \quad i \rightarrow j \text{ not in } \mathcal{D} \Rightarrow \gamma_{ij} = 0. \quad (2)$$

where  $\Gamma \in \mathbb{R}^{|V| \times |V|}$  is a coefficient matrix, and  $\epsilon = (\epsilon_i)$  is a  $|V|$ -dimensional random vector. DAG  $\mathcal{D}$  is associated with the linear SEM in Eq. (1) in the sense that the non-zero entries of  $\Gamma$  correspond to the edges in  $\mathcal{D}$ .

Under causal sufficiency (no latent variables), we assume

$$\{\epsilon_i : i \in V\} \text{ are independent, } \mathbb{E} \epsilon = 0, \quad \mathbb{E} \epsilon \epsilon^\top \succ \mathbf{0}, \quad (3)$$

where for a real symmetric matrix  $A$ ,  $A \succ \mathbf{0}$  means  $A$  is positive definite. The errors  $\{\epsilon_i : i \in V\}$  are not necessarily Gaussian, nor identically distributed.

The law  $P(X)$  is called the *observed distribution*. For a given  $\mathcal{D}$ , we will use  $\mathcal{P}_{\mathcal{D}}$  to denote the set of possible laws of  $X$ , namely the collection of  $P(X)$  as  $\Gamma$  and the error distribution

vary subject to Eqs. (2) and (3). The linear SEM poses certain restrictions on the set of laws  $\mathcal{P}_{\mathcal{D}}$ . Let  $\text{Pa}(i, \mathcal{D})$  denote the set of parents of vertex  $i$ , i.e.,  $\{j : j \rightarrow i \text{ is in } \mathcal{D}\}$ . For any  $P \in \mathcal{P}_{\mathcal{D}}$ , among other constraints, (i)  $P$  factorizes according to  $\mathcal{D}$ , (ii)  $\mathbb{E}[X_i | X_{\text{Pa}(i, \mathcal{D})}]$  is linear in  $X_{\text{Pa}(i, \mathcal{D})}$  and (iii)  $\text{var}[X_i | X_{\text{Pa}(i, \mathcal{D})}]$  is constant in  $X_{\text{Pa}(i, \mathcal{D})}$ .

We observe  $n$  iid samples generated by the model above, namely  $X^{(i)} = (I - \Gamma)^{-\top} \epsilon^{(i)}$  for  $i = 1, \dots, n$ . Note that  $(I - \Gamma)$  is invertible because  $\Gamma$  can be permuted into a lower-triangular matrix by a topological ordering (i.e., causal ordering) of vertices in  $\mathcal{D}$ .

### 3.2 Interventions and total causal effects

The assumed linear SEM also dictates the behavior of the system under interventions. Let  $A \subseteq V$  be a set of vertices indexing treatment variables  $X_A$ . We use  $\text{do}(X_A = x_A)$  to denote *intervening* on variables  $X_A$  and forcing them to take values  $x_A$  (Pearl, 1995). We call this a *point* intervention if  $A$  is a singleton, and a *joint* intervention if  $A$  consists of several vertices, which correspond to the case of multiple treatments. While  $X_A$  is fixed to  $x_A$ , the remaining variables are generated by their corresponding structural equations Eq. (1), with each  $X_i$  for  $i \in A$  appearing in the equations replaced by the corresponding enforced value  $x_i$  (Strotz and Wold, 1960). This generating mechanism defines the *interventional distribution*, denoted by  $P(X | \text{do}(X_A = x_A))$ , where the conditional probability notation is only conventional. More formally, the interventional distribution is expressed as

$$P(X | \text{do}(X_A = x_A)) = \prod_{j \in A} \delta_{x_j}(X_j) \prod_{i \notin A} P(X_i | X_{\text{Pa}(i, \mathcal{D})}), \quad (4)$$

where  $\delta$  denotes a Dirac measure. Factor  $P(X_i | X_{\text{Pa}(i)})$  is defined by the structural equation for  $X_i$ . Eq. (4) is known as the truncated factorization formula (Pearl, 2009), manipulated density formula (Spirtes et al., 2000) or the g-formula (Robins, 1986).

**Definition 1 (Total causal effect, Pearl, 2009; Nandy et al., 2017)** *Let  $X_A$  be a vector of treatment variables and  $X_Y$  with  $Y \in V \setminus A$  be an outcome variable. The total causal effect of  $X_A$  on  $X_Y$  is defined as the vector  $\tau_{AY} \in \mathbb{R}^{|A|}$ , where*

$$(\tau_{AY})_i = \frac{\partial}{\partial x_{A_i}} \mathbb{E}[X_Y | \text{do}(X_A = x_A)], \quad i = 1, \dots, |A|.$$

That is,  $\tau_{AY}$  is the gradient of the linear map  $x_A \mapsto \mathbb{E}[X_Y | \text{do}(X_A = x_A)]$ . When multiple outcomes  $Y = \{Y_1, \dots, Y_k\}$ ,  $k > 1$ , are considered, the total causal effect of  $X_A$  on  $X_{Y_1}, \dots, X_{Y_k}$  can be defined by concatenating  $\tau_{AY_1}, \dots, \tau_{AY_k}$ . Therefore, throughout, we assume the outcome variable is a singleton without loss of generality. Each coordinate of the total causal effect  $\tau_{AY}$  can be expressed as a sum-product of the underlying linear SEM coefficients along certain causal paths from  $A$  to  $Y$  in  $\mathcal{D}$ , that is, certain paths of the form  $A_1 \rightarrow \dots \rightarrow Y_i$  for  $A_1 \in A$ ; see also Wright (1934); Sullivant et al. (2010).

### 3.3 Causal graphs

Two different linear SEMs on the same set of variables can define the same observed distribution. For example, under Gaussian errors, linear SEMs associated with DAGs  $A \rightarrow Y$  and

$A \leftarrow Y$ , define the same set of observed distributions, namely the set of centered bivariate Gaussian distributions. Without making additional assumptions on the error distribution, such as non-Gaussianity (Shimizu et al., 2006), partial non-Gaussianity (Hoyer et al., 2008), or equal variance of errors (Peters and Bühlmann, 2014; Chen et al., 2019), the underlying causal DAG can only be learned from the observed distribution up to its Markov equivalence class (Pearl and Verma, 1995; Chickering, 2002).

**CPDAGs** Two DAGs on the same set of vertices are Markov equivalent if they encode the same set of d-separation relations between the vertices. The d-separations between the vertices, prescribe conditional independences between the corresponding variables (known as the Markov condition (Lauritzen, 1996, §3.2.2)); see Appendix C for the definition of d-separation and more background. This equivalence relation defines a Markov equivalence class, which consists of DAGs as elements. A Markov equivalence class can be uniquely represented by a *completed partially directed acyclic graph* (CPDAG), also known as an *essential graph* (Meek, 1995; Andersson et al., 1997). A CPDAG  $\mathcal{C}$  is a graph on the same set of vertices, that can contain both directed and undirected edges. We use  $[\mathcal{C}]$  to denote the Markov equivalence class represented by CPDAG  $\mathcal{C}$ . A directed edge  $i \rightarrow j$  in  $\mathcal{C}$  implies  $i \rightarrow j$  is in every  $\mathcal{D} \in [\mathcal{C}]$ , whereas an undirected edge  $i - j$  in  $\mathcal{C}$  implies there exist  $\mathcal{D}_1, \mathcal{D}_2 \in [\mathcal{C}]$  such that  $i \rightarrow j$  in  $\mathcal{D}_1$  but  $i \leftarrow j$  in  $\mathcal{D}_2$ . Given a DAG  $\mathcal{D}$ , the CPDAG  $\mathcal{C}$  representing the Markov equivalence class of  $\mathcal{D}$  can be drawn by keeping the skeleton of  $\mathcal{D}$ , adding all the unshielded colliders from  $\mathcal{D}$  and completing the orientation rules R1–R3 of Meek (1995); see Fig. 6 in Appendix C. For example, DAGs  $A \rightarrow Y$  and  $A \leftarrow Y$  are represented by CPDAG  $A - Y$ . To slightly abuse the notation, for a distribution  $Q$ , we write  $Q \in [\mathcal{C}]$  if  $Q$  factorizes according to some DAG  $\mathcal{D} \in [\mathcal{C}]$ ; see Lauritzen (1996, §3.2.2).

There are various *structure learning* algorithms that can be used to uncover CPDAG  $\mathcal{C}$  from observational data. Some well-known examples are the PC algorithm (Spirtes et al., 2000) and the greedy equivalence search (Chickering, 2002). Choosing an appropriate algorithm for the data set at hand is beyond the scope of this paper; the reader is referred to Drton and Maathuis (2017, §4) for a recent overview.

**MPDAGs** Certain background knowledge, if present, can be used to further orient some undirected edges in a CPDAG  $\mathcal{C}$ . Typically, knowledge of temporal orderings can inform the orientation of certain undirected edges; see Spirtes et al. (2000, §5.8.4) for an example. Adding these background-knowledge orientations and the additionally implied orientations based on the orientation rules of Meek (1995) to  $\mathcal{C}$  results in a *maximally oriented partially directed graph* (MPDAG)  $\mathcal{G}$ . See Fig. 6 and Algorithm 1 in Appendix C. MPDAGs are a rather general class of graphs that subsumes both DAGs and CPDAGs. An MPDAG  $\mathcal{G}$  represents a *restricted* Markov equivalence class of DAGs, which we also denote by  $[\mathcal{G}]$ . Analogously to the case of a CPDAG,  $i \rightarrow j$  in  $\mathcal{G}$  implies  $i \rightarrow j$  is in every  $\mathcal{D} \in [\mathcal{G}]$ , and  $i - j$  in  $\mathcal{G}$  implies there exist  $\mathcal{D}_1, \mathcal{D}_2 \in [\mathcal{G}]$  such that  $i \rightarrow j$  in  $\mathcal{D}_1$  but  $i \leftarrow j$  in  $\mathcal{D}_2$ .

For the rest of the paper, we will assume that we have access to an MPDAG  $\mathcal{G}$  that represents our structural knowledge about the underlying DAG  $\mathcal{D}$ . That is,

$$\text{causal DAG } \mathcal{D} \in [\mathcal{G}], \quad \mathcal{G} \text{ is an MPDAG,} \tag{5}$$

where  $[\mathcal{G}]$  represents a collection of DAGs that are Markov equivalent, but can be strictly smaller than the corresponding Markov equivalence class due to background knowledge.

### 3.4 Causal effect identification

Throughout the paper we will use the following notations. Given treatment variables  $X_A$  and an outcome variable  $X_Y$  such that  $Y \notin A$ , we are interested in learning the total causal effect  $\tau_{AY}$ . We assume that we have access to an MPDAG  $\mathcal{G}$ , and to observational data that are generated as iid samples from a linear SEM defined by Eqs. (2) and (3), where the causal DAG  $\mathcal{D}$  is in  $[\mathcal{G}]$ . Before estimation can be performed, we need to make sure that  $\tau_{AY}$  can be identified from observational data. That is, we need to ensure that  $\tau_{AY}$  can be expressed as a functional of the observed distribution that is the same for *every* DAG in  $[\mathcal{G}]$ . We have the following graphical criterion.

**Theorem 2 (Perković, 2020)** *The total causal effect  $\tau_{AY}$  of  $X_A$  on  $X_Y$  is identified given an MPDAG  $\mathcal{G}$  if and only if there is no proper, possibly causal path from  $A$  to  $Y$  in  $\mathcal{G}$  that starts with an undirected edge.*

Theorem 2 is Proposition 3.2 of Perković (2020), which holds for nonparametric causal graphical models. It does not require that the data is generated by a linear SEM. However, Perković (2020) proves that when the criterion fails, then two linear SEMs with Gaussian errors can be constructed such that their observed distributions coincide but their  $\tau_{AY}$ 's are different. Hence, even if we restrict ourselves to linear SEMs, Theorem 2 still holds.

A few terms need some explanation. A *path* from  $A$  to  $Y$  in  $\mathcal{G}$  is a sequence of distinct vertices  $\langle v_1, \dots, v_k \rangle$  for  $k > 1$  with  $v_1 \in A$  and  $v_k = Y$ , such that every pair of successive vertices are adjacent in  $\mathcal{G}$ . The path is *proper* when only its first vertex is in  $A$ . The path is *possibly causal* if no edge  $v_l \leftarrow v_r$  is in  $\mathcal{G}$  for  $1 \leq l < r \leq k$ . The reader is referred to Appendix C for more graphical preliminaries. When  $\mathcal{G}$  satisfies Theorem 2 relative to vertex sets  $A$  and  $Y$ , the interventional distribution  $P(X_Y | \text{do}(X_A = x_A))$ , and hence the total effect, can be computed from the observed distribution  $P(X)$ . To express the identification formula, we require the following concepts.

#### 3.4.1 BUCKETS AND BUCKET DECOMPOSITION

Let  $\mathcal{G} = (V, E, U)$  be a partially directed graph, where  $V$  is the set of vertices, and  $E$  and  $U$  are sets of directed and undirected edges respectively. Let  $B_1, \dots, B_K$  be the *maximal connected components* of the undirected graph  $\mathcal{G}_U := (V, \emptyset, U)$ . Then  $V = B_1 \dot{\cup} \dots \dot{\cup} B_K$ , where symbol  $\dot{\cup}$  denotes disjoint union. Note that all the directed edges within each  $B_i$  are due to background knowledge. If we ignore the distinction between directed and undirected edges, then the subgraph induced by each  $B_i$  is chordal (Andersson et al., 1997, §4).

Suppose the connected components are ordered such that

$$i \rightarrow j \in E, i \in B_i, j \in B_j \quad \Rightarrow \quad i < j. \quad (6)$$

One can show that such a *partial causal ordering* always exists, though it may not be unique; see Algorithm 2 in Appendix C to obtain such an ordering. Our result does not depend on the particular choice of partial causal ordering. We call  $B_1, \dots, B_K$  the *bucket decomposition* of  $V$  and call each  $B_k$  for  $k = 1, \dots, K$  a *bucket*; see Fig. 1(a) for an example. If it is clear which graph  $\mathcal{G}$  is being referred to, we will shorten  $\text{Pa}(j, \mathcal{G})$  as  $\text{Pa}(j)$  to reduce clutter. For a set of vertices  $C$  in  $\mathcal{G}$ , we use  $\text{Pa}(C) := \cup_{i \in C} \text{Pa}(i) \setminus C$  to denote the set of their *external parents*. Clearly,  $\text{Pa}(B_k) \subseteq B_{[k-1]}$ , where  $B_{[k-1]} := B_1 \cup \dots \cup B_{k-1}$ .

**Lemma 3** *Let  $i$  and  $j$  be two distinct vertices in MPDAG  $\mathcal{G} = (V, E, U)$  such that  $i \rightarrow j \in E$ . Suppose that there is no undirected path from  $i$  to  $j$  in  $\mathcal{G}$ . If there is a vertex  $k$ , and an undirected path  $j - \dots - k$  in  $\mathcal{G}$ , then  $i \rightarrow k \in E$ .*

By definition of the parent set above we have that  $\text{Pa}(B_k) = \cup_{i \in B_k} \text{Pa}(i) \setminus B_k$ ,  $k = 1, \dots, K$ . However, since a bucket  $B_k$  is a maximal subset of  $V$  that is connected by undirected edges in  $\mathcal{G}$ , Theorem 3 implies the following important property.

**Corollary 4 (Restrictive property)** *Let  $B_1, \dots, B_K$  be the bucket decomposition of  $V$  in MPDAG  $\mathcal{G} = (V, E, U)$ . Then, all vertices in the same bucket have the same set of external parents, namely*

$$\text{Pa}(B_k) = \text{Pa}(i) \setminus B_k, \quad \text{for any } i \in B_k, k = 1, \dots, K.$$

The causal identification formula for  $P(X_Y | \text{do}(X_A = x_A))$  of Perković (2020) relies on a decomposition of certain *ancestors* of  $Y$  in MPDAG  $\mathcal{G}$  according to the buckets. We call vertex  $i$  an ancestor of vertex  $j$  in  $\mathcal{G}$  if there exists a directed path  $i \rightarrow \dots \rightarrow j$  in  $\mathcal{G}$ ; we use the convention that  $j$  is an ancestor of itself. We denote the set of ancestors of  $j$  in  $\mathcal{G}$  as  $\text{An}(j, \mathcal{G})$ , or shortened as  $\text{An}(j)$ .

Let  $\mathcal{G}_{V \setminus A} = (V \setminus A, E', U')$  denote the subgraph of  $\mathcal{G}$  induced by the vertices  $V \setminus A$ , where  $E'$  includes those edges in  $E$  that are between vertices in  $V \setminus A$ , and similarly for  $U'$ . Consider the set of ancestors of  $Y$  in  $\mathcal{G}_{V \setminus A}$ , denoted as

$$D := \text{An}(Y, \mathcal{G}_{V \setminus A}). \tag{7}$$

The bucket decomposition  $D_1, \dots, D_K$  of  $D$ , induced by the bucket decomposition of  $V$ , is simply

$$D = \bigcup_{k=1}^K D_k, \quad D_k = D \cap B_k, \quad i = 1, \dots, K. \tag{8}$$

**Lemma 5** *When the criterion in Theorem 2 is satisfied, we have  $\text{Pa}(D_k, \mathcal{G}) = \text{Pa}(B_k, \mathcal{G})$  for every nonempty  $D_k$ .*

Proofs of Theorems 3 and 5 are left to Appendix B.

**Theorem 6 (Perković, 2020)** *Suppose the criterion in Theorem 2 is satisfied for  $A, Y$  in MPDAG  $\mathcal{G} = (V, E, U)$  such that  $Y \notin A$ . Let  $P(X)$  be the observed distribution. Let  $D = \text{An}(Y, \mathcal{G}_{V \setminus A})$  and  $D_1, \dots, D_K$  be the bucket decomposition of  $D$  as in Eq. (8). Then the interventional distribution  $P(X_Y | \text{do}(X_A = x_A))$  can be identified as*

$$P(X_Y | \text{do}(X_A = x_A)) = \int \left\{ \prod_{k=1}^K P(X_{D_k} | X_{\text{Pa}(D_k)}) \right\} dX_{D \setminus Y} \tag{9}$$

for values  $X_{\text{Pa}(D_k)}$  in agreement with  $x_A$ , where  $P(X_{D_k} | X_{\text{Pa}(D_k)}) \equiv 1$  if  $D_k = \emptyset$ .

The expression in Eq. (9) above is a generalization of the truncated factorization Eq. (4) from DAGs to MPDAGs. Theorem 6 holds generally even when an underlying linear SEM is not assumed.



#### 4. Block-recursive representation

In this section, we express the observed distribution  $P(X)$  induced by a linear SEM compatible with MPDAG  $\mathcal{G} = (V, E, U)$  in a block-recursive form. Each block corresponds to a bucket in the bucket decomposition of  $V$ . Such a reparameterization is necessitated by the fact that the causal ordering of  $\mathcal{D}$  is unknown, whereas the buckets can be arranged into a valid partial causal ordering as in Eq. (6). We will use this representation to compute an estimator for the total causal effect.

Recall that  $\mathcal{P}_{\mathcal{D}}$  denotes the family of laws of  $X$  arising from a linear SEM Eqs. (2) and (3) compatible with DAG  $\mathcal{D}$ . Let  $\mathcal{P}_{\mathcal{G}} := \cup_{\mathcal{D} \in [\mathcal{G}]} \mathcal{P}_{\mathcal{D}}$ , which denotes the family of laws of  $X$  arising from a linear SEM compatible with a DAG in  $[\mathcal{G}]$ .

**Proposition 7 (Block-recursive form)** *Let  $\mathcal{D}$  be the causal DAG associated with the linear SEM and  $\mathcal{G}$  an MPDAG such that  $\mathcal{D} \in [\mathcal{G}]$ . Further, let  $B_1, \dots, B_K$  be the bucket decomposition of  $V$  in  $\mathcal{G}$ . Then the linear SEM Eqs. (2) and (3) can be rewritten as*

$$X = \Lambda^\top X + \varepsilon,$$

for some matrix of coefficients  $\Lambda = (\lambda_{ij}) \in \mathbb{R}^{|V| \times |V|}$  and random vector  $\varepsilon = (\varepsilon_i) \in \mathbb{R}^{|V|}$  such that

$$j \in B_l, i \notin \text{Pa}(B_l, \mathcal{G}) \Rightarrow \lambda_{ij} = 0, \quad (10)$$

$$\mathbb{E} \varepsilon = 0, \quad \mathbb{E} \varepsilon_{B_k} \varepsilon_{B_k}^\top \succ \mathbf{0}, \quad (k = 1, \dots, K), \quad \varepsilon_{B_1}, \dots, \varepsilon_{B_K} \text{ are mutually independent,} \quad (11)$$

and

$$\text{law of } (\varepsilon_{B_k}) \in \mathcal{P}_{\mathcal{G}_{B_k}}, \quad k = 1, \dots, K, \quad (12)$$

where  $\mathcal{G}_{B_k}$  is the subgraph of  $\mathcal{G}$  induced by  $B_k$ .

Note that in contrast to symbol  $\epsilon$  used in Eqs. (2) and (3), symbol  $\varepsilon$  is used here to denote the errors in the block-recursive form. The coordinates within each  $\varepsilon_{B_k}$  may be *dependent*.

**Proof** For  $k = 2, \dots, K$ , by Eq. (2) and the restrictive property (Theorem 4), we have

$$X_{B_k} = \Gamma_{\text{Pa}(B_k), B_k}^\top X_{\text{Pa}(B_k)} + \Gamma_{B_k}^\top X_{B_k} + \epsilon_{B_k},$$

where  $\text{Pa}(B_k) = \text{Pa}(B_k, \mathcal{G})$ . The expression can be rewritten as

$$\begin{aligned} X_{B_k} &= (I - \Gamma_{B_k})^{-\top} \Gamma_{\text{Pa}(B_k), B_k}^\top X_{\text{Pa}(B_k)} + (I - \Gamma_{B_k})^{-\top} \epsilon_{B_k} \\ &= \Lambda_{\text{Pa}(B_k), B_k}^\top X_{\text{Pa}(B_k)} + \varepsilon_{B_k}, \end{aligned}$$

where  $\varepsilon_{B_k} := (I - \Gamma_{B_k})^{-\top} \epsilon_{B_k}$  for  $k = 1, \dots, K$  (note that  $X_{B_1} = \varepsilon_{B_1}$ ). Additionally,  $\Lambda_{\text{Pa}(B_k), B_k} = \Gamma_{\text{Pa}(B_k), B_k} (I - \Gamma_{B_k})^{-1}$  for  $k = 2, \dots, K$ .

Matrix  $\Lambda \in \mathbb{R}^{|V| \times |V|}$  in the statement of the proposition is defined by blocks  $\Lambda_{\text{Pa}(B_k), B_k}$  for  $k = 2, \dots, K$  and zero entries otherwise. Therefore,  $\lambda_{ij} = 0$  if  $j \in B_l$  and  $i \notin \text{Pa}(B_l)$  for some  $l = 1, \dots, K$ . Hence, by putting the blocks together, the model can be written as  $X = \Lambda^\top X + \varepsilon$ .

The “new” errors  $\varepsilon$  satisfy

$$\varepsilon_{B_k} = \Gamma_{B_k}^\top \varepsilon_{B_k} + \epsilon_{B_k}, \quad k = 1, \dots, K.$$

It then follows from Eqs. (2) and (3) that for every  $k$ ,

$$\text{law of } \varepsilon_{B_k} \in \mathcal{P}_{\mathcal{D}_{B_k}} \subset \mathcal{P}_{\mathcal{G}_{B_k}},$$

since  $D \in [\mathcal{G}]$ . Moreover, for every  $k$ ,

$$\mathbb{E} \varepsilon_{B_k} = \mathbf{0}, \quad \mathbb{E} \varepsilon_{B_k} \varepsilon_{B_k}^\top = (I - \Gamma_{B_k})^{-\top} \mathbb{E} \varepsilon_{B_k} \varepsilon_{B_k}^\top (I - \Gamma_{B_k})^{-1} \succ \mathbf{0},$$

where both  $(I - \Gamma_{B_k})$  and  $\mathbb{E} \varepsilon_{B_k} \varepsilon_{B_k}^\top$  are full rank, because  $\Gamma_{B_k}$  can be permuted into an upper-triangular matrix and  $\mathbb{E} \varepsilon \varepsilon^\top \succ \mathbf{0}$  by Eq. (3).  $\blacksquare$

**Corollary 8** *Under the same conditions as Theorem 7, it holds that*

$$\begin{aligned} X_{B_1} &= \varepsilon_{B_1}, \\ X_{B_k} &= \Lambda_{\text{Pa}(B_k), B_k}^\top X_{\text{Pa}(B_k)} + \varepsilon_{B_k}, \quad \varepsilon_{B_k} \perp\!\!\!\perp X_{\text{Pa}(B_k)}, \quad k = 2, \dots, K, \end{aligned} \quad (13)$$

where  $\text{Pa}(B_k) = \text{Pa}(B_k, \mathcal{G})$ .

Next, we show that if the total causal effect  $\tau_{AY}$  is identifiable from MPDAG  $\mathcal{G}$  (Theorem 2), then it can be calculated from  $\Lambda$  in the block-recursive representation of Proposition 7. Therefore, the distribution of  $\varepsilon$  is a nuisance relative to estimating  $\tau_{AY}$ .

**Proposition 9** *Suppose the criterion in Theorem 2 is satisfied for  $A, Y$  in MPDAG  $\mathcal{G} = (V, E, U)$  such that  $Y \notin A$ . Let  $\Lambda$  be the block-recursive coefficient matrix given by Theorem 7. The total causal effect of  $X_A$  on  $X_Y$  is identified as*

$$\tau_{AY} = \Lambda_{A,D} [(I - \Lambda_{D,D})^{-1}]_{D,Y}, \quad (14)$$

where  $D = \text{An}(Y, \mathcal{G}_{V \setminus A})$  and the last subscript denotes the column corresponding to  $Y \in D$ .

**Proof** We derive this result using Theorem 6. Recall that  $D_1, \dots, D_K$  is a partition of  $D$  induced by the bucket decomposition  $B_1, \dots, B_K$  of  $V$  in the sense that  $D_k = D \cap B_k$  for  $k = 1, \dots, K$ . When  $D_k = \emptyset$ , we use the convention that  $P(X_{D_k} | X_{\text{Pa}(D_k)}) \equiv 1$ . By definition of  $D = \text{An}(Y, \mathcal{G}_{V \setminus A})$  and Eq. (6), observe that a vertex in  $\text{Pa}(D_k) = \text{Pa}(D_k, \mathcal{G})$  is either in  $D_1 \cup \dots \cup D_{k-1}$  or in  $A$ . Let  $F_k := A \cap \text{Pa}(D_k)$ . In Eq. (9), we note that the joint interventional distribution of  $X_D$  is given by

$$P(X_D | \text{do}(X_A = x_A)) = \prod_{k=1}^K P(X_{D_k} | X_{\text{Pa}(D_k)}) = \prod_{k=1}^K P(X_{D_k} | X_{\text{Pa}(D_k) \setminus F_k}, X_{F_k} = x_{F_k}),$$

where  $x_{F_k}$  is fixed by the  $\text{do}(X_A = x_A)$  operation. Further, fix a factor  $i \in \{1, \dots, K\}$ . By Theorem 5,  $\text{Pa}(D_i) = \text{Pa}(B_i)$ . By Eq. (13) and  $\varepsilon_{D_i} \perp\!\!\!\perp X_{\text{Pa}(B_i)}$ , we have

$$\begin{aligned} X_{D_i} | \{X_{\text{Pa}(D_i) \setminus F_i}, X_{F_i} = x_{F_i}\} &= \Lambda_{\text{Pa}(D_i) \setminus F_i, D_i}^\top X_{\text{Pa}(D_i) \setminus F_i} + \Lambda_{F_i, D_i} x_{F_i} + \varepsilon_{D_i} \\ &= \Lambda_{\text{Pa}(D_i) \cap D, D_i}^\top X_{\text{Pa}(D_i) \cap D} + \Lambda_{\text{Pa}(D_i) \cap A, D_i} x_{\text{Pa}(D_i) \cap A} + \varepsilon_{D_i} \end{aligned}$$

The fact that the display above holds for every  $i = 1, \dots, K$  implies that the joint interventional distribution  $P(X_D | \text{do}(X_A = x_A))$  satisfies

$$X_D = \Lambda_{D,D}^T X_D + \Lambda_{A,D}^T x_A + \varepsilon_D.$$

It follows that  $X_D = (I - \Lambda_{D,D})^{-\top} (\Lambda_{A,D}^T x_A + \varepsilon_D)$  and hence

$$\mathbb{E}[X_D | \text{do}(X_A = x_A)] = (I - \Lambda_{D,D})^{-\top} \Lambda_{A,D}^T x_A.$$

Since  $Y \in D$ , by Theorem 1 we have

$$\tau_{AY} = \frac{\partial}{\partial x_A} \mathbb{E}[X_Y | \text{do}(X_A = x_A)] = \Lambda_{A,D} [(I - \Lambda_{D,D})^{-1}]_{D,Y}.$$

■

We say vertex  $j$  is a possible descendant of  $i$ , denoted as  $j \in \text{PossDe}(i)$ , if there exists a possibly causal path from  $i$  to  $j$ . For a set of vertices  $A$ , define  $\text{PossDe}(A) := \cup_{i \in A} \text{PossDe}(i)$ . See Appendix C for more details.

**Corollary 10** *If  $Y \notin \text{PossDe}(A)$ , then  $\tau_{AY} = 0$ .*

**Proof** Since  $D = \text{An}(Y, \mathcal{G}_{V \setminus A})$  and  $Y \notin \text{PossDe}(A)$ ,  $\Lambda_{A,D} = \mathbf{0}$ . ■

## 5. Recursive least squares

Consider the *special case* when the errors in the linear SEM Eq. (1) are jointly Gaussian. In this case, by the standard maximum likelihood theory, the Cramér–Rao bound is achieved by the maximum likelihood estimator (MLE) of the total causal effect, which can be obtained by plugging in the MLE for  $\Lambda$  in the block-recursive form (Theorem 7) into the formula Eq. (14). We now compute the MLE for  $\Lambda$  given an MPDAG  $\mathcal{G}$ .

When  $\varepsilon$  is multivariate Gaussian, the block-recursive form in Theorem 7 is a linear Gaussian model parameterized by  $\{(\Lambda_k)_{k=2}^K, (\Omega_k)_{k=1}^K\}$ , where  $\Lambda_k := \Lambda_{\text{Pa}(B_k), B_k}$  and  $\Omega_k$  is the covariance for  $\varepsilon_{B_k}$ . Because  $\varepsilon$  are independent between blocks (Proposition 7), the likelihood factorizes as

$$\mathcal{L}((\Lambda_k)_k, (\Omega_k)_k) = \prod_{k=1}^K \mathcal{N}(X_{B_k} - \Lambda_k^\top X_{\text{Pa}(B_k)}; \mathbf{0}, \Omega_k). \quad (15)$$

Denote the MLE of  $\Lambda$  by  $\widehat{\Lambda}^{\mathcal{G}}$ , which consists of blocks  $(\widehat{\Lambda}_k^{\mathcal{G}})_{k=2}^K$  and zero values elsewhere, and the MLE of  $\Omega$  by  $\widehat{\Omega}^{\mathcal{G}} = (\widehat{\Omega}_k^{\mathcal{G}})_{k=1}^K$ . The superscripts highlight the dependence on MPDAG  $\mathcal{G}$ . The MLE maximizes  $\mathcal{L}((\Lambda_k)_k, (\Omega_k)_k)$  subject to Eq. (12), namely

$$\mathcal{N}(\mathbf{0}, \Omega_k) \in \mathcal{P}_{\mathcal{G}_{B_k}}, \quad k = 1, \dots, K,$$

where  $\mathcal{G}_{B_k}$  is the subgraph of  $\mathcal{G}$  induced by  $B_k$ . This further translates to a set of algebraic constraints on  $(\Omega_k)_{k=1}^K$ , namely for  $k = 1, \dots, K$ ,

$$\det [(\Omega_k)_{\{i\} \cup C, \{j\} \cup C}] = 0, \text{ if } i \text{ and } j \text{ are d-separated by } C \text{ in } \mathcal{G}_{B_k}; \quad (16)$$

see, e.g., Drton et al. (2008, §3.1). Although the constraints Eq. (16) may seem daunting, we will show that they do not affect the MLE for  $\Lambda$ .

Let the sample covariance matrix be computed with respect to mean zero, i.e.,

$$\widehat{\Sigma}^{(n)} := \frac{1}{n} \sum_{i=1}^n X^{(i)} X^{(i)\top}, \quad (17)$$

where  $n$  is the sample size, and the superscripts are reserved to index samples. To reduce clutter, for a set of indices  $C$ , we often abbreviate  $\Sigma_{C,C}$  as  $\Sigma_C$ .

**Lemma 11** *Suppose  $X^{(i)} : i = 1, \dots, n$  is generated iid from a linear SEM Eqs. (2) and (3) associated with an unknown causal DAG  $\mathcal{D}$ . Suppose the error  $\epsilon$  is distributed as multivariate Gaussian. Suppose  $\mathcal{D} \in [\mathcal{G}]$  for a known MPDAG  $\mathcal{G}$ . Let  $\widehat{\Sigma}^{(n)}$  be the sample covariance as defined in Eq. (17). The MLE for  $\Lambda_k = \Lambda_{\text{Pa}(B_k), B_k}$  in the block-recursive form is given by*

$$\widehat{\Lambda}_k^{\mathcal{G}} = \left( \widehat{\Sigma}_{\text{Pa}(B_k)}^{(n)} \right)^{-1} \widehat{\Sigma}_{\text{Pa}(B_k), B_k}^{(n)}, \quad k = 2, \dots, K. \quad (18)$$

**Proof** By factorization in Eq. (15), MLE  $(\widehat{\Lambda}_k^{\mathcal{G}}, \widehat{\Omega}_k^{\mathcal{G}})$  is the maximizer of log-likelihood

$$\begin{aligned} & \ell_n(\Lambda_k, \Omega_k) \\ &= -\frac{1}{2} \sum_{i=1}^n \left( X_{B_k}^{(i)} - \Lambda_k^\top X_{\text{Pa}(B_k)}^{(i)} \right)^\top \Omega_k^{-1} \left( X_{B_k}^{(i)} - \Lambda_k^\top X_{\text{Pa}(B_k)}^{(i)} \right) - \frac{n}{2} \log \det(\Omega_k) \\ &= -\frac{1}{2} \text{Tr} \left( \sum_{i=1}^n \Omega_k^{-1} (X_{B_k}^{(i)} - \Lambda_k^\top X_{\text{Pa}(B_k)}^{(i)}) (X_{B_k}^{(i)} - \Lambda_k^\top X_{\text{Pa}(B_k)}^{(i)})^\top \right) - \frac{n}{2} \log \det(\Omega_k), \end{aligned}$$

subject to Eq. (16). Taking a derivative with respect to  $\Lambda_k \in \mathbb{R}^{|\text{Pa}(B_k)| \times |B_k|}$ , we have

$$\frac{\partial \ell_n(\Lambda_k, \Omega_k)}{\partial \Lambda_k} = -2 \sum_{i=1}^n X_{\text{Pa}(B_k)}^{(i)} X_{B_k}^{(i)\top} \Omega_k^{-1} + 2 \sum_{i=1}^n X_{\text{Pa}(B_k)}^{(i)} X_{\text{Pa}(B_k)}^{(i)\top} \Lambda_k \Omega_k^{-1}.$$

For any positive definite  $\Omega_k$  satisfying Eq. (16), setting the derivative  $\ell_n(\Lambda_k, \Omega_k)/\partial \Lambda_k$  to zero yields the estimate

$$\widehat{\Lambda}_k^{\mathcal{G}} = \left( \frac{1}{n} \sum_{i=1}^n X_{\text{Pa}(B_k)}^{(i)} X_{\text{Pa}(B_k)}^{(i)\top} \right)^{-1} \left( \frac{1}{n} \sum_{i=1}^n X_{\text{Pa}(B_k)}^{(i)} X_{B_k}^{(i)\top} \right) = \left( \widehat{\Sigma}_{\text{Pa}(B_k)}^{(n)} \right)^{-1} \widehat{\Sigma}_{\text{Pa}(B_k), B_k}^{(n)}. \quad \blacksquare$$

**Remark 12** *Because of the restrictive property (Theorem 4), each  $\widehat{\Lambda}_k^{\mathcal{G}}$  is computed by optimizing over the space of  $|\text{Pa}(B_k)| \times |B_k|$  matrices and the resulting MLE takes the simple form as above; see also Anderson and Olkin (1985, §5) and Amemiya (1985, §6.4) for earlier discussions of this phenomenon.*

*However, such a simple form is unavailable in general, when the zero constraints on  $\Lambda$  do not obey the restrictive property, even if we ignore the algebraic constraints Eq. (16) on  $\Omega$ . In fact, the likelihood function can be multimodal; see also Drton and Richardson (2004); Drton (2006); Drton et al. (2009) on seemingly unrelated regressions.*

Since  $\widehat{\Lambda}^{\mathcal{G}}$  is obtained by simply regressing each  $B_i$  onto  $\text{Pa}(B_i, \mathcal{G})$  using ordinary least squares, we call this specific recursive least squares  $\mathcal{G}$ -regression. The resulting MLE for an identified total causal effect is a plugin estimator using the formula in Theorem 9.

**Definition 13 ( $\mathcal{G}$ -regression estimator)** *Suppose  $X^{(i)} : i = 1, \dots, n$  is generated iid from a linear SEM Eqs. (2) and (3) associated with an unknown causal DAG  $\mathcal{D}$ . Suppose  $\mathcal{D} \in [\mathcal{G}]$  for a known MPDAG  $\mathcal{G}$ . Further, suppose for  $A \subset V$ ,  $Y \in V \setminus A$ ,  $\tau_{AY}$  is identified under the criterion of Theorem 2. The  $\mathcal{G}$ -regression estimator for the total causal effect  $\tau_{AY}$  is defined as*

$$\widehat{\tau}_{AY}^{\mathcal{G}} = \widehat{\Lambda}_{A,D}^{\mathcal{G}} \left[ (I - \widehat{\Lambda}_{D,D}^{\mathcal{G}})^{-1} \right]_{D,Y}, \quad (19)$$

where  $\widehat{\Lambda}^{\mathcal{G}}$  is given by Eq. (18).

## 6. Efficiency theory

In this section, we establish the asymptotic efficiency of our  $\mathcal{G}$ -regression estimator, when the errors in the generating linear SEM are *not* necessarily Gaussian, among a reasonably large class of estimators—all regular estimators that only depend on the sample covariance. This class of estimators, despite not covering all the estimators considered in the standard semiparametric efficiency theory, includes many in the literature:

1. Optimal covariate adjustment (Henckel et al., 2022; Witte et al., 2020):  $\widehat{\tau}_{AY}^{\text{adj}\cdot 0}$  is the coefficient of  $A$  in the least squares regression of  $Y \sim A + O$ , where  $O$  is the optimal adjustment set, which minimizes asymptotic variance among all valid adjustment sets.
2. Recursive least squares (Nandy et al., 2017; Gupta et al., 2020): Under point intervention ( $|A| = 1$ ), total effect  $\widehat{\tau}_{AY}$  is taken as the coefficient of  $A$  in the least square regression of  $Y \sim A + \text{Pa}(A)$ . Nandy et al. (2017, §3.1) consider a recursive relation that expresses  $\widehat{\tau}_{AY}$  under joint intervention ( $|A| > 1$ ) as a polynomial function of a collection of (intermediate) point intervention total effects.
3. Modified Cholesky decomposition (Nandy et al., 2017): If the causal ordering is known, then  $\Gamma$  can be recovered by regressing each variables on all the preceding variables, which is equivalent to Cholesky decomposition of the covariance  $\Sigma = LDL^{\top}$ . In fact,  $L^{-1}$  is a lower-triangular matrix filled with negative values of these regression coefficients and ones on the diagonal. Since the causal ordering is unknown, Nandy et al. (2017, §3.2) consider an iterative procedure that alters between Cholesky decomposition and its reverse.

**Definition 14** *Consider an estimator  $\widehat{\theta}_n$  of  $\theta$ ,  $\theta \in \mathbb{R}^k$ . We say that the asymptotic covariance of  $\widehat{\theta}_n$  is  $S$ , and write  $\text{acov}\widehat{\theta}_n = S$ , if  $\sqrt{n}(\widehat{\theta}_n - \theta) \rightarrow_d \mathcal{N}(\mathbf{0}, S)$ . When  $k = 1$ , we write  $\widehat{\text{avar}}\widehat{\theta}_n$  for asymptotic variance.*

For real symmetric matrices  $A$  and  $B$ , we say  $A \succeq B$  if  $A - B$  is positive semidefinite. We now state our main result.

**Theorem 15 (Asymptotic efficiency of the  $\mathcal{G}$ -regression estimator)** *Suppose data is generated iid from a linear SEM Eqs. (2) and (3) associated with an unknown causal DAG  $\mathcal{D}$ . Suppose  $\mathcal{D} \in [\mathcal{G}]$  for a known MPDAG  $\mathcal{G}$ . Further, suppose for  $A \subset V$ ,  $Y \in V \setminus A$ ,  $\tau_{AY}$  is identified under the criterion of Theorem 2. Let  $\hat{\tau}_{AY}^{\mathcal{G}}$  be the  $\mathcal{G}$ -regression estimator of  $\tau_{AY}$  (Theorem 13). Consider any consistent estimator  $\hat{\tau}_{AY} = \hat{\tau}_{AY}(\hat{\Sigma}^{(n)})$  that is a differentiable function of the sample covariance. It holds that*

$$\text{acov}(\hat{\tau}_{AY}) \succeq \text{acov}(\hat{\tau}_{AY}^{\mathcal{G}}).$$

It is clear from definitions that both  $\hat{\tau}_{AY}^{\mathcal{G}}$  and  $\hat{\tau}_{AY}$  are asymptotically linear. Therefore, their asymptotic covariances are well-defined. To prove Theorem 15, it suffices to show that for every  $w \in \mathbb{R}^{|A|}$

$$\text{avar}\left(w^\top \hat{\tau}_{AY}\right) \geq \text{avar}\left(w^\top \hat{\tau}_{AY}^{\mathcal{G}}\right).$$

To this end, for any fixed  $w \in \mathbb{R}^{|A|}$  we define  $\tau_w$  as

$$\tau_w := w^\top \tau_{AY} = \tau_w(\Lambda), \tag{20}$$

which is a smooth function of  $\Lambda$ . The corresponding  $\mathcal{G}$ -regression estimator  $\hat{\tau}_w^{\mathcal{G}} := w^\top \hat{\tau}_{AY}^{\mathcal{G}} = \tau_w(\hat{\Lambda}^{\mathcal{G}})$  is still a plugin estimator (now of  $\tau_w$ ). Additionally, for a consistent estimator  $\hat{\tau}_{AY}$  of  $\tau_{AY}$ , the corresponding  $\hat{\tau}_w := w^\top \hat{\tau}_{AY} = \hat{\tau}_w(\hat{\Sigma}^{(n)})$  is a consistent estimator of  $\tau_w$ , in the form of a differentiable function of the sample covariance. It suffices to show  $\text{avar} \hat{\tau}_w \geq \text{avar} \hat{\tau}_w^{\mathcal{G}}$  for every  $w \in \mathbb{R}^{|A|}$ .

The rest of this section is devoted to proving Theorem 15. First, we introduce graph  $\bar{\mathcal{G}}$  as a saturated version of  $\mathcal{G}$  (Theorem 16). In Section 6.1, we show that  $\mathcal{G}$ -regression with  $\mathcal{G}$  replaced by  $\bar{\mathcal{G}}$ , aptly named  $\bar{\mathcal{G}}$ -regression, is a diffeomorphism between the space of covariance matrices and the space of parameters. In Section 6.2, we characterize the class of estimators relative to which  $\mathcal{G}$ -regression is optimal. To prove Theorem 15, we establish an efficiency bound for this class of estimators in Section 6.4 and verify that  $\mathcal{G}$ -regression achieves this bound in Section 6.5. Some of the proofs are left to Appendix A. See also Fig. 5 in for an overview of the dependency structure of our results in this section.

### 6.1 $\bar{\mathcal{G}}$ -regression as a diffeomorphism

**Proposition 16 (Saturated MPDAG  $\bar{\mathcal{G}}$ )** *For MPDAG  $\mathcal{G} = (V, E, U)$ , an associated saturated MPDAG is  $\bar{\mathcal{G}} = (V, \bar{E}, U)$ , such that  $\text{Pa}(B_k, \bar{\mathcal{G}}) = B_{[k-1]}$  for  $k = 2, \dots, K$ , where  $(B_1, \dots, B_K)$  is a bucket decomposition of  $V$  in both  $\mathcal{G}$  and  $\bar{\mathcal{G}}$ .*

The proof can be found in Appendix B. In words, to create the saturated MPDAG  $\bar{\mathcal{G}}$ , we add all the possible directed edges between buckets  $B_1, \dots, B_K$  subject to the ordering  $B_1, \dots, B_K$ . By construction,  $\bar{\mathcal{G}}$  also satisfies the restrictive property in Theorem 4. See Fig. 1 for an example.

In the following, we introduce  $\bar{\mathcal{G}}$ -regression as a technical tool for establishing a diffeomorphism between the space of sample covariance matrices and the space of parameters in our semiparametric model. This link is the key to analyzing the efficiency of the estimators under consideration.




Figure 1: (a) MPDAG  $\mathcal{G} = (V, E, U)$  with buckets  $B_1 = \{1\}$ ,  $B_2 = \{2, 3, 4\}$  and  $B_3 = \{5, 6\}$  and (b) its associated saturated MPDAG  $\bar{\mathcal{G}} = (V, \bar{E}, U)$ . The new edges in  $\bar{E} \setminus E$  are drawn as dashed. Both  $\mathcal{G}$  and  $\bar{\mathcal{G}}$  satisfy the restrictive property in Corollary 4.

Recall that  $\mathcal{P}_{\mathcal{G}}$  is the set of observed distributions generated by some linear SEM associated with a causal DAG  $\mathcal{D} \in [\mathcal{G}]$ , which is characterized by Theorem 7. More explicitly, let  $Q_k$  be the law of  $\varepsilon_{B_k}$  for  $k = 1, \dots, K$ . The set of laws is explicitly prescribed as

$$\mathcal{P}_{\mathcal{G}} = \left\{ Q_1(X_{B_1}) \prod_{k=2}^K Q_k \left( X_{B_k} - \Lambda_{B_{[k-1]}, B_k}^\top X_{B_{[k-1]}} \right) : Q_k \in \mathcal{P}_{\mathcal{G}_{B_k}}, i \rightarrow j \text{ not in } \mathcal{G} \Rightarrow \lambda_{ij} = 0 \right\} \quad (21)$$

where the law is indexed by  $\Lambda = (\lambda_{ij})$  and  $(Q_k)_{k=1}^K$ . This is a *semiparametric* model and  $(Q_k)_k$  is an infinite-dimensional nuisance (van der Vaart, 2000, Chap. 25).

Consider the set of laws  $\mathcal{P}_{\bar{\mathcal{G}}}$  associated with the saturated graph. Let  $\Omega_k := \mathbb{E}_{Q_k} \varepsilon \varepsilon^\top$  be the covariance of  $Q_k$  for  $k = 1, \dots, K$ . Let  $\mathbb{R}_{PD}^{n \times n}$  denote the set of  $n \times n$  symmetric, positive definite matrices. By our assumption,  $\Omega_k \in \mathbb{R}_{PD}^{|B_k| \times |B_k|}$ . Also, consider the coefficients  $\Lambda = (\lambda_{ij})$  such that  $\lambda_{ij} \neq 0$  only if  $i \rightarrow j$  in  $\bar{\mathcal{G}}$ , or equivalently,  $i \in B_l$  and  $j \in B_m$  for  $l < m$ . Then, the covariance of  $X$ , denoted as  $\Sigma$ , under any  $P \in \mathcal{P}_{\bar{\mathcal{G}}}$  is determined from  $(\Omega_k)_k$  and  $\Lambda$ . Let us write this *covariance map* as

$$\Sigma = \phi_{\bar{\mathcal{G}}} \left( (\Lambda_k)_{k=2}^K, (\Omega_k)_{k=1}^K \right),$$

where  $\Lambda_k = \Lambda_{B_{[k-1]}, B_k}$  is of dimension  $(|B_1| + \dots + |B_{k-1}|) \times |B_k|$ . It follows from Theorem 8 that the covariance map  $\phi_{\bar{\mathcal{G}}}$  is explicitly given by

$$\Sigma_{B_1} = \Omega_1, \quad \Sigma_{B_k} = \Lambda_k^\top \Sigma_{B_{[k-1]}} \Lambda_k + \Omega_k, \quad \Sigma_{B_{[k-1]}, B_k} = \Sigma_{B_{[k-1]}} \Lambda_k, \quad k = 2, \dots, K. \quad (22)$$

Further, the covariance map  $\phi_{\bar{\mathcal{G}}}$  is a *diffeomorphism* between its domain and the set of  $|V| \times |V|$  positive definite matrices.

**Lemma 17** *Covariance map  $\phi_{\bar{\mathcal{G}}}$  given by Eq. (22) is invertible. Further,  $((\Lambda_k)_{k=2}^K, (\Omega_k)_{k=1}^K) \leftrightarrow \Sigma$  given by  $\phi_{\bar{\mathcal{G}}}$  and its inverse  $\phi_{\bar{\mathcal{G}}}^{-1}$  is a diffeomorphism between  $\left( \times_{k=2}^K \mathbb{R}^{(|B_1| + \dots + |B_{k-1}|) \times |B_k|} \right) \times \left( \times_{k=1}^K \mathbb{R}_{PD}^{|B_k| \times |B_k|} \right)$  and  $\mathbb{R}_{PD}^{|V| \times |V|}$ .*

**Proof** By definition, covariance map  $\phi_{\bar{\mathcal{G}}}$  is differentiable. To show diffeomorphism, we need to show that  $\phi_{\bar{\mathcal{G}}}^{-1}(\Sigma)$  exists for every  $\Sigma \in \mathbb{R}_{PD}^{|V| \times |V|}$  and that  $\phi_{\bar{\mathcal{G}}}^{-1}$  is differentiable. For any positive definite  $\Sigma$ , the inverse covariance map  $\phi_{\bar{\mathcal{G}}}^{-1}(\Sigma)$  is explicitly given by

$$\Lambda_k = \left( \Sigma_{B_{[k-1]}} \right)^{-1} \Sigma_{B_{[k-1]}, B_k}, \quad k = 2, \dots, K, \quad (23)$$

and

$$\Omega_k = \Sigma_{B_k \cdot B_{[k-1]}} = \Sigma_{B_k} - \Sigma_{B_{[k-1]}, B_k}^\top \Sigma_{B_{[k-1]}}^{-1} \Sigma_{B_{[k-1]}, B_k}, \quad k = 1, \dots, K, \quad (24)$$

where  $\Sigma_{B_k \cdot B_{[k-1]}}$  is the Schur complement of block  $B_k$  with respect to block  $B_{[k-1]}$ . Because  $\Sigma$  is positive definite, Schur complement  $\Omega_k$  is also positive definite (Horn and Johnson, 2012, page 495). Clearly, the map  $\phi_{\bar{\mathcal{G}}}^{-1}(\cdot)$  is differentiable.  $\blacksquare$

By Eqs. (23) and (24),  $\Lambda_k$  is the matrix of population least squares coefficients in a regression of  $X_{B_k}$  onto  $X_{B_1 \cup \dots \cup B_{k-1}}$  according to  $\bar{\mathcal{G}}$ , and  $\Omega_k$  is the corresponding covariance of regression residuals. Hence,  $\phi_{\bar{\mathcal{G}}}^{-1}(\Sigma)$  is called “ $\bar{\mathcal{G}}$ -regression”.

**Remark 18** *In the special case when  $\mathcal{G}$  is a DAG such that every bucket  $B_i$  is a singleton, Theorem 17 reduces to  $(\Lambda, \omega) \leftrightarrow \Sigma$  given by  $(\phi_{\bar{\mathcal{G}}}, \phi_{\bar{\mathcal{G}}}^{-1})$  being a diffeomorphism between*

$$\left\{ \Lambda \in \mathbb{R}^{|V| \times |V|} : \Lambda \text{ is upper-triangular} \right\} \times \left\{ \omega \in \mathbb{R}^{|V|} : \omega_i > 0, i = 1, \dots, |V| \right\} \longleftrightarrow \mathbb{R}_{PD}^{|V| \times |V|}.$$

*The covariance map is  $\Sigma = \phi_{\bar{\mathcal{G}}}(\Lambda, \omega) = (I - \Lambda)^{-\top} \text{diag}(\omega)(I - \Lambda)^{-1}$ , and the inverse covariance map  $\phi_{\bar{\mathcal{G}}}^{-1}$  is given by the unique LDL decomposition of  $\Sigma^{-1}$ . Theorem 17 is a generalization of Drton (2018, Theorem 7.2).*

## 6.2 Covariance-based, consistent estimators

We now characterize the class of estimators relative to which the optimality of our estimator is established. Recall that under  $P \in \mathcal{P}_{\mathcal{G}}$ ,  $\widehat{\Sigma} = \widehat{\Sigma}^{(n)}$  is the sample covariance,  $\Sigma$  is the population covariance and  $\tau_w = w^\top \tau_{AY}$ . We assume that  $n > \max_k \{|B_k| + |\text{Pa}(B_k, \mathcal{G})|\}$  such that  $\widehat{\Sigma}^{(n)}$  is positive definite almost surely (Drton and Eichler, 2006, Sec. 3.1). For simplicity, the superscript  $(n)$  is often omitted.

**Definition 19** *The class of estimators for  $\tau_w$  under consideration is*

$$\mathcal{T}_w := \left\{ \widehat{\tau}_w \left( \widehat{\Sigma}^{(n)} \right) : \mathbb{R}_{PD}^{|V| \times |V|} \rightarrow \mathbb{R} : \right. \\ \left. \widehat{\tau}_w \text{ differentiable, } \widehat{\tau}_w \left( \widehat{\Sigma}^{(n)} \right) \rightarrow_p \tau_w(P) \text{ as } n \rightarrow \infty \text{ under every } P \in \mathcal{P}_{\mathcal{G}} \right\}. \quad (25)$$

By definition, in particular,  $\mathcal{T}_w$  includes all regular estimators computable with least squares operations.



6.2.1 CHARACTERIZING  $\mathcal{T}_w$ 

Let  $(\widehat{\Lambda}_k^{\bar{\mathcal{G}}})_{k=2}^K, (\widehat{\Omega}_k^{\bar{\mathcal{G}}})_{k=1}^K$  be the image of  $\widehat{\Sigma}$  under  $\phi_{\bar{\mathcal{G}}}^{-1}$ . Recall that  $(\Lambda_k)_{k=2}^K, (\Omega_k)_{k=1}^K$  is the image of  $\Sigma$  under  $\phi_{\bar{\mathcal{G}}}^{-1}$ . For a matrix  $C$ , let  $\text{vec } C$  denote vectorizing  $C$  by concatenating its columns. Each  $\text{vec } \widehat{\Lambda}_k^{\bar{\mathcal{G}}}$  can be split by *coordinates* into vectors

$$\widehat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}} = \left( \widehat{\lambda}_{ij}^{\bar{\mathcal{G}}} : j \in B_k, i \in \text{Pa}(B_k, \mathcal{G}) \right), \quad \widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}} = \left( \widehat{\lambda}_{ij}^{\bar{\mathcal{G}}} : j \in B_k, i \in \text{Pa}(B_k, \bar{\mathcal{G}}) \setminus \text{Pa}(B_k, \mathcal{G}) \right), \quad (26)$$

where  $\left( \widehat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}} \right)_k$  corresponds to between-bucket edges in  $\mathcal{G}$  and  $\left( \widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}} \right)_k$  corresponds to between-bucket edges in  $\bar{\mathcal{G}}$  but not in  $\mathcal{G}$ . In the example of Fig. 1, we have  $\widehat{\Lambda}_{2,\mathcal{G}}^{\bar{\mathcal{G}}} = (\widehat{\lambda}_{12}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{13}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{14}^{\bar{\mathcal{G}}})^\top$ ,  $\widehat{\Lambda}_{3,\mathcal{G}}^{\bar{\mathcal{G}}} = (\widehat{\lambda}_{45}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{46}^{\bar{\mathcal{G}}})^\top$  and  $\widehat{\Lambda}_{2,\mathcal{G}^c}^{\bar{\mathcal{G}}} = \text{NULL}$ ,  $\widehat{\Lambda}_{3,\mathcal{G}^c}^{\bar{\mathcal{G}}} = (\widehat{\lambda}_{15}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{16}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{25}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{26}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{35}^{\bar{\mathcal{G}}}, \widehat{\lambda}_{36}^{\bar{\mathcal{G}}})^\top$ . Similarly,  $\text{vec } \Lambda_k$  can be split into  $\Lambda_{k,\mathcal{G}}$  and  $\Lambda_{k,\mathcal{G}^c}$  for  $k = 2, \dots, K$ .

The following lemma directly follows from Theorem 19 and Theorem 17.

**Lemma 20** *An estimator  $\widehat{\tau}_w \in \mathcal{T}_w$  can be written as*

$$\widehat{\tau}_w \left( \widehat{\Sigma}^{(n)} \right) = \widehat{\tau}_w \left( (\widehat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}})_{k=2}^K, (\widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}})_{k=2}^K, (\widehat{\Omega}_k^{\bar{\mathcal{G}}})_{k=1}^K \right)$$

for function  $\widehat{\tau}_w \left( (\widehat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}})_{k=2}^K, (\widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}})_{k=2}^K, (\widehat{\Omega}_k^{\bar{\mathcal{G}}})_{k=1}^K \right)$  that is differentiable in its arguments.

The consistency of  $\widehat{\tau}_w$  implies the following two results.

**Lemma 21** *For any  $\widehat{\tau}_w \in \mathcal{T}_w$ , it holds that*

$$\widehat{\tau}_w \left( (\Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k)_{k=1}^K \right) \equiv \tau_w \left( (\Lambda_{k,\mathcal{G}})_{k=2}^K \right) \quad (27)$$

for all  $(\Lambda_{k,\mathcal{G}})_k$  and all positive definite  $(\Omega_k)_k$ .

**Proof** Under any  $P \in \mathcal{P}_{\mathcal{G}}$ , since  $\widehat{\Sigma} \rightarrow_p \Sigma$  as  $n \rightarrow \infty$  by the law of large numbers, by Theorem 17 and the continuous mapping theorem (van der Vaart, 2000, page 11), we have  $\widehat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}} \rightarrow_p \Lambda_{k,\mathcal{G}}$ ,  $\widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}} \rightarrow_p \mathbf{0}$  and  $\widehat{\Omega}_k^{\bar{\mathcal{G}}} \rightarrow_p \Omega_k$  for  $k = 2, \dots, K$ . By Theorem 20 and continuous mapping again,  $\widehat{\tau}_w \rightarrow_p \tau_w \left( (\Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k)_{k=1}^K \right)$ . The result then follows from the consistency of  $\widehat{\tau}_w$  under every  $P \in \mathcal{P}_{\mathcal{G}}$ .  $\blacksquare$

**Corollary 22** *For  $\widehat{\tau}_w \in \mathcal{T}_w$ , at any  $((\Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k)_{k=1}^K)$ , it holds that*

$$\frac{\partial \widehat{\tau}_w}{\partial \Lambda_{k,\mathcal{G}}} = \frac{\partial \tau_w}{\partial \Lambda_{k,\mathcal{G}}} \quad (k = 2, \dots, K), \quad \frac{\partial \widehat{\tau}_w}{\partial \Omega_k} = \mathbf{0} \quad (k = 1, \dots, K). \quad (28)$$

**Proof** Let symbol  $\langle \cdot, \cdot \rangle$  denote inner product. Since  $\widehat{\tau}_w$  is differentiable (Theorem 20), by a Taylor expansion at  $((\Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k)_{k=1}^K)$ , we have

$$\begin{aligned} & \widehat{\tau}_w \left( (\Lambda_{k,\mathcal{G}} + \Delta \Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k + \Delta \Omega_k)_{k=1}^K \right) - \widehat{\tau}_w \left( (\Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k)_{k=1}^K \right) \\ &= \sum_{k=2}^K \left( \left\langle \frac{\partial \widehat{\tau}_w}{\partial \Lambda_{k,\mathcal{G}}}, \Delta \Lambda_{k,\mathcal{G}} \right\rangle + o(\|\Delta \Lambda_{k,\mathcal{G}}\|) \right) + \sum_{k=1}^K \left( \left\langle \frac{\partial \widehat{\tau}_w}{\partial \Omega_k}, \Delta \Omega_k \right\rangle + o(\|\Delta \Omega_k\|) \right), \end{aligned}$$

which by Theorem 21 must equal  $\tau_w((\Lambda_{k,\mathcal{G}} + \Delta\Lambda_{k,\mathcal{G}})_{k=2}^K) - \tau_w((\Lambda_{k,\mathcal{G}})_{k=2}^K)$ . The result then follows from the differentiability of  $\tau_w(\cdot)$  and the definition of derivatives.  $\blacksquare$

Note that Theorem 22 is similar to the conditions imposed on influence functions in standard semiparametric efficiency theory; see, e.g., Tsiatis (2006, Corollary 1, §3.1). However, the gradients  $\partial\widehat{\tau}_w/\partial\widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}}$  for  $k = 2, \dots, K$  are *free to vary* because  $\widehat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}} \rightarrow_p \mathbf{0}$ . That is, an estimator  $\widehat{\tau}_w \in \mathcal{T}_w$  can take arbitrary values as its second argument varies in the vicinity of zero, as long as differentiability is maintained.

### 6.3 Asymptotic covariance of least squares coefficients

We use this section to derive some asymptotic results that will be used to prove Theorem 15.

Consider a vertex  $j \in B_k$  for  $k \in \{2, \dots, K\}$  and a set of vertices  $C$  such that  $\text{Pa}(B_k, \mathcal{G}) \subseteq C \subseteq \text{Pa}(B_k, \bar{\mathcal{G}})$ . Let  $\widehat{\lambda}_{C,j}^{(n)} \in \mathbb{R}^{|C|}$  be the least squares coefficients from regressing  $X_j$  onto  $X_C$  under sample size  $n$ . Let  $\lambda_{C,j}$  be the corresponding true edge coefficient vector from  $\Lambda$  in Theorem 7. Then  $\lambda_{C,j}$  has non-zero coordinates only for those indices in  $\text{Pa}(B_k, \mathcal{G})$ . Because  $X_j = \lambda_{C,j}^\top X_C + \varepsilon_j$  with  $\varepsilon_j \perp\!\!\!\perp X_C$  by Theorem 8, we have  $\widehat{\lambda}_{C,j}^{(n)} \rightarrow_p \lambda_{C,j}$  under every  $P \in \mathcal{P}_{\mathcal{G}}$ . Moreover, we have the following asymptotic linear expansion.

**Lemma 23** *Let  $j$  be a vertex in bucket  $B_k$  for  $k \in \{2, \dots, K\}$ . Let  $C$  be a set of vertices such that  $\text{Pa}(B_k, \mathcal{G}) \subseteq C \subseteq \text{Pa}(B_k, \bar{\mathcal{G}})$ . Under any  $P \in \mathcal{P}_{\mathcal{G}}$ , it holds that*

$$\widehat{\lambda}_{C,j}^{(n)} - \lambda_{C,j} = \frac{1}{n} \sum_{i=1}^n (\Sigma_C)^{-1} X_C^{(i)} \varepsilon_j^{(i)} + O_p(n^{-1}),$$

where  $\Sigma = \mathbb{E}_P X X^\top$ ,  $\widehat{\lambda}_{C,j}^{(n)}$  is the vector of least squares coefficients from regressing  $X_j$  onto  $X_C$  under sample size  $n$ , and  $\lambda_{C,j}$  is the vector of true coefficients in Theorem 7.

We now use Theorem 23 to obtain the covariance structure of  $\bar{\mathcal{G}}$ -regression coefficients  $(\widehat{\Lambda}_k^{\bar{\mathcal{G}}})_{k=2}^K$ . Recall that  $\widehat{\Lambda}_k^{\bar{\mathcal{G}}} \in \mathbb{R}^{|B_{[k-1]}| \times |B_k|}$  with  $B_{[k-1]} = B_1 \cup \dots \cup B_{k-1}$  and

$$\left( (\widehat{\Lambda}_k^{\bar{\mathcal{G}}})_{k=2}^K, (\widehat{\Omega}_k^{\bar{\mathcal{G}}})_{k=1}^K \right) = \phi_{\bar{\mathcal{G}}}^{-1} \left( \widehat{\Sigma}^{(n)} \right),$$

as given by Eqs. (23) and (24). For matrices  $A \in \mathbb{R}^{m \times n}$ ,  $B \in \mathbb{R}^{p \times q}$ , the Kronecker product  $A \otimes B$  is an  $mp \times nq$  matrix given by

$$A \otimes B = \begin{pmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{pmatrix}.$$

**Lemma 24** *Let  $(\widehat{\Lambda}_k^{\bar{\mathcal{G}}})_{k=2}^K$  be the  $\bar{\mathcal{G}}$ -regression coefficients under sample size  $n$ . Under any  $P \in \mathcal{P}_{\mathcal{G}}$ , it holds that*

$$\sqrt{n} \begin{pmatrix} \text{vec}(\widehat{\Lambda}_2^{\bar{\mathcal{G}}} - \Lambda_2) \\ \vdots \\ \text{vec}(\widehat{\Lambda}_K^{\bar{\mathcal{G}}} - \Lambda_K) \end{pmatrix} \rightarrow_d \mathcal{N} \left( \mathbf{0}, \text{diag} \left\{ \Omega_2 \otimes (\Sigma_{B_{[1]}})^{-1}, \dots, \Omega_K \otimes (\Sigma_{B_{[K-1]}})^{-1} \right\} \right).$$

**Remark 25**  $\sqrt{n} \text{vec}(\widehat{\Lambda}_k^{(n)} - \Lambda_k) \rightarrow_d \mathcal{N}\left(\mathbf{0}, \Omega_k \otimes \left(\Sigma_{B_{[k-1]}}\right)^{-1}\right)$  is equivalent to

$$\sqrt{n}(\widehat{\Lambda}_k^{(n)} - \Lambda_k) \rightarrow_d \mathcal{MN}\left(\mathbf{0}, \left(\Sigma_{B_{[k-1]}}\right)^{-1}, \Omega_k\right),$$

where the RHS is a centered matrix normal distribution with row covariance  $(\Sigma_{B_{[k-1]}})^{-1}$  and column covariance  $\Omega_k$ ; see Dawid (1981).

Similarly, we can compute the asymptotic covariance of the  $\mathcal{G}$ -regression coefficients. To obtain the result below, we rely on the restrictive property of  $\mathcal{G}$  (Theorem 4).

**Lemma 26** Let  $(\widehat{\Lambda}_k^{\mathcal{G}})_{k=2}^K$  be the  $\mathcal{G}$ -regression coefficients as defined in Theorem 11 under sample size  $n$ . Under any  $P \in \mathcal{P}_{\mathcal{G}}$ , it holds that

$$\sqrt{n} \begin{pmatrix} \text{vec}(\widehat{\Lambda}_2^{\mathcal{G}} - \Lambda_2) \\ \vdots \\ \text{vec}(\widehat{\Lambda}_K^{\mathcal{G}} - \Lambda_K) \end{pmatrix} \rightarrow_d \mathcal{N}\left(\mathbf{0}, \text{diag}\left\{\Omega_2 \otimes (\Sigma_{\text{Pa}(B_2, \mathcal{G})})^{-1}, \dots, \Omega_K \otimes (\Sigma_{\text{Pa}(B_K, \mathcal{G})})^{-1}\right\}\right).$$

#### 6.4 Efficiency bound

We first notice a simple fact of the quadratic form and a property of the Kronecker product.

**Lemma 27** Let  $S \in \mathbb{R}_{PD}^{n \times n}$ ,  $x \in \mathbb{R}^n$  and suppose that  $(A, B)$  is a partition of the set  $\{1, \dots, n\}$ . For any fixed  $x_A$ , it holds that

$$x^\top S x \geq x_A^\top (S_{A \cdot B}) x_A,$$

where  $S_{A \cdot B} = S_{A, A} - S_{A, B} S_{B, B}^{-1} S_{B, A}$ . The equality holds if and only if  $x_B = -S_{B, B}^{-1} S_{B, A} x_A$ .

**Lemma 28 (Liu (1999, Theorem 1))** Let  $A \in \mathbb{R}^{m \times m}$  and  $C \in \mathbb{R}^{n \times n}$  be non-singular. Suppose  $\alpha \subset [m]$ ,  $\beta \subset [n]$ . Let  $\alpha^c$ ,  $\beta^c$  denote their respective complements. Let  $\gamma^c = \{n(i-1) + j : i \in \alpha^c, j \in \beta^c\}$  and  $\gamma = [mn] \setminus \gamma^c$ . We have

$$A_{\alpha^c \cdot \alpha} \otimes C_{\beta^c \cdot \beta} = (A \otimes C)_{\gamma^c \cdot \gamma}.$$

**Lemma 29** Suppose the assumptions of Theorem 15 hold. Fix  $w \in \mathbb{R}^{|A|}$  and let  $\tau_w = w^\top \tau_{AY} = \tau_w((\Lambda_{k, \mathcal{G}})_{k=2}^K)$  as in Eq. (20). Consider any estimator  $\widehat{\tau}_w \in \mathcal{T}_w$  given by Theorem 19. Then under any  $P \in \mathcal{P}_{\mathcal{G}}$ , it holds that

$$\text{avar}(\widehat{\tau}_w) \geq \sum_{k=2}^K h_k^\top \Omega_k \otimes (\Sigma_{\text{Pa}(B_k, \mathcal{G})})^{-1} h_k, \quad (29)$$

where  $(\Omega_k)_{k=2}^K$  and  $\Sigma$  are determined by  $P$ , and the gradient vectors  $h_k = \partial \tau_w((\Lambda_{k, \mathcal{G}})_k) / \partial \Lambda_{k, \mathcal{G}}$  for  $k = 2, \dots, K$  evaluated at  $(\Lambda_{k, \mathcal{G}})_k$  are determined by  $\tau_w(\cdot)$  and  $P$ .

**Proof** By Theorem 20, estimator  $\hat{\tau}_w \in \mathcal{T}_w$  can be written as

$$\hat{\tau}_w = \hat{\tau}_w \left( (\hat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}})_{k=2}^K, (\hat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}})_{k=2}^K, (\hat{\Omega}_k^{\bar{\mathcal{G}}})_{k=1}^K \right),$$

where the arguments correspond to the image of  $\hat{\Sigma}$  under  $\phi_{\bar{\mathcal{G}}}^{-1}$ ; see Eq. (26). Estimator  $\hat{\tau}_w \in \mathcal{T}_w$  is asymptotically normal. By the delta method (Shorack, 2000, Sec 11.2), we have

$$\text{avar}(\hat{\tau}_w) = \begin{pmatrix} \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}})_{k=2}^K \\ \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}^c})_{k=2}^K \\ \partial \hat{\tau}_w / \partial (\Omega_k)_{k=1}^K \end{pmatrix}^\top \text{acov} \begin{Bmatrix} \text{vec} (\hat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}})_{k=2}^K \\ \text{vec} (\hat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}})_{k=2}^K \\ \text{vec} (\hat{\Omega}_k^{\bar{\mathcal{G}}})_{k=1}^K \end{Bmatrix} \begin{pmatrix} \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}})_{k=2}^K \\ \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}^c})_{k=2}^K \\ \partial \hat{\tau}_w / \partial (\Omega_k)_{k=1}^K \end{pmatrix},$$

where the partial derivatives of  $\hat{\tau}_w(\cdot)$  are evaluated at  $((\Lambda_{k,\mathcal{G}})_{k=2}^K, (\mathbf{0})_{k=2}^K, (\Omega_k)_{k=1}^K)$ , the image of  $\Sigma$  under  $\phi_{\bar{\mathcal{G}}}^{-1}$ .

Using  $\partial \hat{\tau}_w / \partial \Omega_k = \mathbf{0}$  for  $k = 1, \dots, K$  from Theorem 22, it follows that

$$\begin{aligned} \text{avar}(\hat{\tau}_w) &= \begin{pmatrix} \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}})_{k=2}^K \\ \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}^c})_{k=2}^K \end{pmatrix}^\top \text{acov} \begin{Bmatrix} \text{vec} (\hat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}})_{k=2}^K \\ \text{vec} (\hat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}})_{k=2}^K \end{Bmatrix} \begin{pmatrix} \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}})_{k=2}^K \\ \partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}^c})_{k=2}^K \end{pmatrix} \\ &= \sum_{k=2}^K \begin{pmatrix} \partial \hat{\tau}_w / \partial \Lambda_{k,\mathcal{G}} \\ \partial \hat{\tau}_w / \partial \Lambda_{k,\mathcal{G}^c} \end{pmatrix}^\top \text{acov} \begin{Bmatrix} \hat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}} \\ \hat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}} \end{Bmatrix} \begin{pmatrix} \partial \hat{\tau}_w / \partial \Lambda_{k,\mathcal{G}} \\ \partial \hat{\tau}_w / \partial \Lambda_{k,\mathcal{G}^c} \end{pmatrix}, \end{aligned}$$

where we have used the block-diagonal structure of the asymptotic covariance from Theorem 24. Let

$$S^{(k)} := \text{acov} \begin{Bmatrix} \hat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}} \\ \hat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}} \end{Bmatrix}, \quad k = 2, \dots, K,$$

which equals

$$S^{(k)} = \Omega_k \otimes \left( \Sigma_{B_{[k-1]}} \right)^{-1}, \quad k = 2, \dots, K,$$

by Theorem 24. From Theorem 22, note that  $\partial \hat{\tau}_w / \partial (\Lambda_{k,\mathcal{G}})_k \equiv h_k$  is *fixed* for  $k = 2, \dots, K$ . Then, Theorem 27 yields the lower bound

$$\text{avar}(\hat{\tau}_w) \geq \sum_{k=2}^K h_k^\top S_{\mathcal{G},\mathcal{G}^c}^{(k)} h_k,$$

where indices  $\mathcal{G}$  and  $\mathcal{G}^c$  correspond to the coordinates in  $\hat{\Lambda}_{k,\mathcal{G}}^{\bar{\mathcal{G}}}$  and  $\hat{\Lambda}_{k,\mathcal{G}^c}^{\bar{\mathcal{G}}}$  respectively. Indices  $\mathcal{G}$  correspond to  $\{(i, j) : j \in B_k, i \in \text{Pa}(B_k, \mathcal{G})\}$ ; by construction of  $\bar{\mathcal{G}}$ , indices  $\mathcal{G}^c$  correspond to  $\{(i, j) : j \in B_k, i \in \text{Pa}(B_k, \bar{\mathcal{G}}) \setminus \text{Pa}(B_k, \mathcal{G})\}$ . Now, to abuse the notation slightly, we apply Theorem 28 with

$$A = \Omega_k, \quad C = (\Sigma_{B_{[k-1]}})^{-1}, \quad \alpha = \emptyset, \quad \beta = \text{Pa}(B_k, \bar{\mathcal{G}}) \setminus \text{Pa}(B_k, \mathcal{G}),$$

such that

$$\alpha^c = \{1, \dots, |B_k|\}, \quad \beta^c = \text{Pa}(B_k, \mathcal{G}), \quad \gamma = \mathcal{G}^c, \quad \gamma^c = \mathcal{G}.$$

We obtain

$$S_{\mathcal{G},\mathcal{G}^c}^{(k)} = \Omega_k \otimes \left[ (\Sigma_{B_{[k-1]}})^{-1} \right]_{\beta^c, \beta} = \Omega_k \otimes (\Sigma_{\text{Pa}(B_k, \mathcal{G})})^{-1},$$

where the last step follows from  $(H^{-1})_{\beta^c, \beta} = (H_{\beta^c, \beta^c})^{-1}$  (Horn and Johnson, 2012, §0.8). ■

### 6.5 Efficiency of $\mathcal{G}$ -regression estimator

In Section 5, we have seen that when the errors are Gaussian, the  $\mathcal{G}$ -regression plugin is the MLE and hence achieves the efficiency bound. Here, we show that this is still true relative to the class of estimators we consider, even though the errors are not necessarily Gaussian. We verify that  $\hat{\tau}_w^{\mathcal{G}} = w^\top \hat{\tau}_{AY}^{\mathcal{G}}$  achieves the efficiency bound above.

**Lemma 30** *Let  $\hat{\tau}_w^{\mathcal{G}} := w^\top \hat{\tau}_{AY}^{\mathcal{G}}$ , where  $\hat{\tau}_{AY}^{\mathcal{G}}$  is the  $\mathcal{G}$ -regression estimator (Theorem 13). Under the same assumptions as Theorem 29, it holds that  $\hat{\tau}_w^{\mathcal{G}} \in \mathcal{T}_w$  and  $\hat{\tau}_w^{\mathcal{G}}$  achieves the efficiency bound in Eq. (29) under every  $P \in \mathcal{P}_{\mathcal{G}}$ .*

**Proof** By Theorem 13,  $\hat{\tau}_w^{\mathcal{G}} \in \mathcal{T}_w$ . Further, note that

$$\hat{\tau}_w^{\mathcal{G}} = \tau_w \left( (\hat{\Lambda}_k^{\mathcal{G}})_{k=2}^K \right),$$

where  $(\hat{\Lambda}_k^{\mathcal{G}})_{k=2}^K$  are the  $\mathcal{G}$ -regression coefficients in Eq. (18). Under any  $P \in \mathcal{P}_{\mathcal{G}}$ , we now verify that  $\text{avar } \hat{\tau}_w^{\mathcal{G}}$  matches the RHS of Eq. (29). By the delta method (Shorack, 2000, Sec 11.2), we have

$$\begin{aligned} \text{avar } \hat{\tau}_w^{\mathcal{G}} &= \left( \partial \tau_w / \partial \text{vec} (\Lambda_k)_{k=2}^K \right)^\top \text{acov} \left\{ \text{vec} (\hat{\Lambda}_k^{\mathcal{G}})_{k=2}^K \right\} \left( \partial \tau_w / \partial \text{vec} (\Lambda_k)_{k=2}^K \right) \\ &\stackrel{(i)}{=} \sum_{k=2}^K \left( \partial \tau_w / \partial \text{vec} \Lambda_k \right)^\top \text{acov} \left\{ \text{vec} \hat{\Lambda}_k^{\mathcal{G}} \right\} \left( \partial \tau_w / \partial \text{vec} \Lambda_k \right) \\ &\stackrel{(ii)}{=} \sum_{k=2}^K \left( \partial \tau_w / \partial \text{vec} \Lambda_k \right)^\top \Omega_k \otimes \left( \Sigma_{\text{Pa}(B_k, \mathcal{G})} \right)^{-1} \left( \partial \tau_w / \partial \text{vec} \Lambda_k \right), \end{aligned}$$

which equals the RHS of Eq. (29). The partial derivatives of  $\tau_w(\cdot)$  are evaluated at  $(\Lambda_k)_{k=2}^K$ . Step (i) follows from the block-diagonal structure of the asymptotic covariance of  $\hat{\Lambda}_{\mathcal{G}}$  given by Theorem 26, and (ii) follows from the same lemma.  $\blacksquare$

Finally, we complete the proof of our main result.

**Proof** of Theorem 15. Fix any  $P \in \mathcal{P}_{\mathcal{G}}$ . It suffices to show that for every  $w \in \mathbb{R}^{|A|}$ ,

$$w^\top \text{acov}(\hat{\tau}_{AY}) w \geq w^\top \text{acov}(\hat{\tau}_{AY}^{\mathcal{G}}) w,$$

or equivalently

$$\text{avar} \left( w^\top \hat{\tau}_{AY} \right) \geq \text{avar} \left( w^\top \hat{\tau}_{AY}^{\mathcal{G}} \right).$$

This is true because for every  $\hat{\tau}_{AY}$  in consideration,  $\hat{\tau}_w := w^\top \hat{\tau}_{AY} \in \mathcal{T}_w$  and hence  $\hat{\tau}_w$  is subject to the lower bound in Theorem 29. Meanwhile, by Theorem 30, such a lower bound is achieved by  $\hat{\tau}_w^{\mathcal{G}} = w^\top \hat{\tau}_{AY}^{\mathcal{G}}$ . The proof is complete because the choice of  $w$  is arbitrary.  $\blacksquare$

**Remark 31** *For Theorem 15 to hold, the independence error assumption Eq. (3) of the underlying linear SEM cannot be relaxed to uncorrelated errors. This comes from inspecting*

the proof of Theorem 24 in Appendix A. To show that the  $\bar{\mathcal{G}}$ -regression coefficients are asymptotically independent across buckets, the independence of errors is used to establish that for  $2 \leq k < k' \leq K$ ,  $j \in B_k$ ,  $j' \in B_{k'}$ ,  $\text{cov}(\varepsilon_j X_{B_{[k-1]}}, \varepsilon_{j'} X_{B_{[k'-1]}}) = \mathbf{0}$ .

Suppose for now  $\{\varepsilon_i : i \in V\}$  are only uncorrelated and hence  $\{\varepsilon_{B_k} : k = 1, \dots, K\}$  are only uncorrelated across buckets. Further, suppose  $B_1 = \{1\}$ ,  $B_2 = \{2\}$ ,  $B_3 = \{3\}$  with  $j = k = 2$  and  $j' = k' = 3$ . Then, we have

$$\begin{aligned} \text{cov}\left(\varepsilon_j X_{B_{[k-1]}}, \varepsilon_{j'} X_{B_{[k'-1]}}\right) &= \text{cov}\left(\varepsilon_2 \varepsilon_1, \varepsilon_3 (\varepsilon_1, \gamma_{12} \varepsilon_1 + \varepsilon_2)\right)^\top \\ &= \mathbb{E}[\varepsilon_1 \varepsilon_2 (\varepsilon_1 \varepsilon_3, \gamma_{12} \varepsilon_1 \varepsilon_3 + \varepsilon_2 \varepsilon_3)]^\top, \end{aligned}$$

which may be non-zero.

## 7. Numerical Results

In this section, the finite-sample performance of  $\mathcal{G}$ -regression is evaluated against contending estimators. We use simulations and an *in silico* data set for predicting expression levels in gene knockout experiments. All the numerical experiments were conducted with R v3.6, package `pcalg` v2.6 (Kalisch et al., 2012) and our package `eff2` v0.1.

### 7.1 Simulations

We compare the performance of  $\mathcal{G}$ -regression to several contending estimators under finite samples. We roughly follow the simulation setup of Henckel et al. (2022); Witte et al. (2020). First, we draw a random undirected graph from the Erdős-Rényi model with average degree  $k$ , where  $k$  is drawn from  $\{2, 3, 4, 5\}$  uniformly at random. The graph is converted to a DAG  $\mathcal{D}$  with a random causal ordering and the corresponding CPDAG  $\mathcal{G}$  is recorded. Then we fix a linear SEM by drawing  $\gamma_{ij}$  uniformly from  $[-2, -0.1] \cup [0.1, 2]$  and choosing the error distribution uniformly at random from the following:

1.  $\varepsilon_i \sim \mathcal{N}(0, v_i)$  with  $v_i \sim \text{unif}(0.5, 6)$ ,
2.  $\varepsilon_i / \sqrt{v_i} \sim t_5$  with  $v_i \sim \text{unif}(0.5, 1.5)$ ,
3.  $\varepsilon_i \sim \text{logistic}(0, s_i)$  with  $s_i \sim \text{unif}(0.4, 0.7)$ ,
4.  $\varepsilon_i \sim \text{unif}(-a_i, a_i)$  with  $a_i \sim \text{unif}(1.2, 2.1)$ .

We generate  $n$  iid samples from the model. Treatments  $A$  of a fixed size are randomly selected from the set of vertices with non-empty descendants, and  $Y$  is selected randomly from their descendants; the drawing is repeated until  $\tau_{AY}$  is identified from  $\mathcal{G}$  according to the criterion of Theorem 2. Finally, the data and graph  $\mathcal{G}$  are provided to each estimator of  $\tau_{AY}$ .

We compare to the following three estimators:

- `adj.0`: optimal adjustment estimator (Henckel et al., 2022),
- `IDA.M`: joint-IDA estimator based on modifying Cholesky decompositions (Nandy et al., 2017),
- `IDA.R`: joint-IDA estimator based on recursive regressions (Nandy et al., 2017).

They are implemented in R package `pcalg`. The two joint-IDA estimators use the parents of treatment variables to estimate a causal effect. Both of them reduce to the IDA estimator

of Maathuis et al. (2009) when  $|A| = 1$ . Admittedly, compared to  $\mathcal{G}$ -regression and `adj.0`, the joint-IDA estimators require less knowledge about the graph, namely only  $\text{Pa}(i)$  for each  $i \in A$ .

For each estimator  $\hat{\tau}_{AY}$ , we compute its squared error  $\|\hat{\tau}_{AY} - \tau_{AY}\|_2^2$ . Dividing  $\|\hat{\tau}_{AY} - \tau_{AY}\|_2^2$  by the squared error of  $\mathcal{G}$ -regression, we obtain the *relative squared error* of each contending estimator. We consider  $|A| \in \{1, 2, 3, 4\}$ ,  $|V| \in \{20, 50, 100\}$  and  $n \in \{100, 1000\}$ ; each configuration of  $(|A|, |V|, n)$  is replicated 1,000 times.

Fig. 2 shows the distributions of relative squared errors. In Table 1, we summarize the relative errors with their geometric mean and median. Our estimator dominates all the contending estimators in all cases, and the improvement gets larger as  $|A|$  gets bigger. Even though `adj.0` achieves the minimal asymptotic variance among all adjustment estimators, it can compare less favorably to our estimator by several folds. In general, the IDA estimators have very poor performances. Moreover, the results in Table 1 are computed only from the replications where a contending estimator exists. As mentioned in the Introduction, unlike  $\mathcal{G}$ -regression, none of the contending estimators is guaranteed to exist for every identified effect under joint intervention (`adj.0` always exists for point interventions); see Table 2 for the percentages of instances that are not estimable by contending estimators, even though the effect is identified by Theorem 2 and hence estimable by  $\mathcal{G}$ -regression.

Table 1: Geometric average (brackets: median) of relative squared errors compared to  $\mathcal{G}$ -regression

$ A $	$ V  = 20$				$ V  = 50$				$ V  = 100$			
	$n = 100$		$n = 1000$		$n = 100$		$n = 1000$		$n = 100$		$n = 1000$	
<b>adj.0</b>												
1	1.3	(1.0)	1.3	(1.0)	1.4	(1.0)	1.3	(1.0)	1.5	(1.0)	1.5	(1.0)
2	3.4	(1.7)	4.2	(2.0)	4.7	(2.6)	4.9	(2.6)	4.2	(2.6)	4.5	(2.7)
3	6.3	(3.8)	5.9	(3.4)	7.4	(4.9)	7.2	(4.4)	7.8	(5.7)	8.0	(5.2)
4	9.3	(5.0)	9.3	(5.8)	12	(8.0)	14	(8.7)	12	(8.6)	12	(8.9)
<b>IDA.M</b>												
1	20	(17)	19	(16)	61	(57)	48	(45)	103	(78)	108	(90)
2	62	(48)	65	(51)	220	(153)	182	(120)	293	(205)	356	(272)
3	93	(72)	119	(108)	354	(249)	396	(205)	749	(547)	771	(604)
4	154	(111)	222	(147)	533	(448)	895	(440)	1188	(752)	1604	(1508)
<b>IDA.R</b>												
1	20	(17)	19	(16)	61	(57)	48	(45)	103	(78)	108	(90)
2	33	(29)	38	(29)	121	(96)	113	(89)	176	(140)	199	(168)
3	30	(22)	39	(30)	171	(141)	135	(125)	342	(281)	312	(292)
4	48	(34)	50	(41)	187	(132)	214	(143)	405	(391)	432	(342)

In Appendix D, we report two sets of additional simulation results: (1) The CPDAG is estimated with the greedy equivalence search algorithm (Chickering, 2002) and provided to the estimators. The improvements are more modest but are still typically by several folds. (2) The data generating mechanism in this section is modified by rescaling the coefficients such that the variance of a vertex does not grow with the topological ordering.




Figure 2: Violin plots for the relative squared errors of contending estimators (‘.’: geometric mean, ‘+’: median).

Table 2: Percentage of identified instances not estimable using contending estimators (all estimable with  $\mathcal{G}$ -regression)

Estimator	$ A $	$ V  = 20$	$ V  = 50$	$ V  = 100$
adj.O	1	0%	0%	0%
	2	17%	10%	5%
	3	30%	18%	15%
	4	36%	29%	22%
IDA.M	1	29%	32%	32%
	2	47%	51%	50%
	3	61%	59%	63%
	4	72%	69%	71%
IDA.R	1	29%	32%	32%
	2	47%	51%	50%
	3	61%	59%	63%
	4	72%	69%	71%

## 7.2 Predicting double knockouts in DREAM4 data

The DREAM4 *in silico* network challenge data set (Marbach et al., 2009b) provides a benchmark for evaluating the reverse engineering of gene regulation networks. Here we use the 5th *Size10* data set (Marbach et al., 2009a) as our example, which is a small network of 10 genes. Fig. 3 shows the true gene regulation network, which is constructed based on the networks of living organisms. A stochastic differential equation model was used



to generate the data under wild type (steady state), perturbed steady state and knockout interventions. A task in the challenge is to use data under wild type and perturbed steady state (both are observational data) to predict the steady state expression levels under 5 different joint interventions, each of which knocks out a pair of genes. For our purpose, we also use the true network as input. However, the true network contains one cycle (other networks in DREAM4 contain more than one cycles). In the following, we remove one edge in the cycle and provide the resulting DAG to the estimators. Necessarily, the causal DAG is misspecified. Results are reported under 4 different edge removals.




Figure 3: Gene regulation network from DREAM4 data set, which contains a cycle (blue).

Unfortunately, the wild type data only consists of one sample. To estimate the observational covariance, we use the perturbed steady state data, which consists of 5 segments of time series. A sample covariance is computed from each segment, and the final estimate is taken as their average. For a double knockout of genes  $(i, j)$ , we use  $\mathcal{G}$ -regression to estimate the joint-intervention effect of  $A = (i, j)$  on every other gene. The effect is identified because the DAG is given. For gene  $k$ , let  $s_k$  and  $s_k^{(ij)}$  respectively denote its expression level under wild type and double knockout of genes  $(i, j)$ . The expression level under double knockout is predicted as

$$\hat{s}_k^{(ij)} = \begin{cases} s_k - (s_i, s_j)^\top \hat{\tau}_{ij,k}, & k \notin \{i, j\} \\ 0, & k \in \{i, j\} \end{cases}.$$

The performance is evaluated with normalized squared error

$$\mathcal{E} = \frac{\sum_{(i,j) \in \mathcal{A}} \sum_{k=1}^{10} (\hat{s}_k^{(ij)} - s_k^{(ij)})^2}{\sum_{(i,j) \in \mathcal{A}} \sum_{k=1}^{10} (s_k^{(ij)})^2},$$

where  $\mathcal{A} = \{(6, 8), (7, 8), (8, 10), (8, 5), (8, 9)\}$  consists of 5 double knockouts available in the data set. For comparison, we also evaluate the performance of `adj.0` (optimal adjustment, Henckel et al. (2022)) and `IDA.R` (joint-IDA based on recursive regressions, Nandy et al. (2017)); `IDA.R` is chosen because it outperforms `IDA.M` according to Section 7.1. Unfortunately, `adj.0` is not able to estimate the effect on every  $k$  and a modified metric  $\mathcal{E}^*$  is

computed by only summing over those estimable  $k$ 's; the same metric  $\mathcal{E}^*$  of  $\mathcal{G}$ -regression is also computed for comparison. As a baseline, we also compute  $\mathcal{E}$  from naively estimating  $s_k^{(ij)}$  with just  $s_k$ .

Table 3: Normalized squared errors of predicting gene double knockouts

edge removed from cycle	# adj.0	$\mathcal{E}^*$		$\mathcal{E}$		
		adj.0	$\mathcal{G}$ -reg	IDA.R	$\mathcal{G}$ -reg	baseline
5 $\rightarrow$ 6	36%	43%	<b>35%</b>	46%	<b>30%</b>	81%
6 $\rightarrow$ 8	42%	<b>29%</b>	32%	33%	<b>26%</b>	81%
8 $\rightarrow$ 7	60%	39%	<b>35%</b>	45%	<b>44%</b>	81%
7 $\rightarrow$ 5	46%	40%	<b>33%</b>	45%	<b>34%</b>	81%

Table 3 reports the results, where the column ‘# adj.0’ lists the percentage of effects not estimable by the adjustment estimator. In almost all the cases,  $\mathcal{G}$ -regression dominates all the contending estimators. In this example, even though both the causal graph and the linear SEM are misspecified, one can still witness some usefulness of our estimator.

### 8. Discussion

We have proposed  $\mathcal{G}$ -regression based on recursive least squares to estimate a total causal effect from observational data, under linearity and causal sufficiency assumptions.  $\mathcal{G}$ -regression is applicable to estimating every identified total effect, under either point intervention or joint intervention. Further, via a new semiparametric efficiency theory, we have shown that the estimator achieves the efficiency bound within a restricted, yet reasonably large, class of estimators, including covariate adjustment and other regular estimators based on the sample covariance. Note that the restriction on the class of estimators is motivated by computational simplicity and numerical stability as mentioned in the Introduction. To construct confidence intervals and conduct hypothesis tests, bootstrap can be easily applied to estimate the asymptotic covariance of our estimator. This is implemented in R package `eff2`.




Figure 4: (a) and (b) lead to the same  $\mathcal{G}$ -regression estimator of  $\tau_{AY}$ . The independence between  $O_1$  and  $O_2$  in (a) is dropped in (b).

We conclude with a remark. We have seen that the conditional independence constraints in Eq. (12) play no role for the restricted class of estimators considered — a feature that holds under the restrictive property (Theorem 4). For example, the marginal independence between  $O_1$  and  $O_2$  in Fig. 4(a) can be ignored without changing the  $\mathcal{G}$ -regression estimator of  $\tau_{AY}$ . However, this is no longer true when linearity is dropped. In terms of nonparametric causal graphical models, the asymptotic relative efficiency resulting from ignoring  $O_1 \perp\!\!\!\perp O_2$  can be arbitrarily large; see Rotnitzky and Smucler (2020, Lemma 23).

## Acknowledgments

Authors thank Thomas Richardson for valuable comments and discussions. The first author was supported by ONR Grant N00014-19-1-2446.


**Appendix A. Proofs for asymptotic efficiency**


Figure 5: Dependency structure of proofs in Section 6 of main text.

**A.1 Proof of Theorem 23**

**Proof** For simplicity, we drop the superscripts in  $\widehat{\Sigma}^{(n)}$  and  $\widehat{\lambda}^{(n)}$ . Since  $j \in B_k$  and  $\text{Pa}(B_k, \mathcal{G}) \subseteq C \subseteq \text{Pa}(B_k, \widehat{\mathcal{G}})$ , we have

$$\begin{aligned} \widehat{\lambda}_{C,j} - \lambda_{C,j} &= (\widehat{\Sigma}_C)^{-1} \widehat{\Sigma}_{C,j} - (\Sigma_C)^{-1} \Sigma_{C,j} \\ &= \left( \Sigma_C + \widehat{\Sigma}_C - \Sigma_C \right)^{-1} \left( \widehat{\Sigma}_{C,j} - \Sigma_{C,j} \right) + \left( (\widehat{\Sigma}_C)^{-1} - (\Sigma_C)^{-1} \right) \Sigma_{C,j}. \end{aligned}$$

We compute the two terms separately. The first term becomes

$$\begin{aligned} \left( \Sigma_C + \widehat{\Sigma}_C - \Sigma_C \right)^{-1} \left( \widehat{\Sigma}_{C,j} - \Sigma_{C,j} \right) &= \left( \Sigma_C + O_p(n^{-1/2}) \right)^{-1} \left( \widehat{\Sigma}_{C,j} - \Sigma_{C,j} \right) \\ &= (\Sigma_C)^{-1} \left( \widehat{\Sigma}_{C,j} - \Sigma_{C,j} \right) + O_p(n^{-1}), \end{aligned}$$

where we used the fact that  $\Sigma_C$  is positive definite (Theorem 17) and  $\|\widehat{\Sigma}_{C,j} - \Sigma_{C,j}\| = O_p(n^{-1/2})$ ,  $\|\widehat{\Sigma}_C - \Sigma_C\|_2 = O_p(n^{-1/2})$  by the central limit theorem.

In the second term,

$$\begin{aligned} (\widehat{\Sigma}_C)^{-1} - (\Sigma_C)^{-1} &= \left( \Sigma_C + \widehat{\Sigma}_C - \Sigma_C \right)^{-1} - (\Sigma_C)^{-1} \\ &= \left[ I - \left( I - (\Sigma_C)^{-1} \widehat{\Sigma}_C \right) \right]^{-1} (\Sigma_C)^{-1} - (\Sigma_C)^{-1}. \end{aligned}$$

Since  $\|I - (\Sigma_C)^{-1} \widehat{\Sigma}_C\|_2 = O_p(n^{-1/2})$ , using Neumann series  $(I - H)^{-1} = I + H + H^2 + \dots$  for  $H = I - (\Sigma_C)^{-1} \widehat{\Sigma}_C$  with  $\|H\|_2 \rightarrow_p 0 < 1$ , we have

$$\begin{aligned} (\widehat{\Sigma}_C)^{-1} - (\Sigma_C)^{-1} &= \left[ I + H + O_p(n^{-1}) \right] (\Sigma_C)^{-1} - (\Sigma_C)^{-1} \\ &= H(\Sigma_C)^{-1} + O_p(n^{-1}) \\ &= \left[ I - (\Sigma_C)^{-1} \widehat{\Sigma}_C \right] (\Sigma_C)^{-1} + O_p(n^{-1}). \end{aligned}$$

Combining the two terms, we obtain

$$\begin{aligned}
 \widehat{\lambda}_{C,j} - \lambda_{C,j} &= (\Sigma_C)^{-1} \left( \widehat{\Sigma}_{C,j} - \Sigma_{C,j} \right) + \left[ I - (\Sigma_C)^{-1} \widehat{\Sigma}_C \right] (\Sigma_C)^{-1} \Sigma_{C,j} + O_p(n^{-1}) \\
 &= (\Sigma_C)^{-1} \widehat{\Sigma}_{C,j} - (\Sigma_C)^{-1} \Sigma_{C,j} + (\Sigma_C)^{-1} \Sigma_{C,j} - (\Sigma_C)^{-1} \widehat{\Sigma}_C (\Sigma_C)^{-1} \Sigma_{C,j} + O_p(n^{-1}) \\
 &\stackrel{(i)}{=} (\Sigma_C)^{-1} \left( \widehat{\Sigma}_{C,j} - \widehat{\Sigma}_C \lambda_{C,j} \right) + O_p(n^{-1}) \\
 &= \frac{1}{n} \sum_{i=1}^n (\Sigma_C)^{-1} \left[ X_j^{(i)} X_C^{(i)} - X_C^{(i)} X_C^{(i)\top} \lambda_{C,j} \right] + O_p(n^{-1}) \\
 &= \frac{1}{n} \sum_{i=1}^n (\Sigma_C)^{-1} X_C^{(i)} \left( X_j^{(i)} - \lambda_{C,j}^\top X_C^{(i)} \right) + O_p(n^{-1}) \\
 &\stackrel{(ii)}{=} \frac{1}{n} \sum_{i=1}^n (\Sigma_C)^{-1} X_C^{(i)} \left( X_j^{(i)} - \lambda_{\text{Pa}(B_k, \mathcal{G}), j}^\top X_{\text{Pa}(B_k, \mathcal{G})}^{(i)} \right) + O_p(n^{-1}) \\
 &= \frac{1}{n} \sum_{i=1}^n (\Sigma_C)^{-1} X_C^{(i)} \varepsilon_j^{(i)} + O_p(n^{-1}),
 \end{aligned}$$

where (i) uses  $\lambda_{C,j} = (\Sigma_C)^{-1} \Sigma_{C,j}$  and (ii) follows from Theorem 7 and  $\text{Pa}(B_k, \mathcal{G}) \subseteq C \subseteq \text{Pa}(B_k, \bar{\mathcal{G}})$ .  $\blacksquare$

## A.2 Proof of Theorem 24

**Proof** For each  $k = 2, \dots, K$ , note that for  $C = \text{Pa}(B_k, \bar{\mathcal{G}}) = B_{[k-1]}$ ,  $\text{vec} \widehat{\Lambda}_k^{\bar{\mathcal{G}}} = (\widehat{\lambda}_{C,j}^{(n)})_{j \in B_k}$  by concatenation. By Theorem 23, we have the following asymptotic linear expansion

$$\widehat{\lambda}_{B_{[k-1]}, j}^{(n)} - \lambda_{B_{[k-1]}, j} = \frac{1}{n} \sum_{i=1}^n \left( \Sigma_{B_{[k-1]}} \right)^{-1} X_{B_{[k-1]}}^{(i)} \varepsilon_j^{(i)} + O_p(n^{-1}). \quad (30)$$

By the central limit theorem,

$$\sqrt{n} \begin{pmatrix} \text{vec}(\widehat{\Lambda}_2^{\bar{\mathcal{G}}} - \Lambda_2) \\ \vdots \\ \text{vec}(\widehat{\Lambda}_K^{\bar{\mathcal{G}}} - \Lambda_K) \end{pmatrix}$$

converges to a centered multivariate normal distribution. Further, we claim that the asymptotic covariance must be block-diagonal according to  $k = 2, \dots, K$ . To see this, take  $k < k'$ ,  $j \in B_k$ ,  $j' \in B_{k'}$  and let  $C = B_{[k-1]}$ ,  $C' = B_{[k'-1]}$ . Using Eq. (30), we have

$$\begin{aligned}
 &\lim_{n \rightarrow \infty} n^{-1} \mathbb{E} \left( \widehat{\lambda}_{C,j}^{(n)} - \lambda_{C,j} \right) \left( \widehat{\lambda}_{C',j'}^{(n)} - \lambda_{C',j'} \right)^\top \\
 &= (\Sigma_C)^{-1} \text{cov}(\varepsilon_j X_C, \varepsilon_{j'} X_{C'}) (\Sigma_{C'})^{-1} \\
 &= (\Sigma_C)^{-1} \left\{ \mathbb{E} \left[ \varepsilon_j \varepsilon_{j'}^\top X_C X_{C'}^\top \right] - \mathbb{E}[\varepsilon_j X_C] \mathbb{E}[\varepsilon_{j'} X_{C'}^\top] \right\} (\Sigma_{C'})^{-1}.
 \end{aligned}$$

In the expression above, because  $\varepsilon_{B_k} \perp\!\!\!\perp X_{B_{[k-1]}}$  and  $\varepsilon_{B_{k'}} \perp\!\!\!\perp X_{B_{[k'-1]}}$  by Theorem 8 and  $j \in B_k$ ,  $j' \in B_{k'}$  for  $k < k'$ , we have  $\mathbb{E}[\varepsilon_j \varepsilon_{j'}^\top X_C X_{C'}^\top] = \mathbb{E} \varepsilon_{j'} \mathbb{E}[\varepsilon_j X_C X_{C'}^\top] = \mathbf{0}$ ,  $\mathbb{E} \varepsilon_j X_C = \mathbf{0}$

and  $\mathbb{E} \varepsilon_{j'} X_{C'} = \mathbf{0}$ . It follows that the display above evaluates to  $\mathbf{0}$  and hence the asymptotic covariance matrix is block-diagonal.

It remains to be shown that  $\text{acov} \text{vec}(\widehat{\Lambda}_k^{\bar{\mathcal{G}}} - \Lambda_k) = \Omega_k \otimes (\Sigma_{B_{[k-1]}})^{-1}$  for  $k = 2, \dots, K$ . Fix  $k$ , take any two distinct  $j, j' \in B_k$  and let  $C = B_{[k-1]}$ . Again using Eq. (30), we have

$$\text{acov} \begin{pmatrix} \widehat{\lambda}_{C,j}^{(n)} \\ \widehat{\lambda}_{C,j'}^{(n)} \end{pmatrix} = \begin{pmatrix} H & F \\ F^\top & D \end{pmatrix},$$

where

$$\begin{aligned} H &= (\Sigma_C)^{-1} \text{cov}(\varepsilon_j X_C, \varepsilon_j X_C) (\Sigma_C)^{-1} = \text{var}(\varepsilon_j) (\Sigma_{B_{[k-1]}})^{-1}, \\ F &= (\Sigma_C)^{-1} \text{cov}(\varepsilon_j X_C, \varepsilon_{j'} X_C) (\Sigma_C)^{-1} = \text{cov}(\varepsilon_j, \varepsilon_{j'}) (\Sigma_{B_{[k-1]}})^{-1}, \\ D &= (\Sigma_C)^{-1} \text{cov}(\varepsilon_{j'} X_C, \varepsilon_{j'} X_C) (\Sigma_C)^{-1} = \text{var}(\varepsilon_{j'}) (\Sigma_{B_{[k-1]}})^{-1}. \end{aligned}$$

Noting that  $\Omega_k = \text{cov}(\varepsilon_{B_k})$  and  $\text{vec} \widehat{\Lambda}_k^{\bar{\mathcal{G}}} = (\widehat{\lambda}_{C,j}^{(n)})_{j \in B_k}$ , the result then follows from comparing the expressions above to the definition of Kronecker product for every pair  $j, j' \in B_k$ .  $\blacksquare$

### A.3 Proof of Theorem 26

**Proof** Note that by the restrictive property of  $\mathcal{G}$  (Theorem 4), we have  $\text{vec} \widehat{\Lambda}_k^{\mathcal{G}} = (\widehat{\lambda}_{\text{Pa}(B_k, \mathcal{G}), j}^{(n)})_{j \in B_k}$  for  $k = 2, \dots, K$ . Using Theorem 23 with  $C = \text{Pa}(B_k, \mathcal{G})$ , we have the following asymptotic linear expansion

$$\widehat{\lambda}_{\text{Pa}(B_k, \mathcal{G}), j}^{(n)} - \lambda_{\text{Pa}(B_k, \mathcal{G}), j} = \frac{1}{n} \sum_{i=1}^n (\Sigma_{\text{Pa}(B_k, \mathcal{G})})^{-1} X_{\text{Pa}(B_k, \mathcal{G})}^{(i)} \varepsilon_j^{(i)} + O_p(n^{-1}). \quad (31)$$

The rest of computation follows similarly to the proof of Theorem 24.  $\blacksquare$

### A.4 Proof of Theorem 27

**Proof** Since  $S \in \mathbb{R}_{\text{PD}}^{n \times n}$ , by completing the square, we have

$$\begin{aligned} x^\top S x &= x_A^\top S_{A,A} x_A + x_B^\top S_{B,A} x_A + x_A^\top S_{A,B} x_B + x_B^\top S_{B,B} x_B \\ &\quad - x_A^\top S_{A,B} S_{B,B}^{-1} S_{B,A} x_A + x_A^\top S_{A,B} S_{B,B}^{-1} S_{B,A} x_A \\ &= x_A^\top (S_{A,A} - S_{A,B} S_{B,B}^{-1} S_{B,A}) x_A + (x_B + S_{B,B}^{-1} S_{B,A} x_A)^\top S_{B,B} (x_B + S_{B,B}^{-1} S_{B,A} x_A) \\ &\geq x_A^\top (S_{A,A} - S_{A,B} S_{B,B}^{-1} S_{B,A}) x_A = x_A^\top S_{A,B} x_A, \end{aligned}$$

where the equality holds if and only if  $x_B = -S_{B,B}^{-1} S_{B,A} x_A$ .  $\blacksquare$

## Appendix B. Proofs for graphical results

### B.1 Proof of Theorem 3

**Proof** Let the undirected path between  $j$  and  $k$  be  $p = \langle j = V_1, \dots, V_l = k \rangle$  with  $l > 1$ . First note that  $i$  is not on  $p$  because there is no undirected path between  $i$  and  $j$  in  $\mathcal{G}$ .

Further, since  $i \rightarrow j - V_2$  is in  $\mathcal{G}$ , by Meek rules R1 and R2 (Fig. 6 in Appendix C),  $i - V_2$  or  $i \rightarrow V_2$  is in  $\mathcal{G}$ . Since, by assumption, there is no undirected path from  $i$  to  $j$  in  $\mathcal{G}$ ,  $i - V_2 \notin U$ . Hence,  $i \rightarrow V_2 \in E$  and if  $l = 2$ , the statement of the lemma holds. If  $l > 2$ , we can apply the above reasoning iteratively until we obtain  $i \rightarrow V_l \in E$ . ■

### B.2 Proof of Theorem 5

**Proof** Let  $l \in D_k$ . Since  $D_k \subseteq B_k$ ,  $l \in B_k$ . Then by Theorem 4, we have that  $\text{Pa}(B_k) = \text{Pa}(l) \setminus B_k$ . Therefore,  $\text{Pa}(B_k) \subseteq \cup_{j \in D_k} \text{Pa}(j) \setminus B_k$  and furthermore,  $\text{Pa}(B_k) \subseteq \cup_{j \in D_k} \text{Pa}(j) \setminus D_k = \text{Pa}(D_k)$ . Hence, it suffices to show  $\text{Pa}(D_k) \subseteq \text{Pa}(B_k)$ .

We prove  $\text{Pa}(D_k) \subseteq \text{Pa}(B_k)$  by contradiction. Suppose there exists  $j \in \text{Pa}(D_k) \setminus \text{Pa}(B_k)$ . We know that  $\text{Pa}(D_k) = \cup_{l \in D_k} \text{Pa}(l) \setminus D_k$  and since  $D_k \subseteq B_k$ , for any  $l \in D_k$ ,  $l \in B_k$ . Using Theorem 4, we also know that  $\text{Pa}(B_k) = \text{Pa}(l) \setminus B_k$  for every  $l \in D_k$ . Therefore, if  $j \in \text{Pa}(D_k) \setminus \text{Pa}(B_k)$ , it follows that  $j \in B_k$ .

Additionally, by definition  $D = \text{An}(Y, \mathcal{G}_{V \setminus A})$  and  $D = \cup_{r=1}^K D_r$ . Therefore, if  $k = 1$ , then  $j \in A$ ; if  $k > 1$ ,  $j$  must be contained in  $\cup_{r=1}^{k-1} D_r$  or in  $A$ . If  $j \in A$ , then because condition of Theorem 2 holds,  $j \in A \cap B_k$  leads to a contradiction with Theorem 33 in Appendix C. Suppose  $k > 1$  and  $j \in \cup_{r=1}^{k-1} D_r$ . Because  $\cup_{r=1}^{k-1} D_r \subseteq \cup_{r=1}^{k-1} B_r$  and buckets  $\{B_1, \dots, B_K\}$  are disjoint, we have  $(\cup_{r=1}^{k-1} D_r) \cap B_k = \emptyset$ . However, this contradicts that  $j \in B_k$ . ■

### B.3 Proof of Theorem 16

**Proof** By construction, the undirected component of  $\bar{\mathcal{G}}$  remains the same as that of  $\mathcal{G}$ . Hence,  $\bar{\mathcal{G}}$  has the same bucket decomposition as  $\mathcal{G}$ . We only need to show that  $\bar{\mathcal{G}}$  is an MPDAG. It is enough to show that the edge orientations in  $\bar{\mathcal{G}}$  are closed under rules R1–R4 of Meek (1995) that are displayed in Fig. 6 of Appendix C. Note that since  $\mathcal{G}$  is an MPDAG it is closed under R1–R4. So if any of the left-hand-side graphs in Figure 6 are induced subgraphs of  $\bar{\mathcal{G}}$ , then at least one of the directed edges in these induced subgraphs must have been added in the construction of  $\bar{\mathcal{G}}$ .

Since the construction of  $\bar{\mathcal{G}}$  does not involve adding directed edges within a bucket, the left-hand-side of rules R3 and R4 in Figure 6 cannot appear as induced subgraphs of  $\bar{\mathcal{G}}$ . Hence, edge orientations in  $\bar{\mathcal{G}}$  are complete under rules R3 and R4.

Consider the left-hand-side of rule R1 in Figure 6,  $A \rightarrow B - C$ , for some  $A, B, C \in V$ . For  $A \rightarrow B - C$  to be an induced subgraph of  $\bar{\mathcal{G}}$ ,  $A \rightarrow B$  must have been added in the construction of  $\bar{\mathcal{G}}$  from  $\mathcal{G}$ . Hence,  $A$  and  $B$  would need to be in different buckets in  $V$  in  $\mathcal{G}$ . Since  $B$  and  $C$  are in the same bucket because of edge  $B - C$ ,  $A \rightarrow C$  would also be added to  $\mathcal{G}$  in the construction of  $\bar{\mathcal{G}}$ . Hence,  $A \rightarrow B - C$  will also not appear as an induced subgraph of  $\bar{\mathcal{G}}$  and edge orientations in  $\bar{\mathcal{G}}$  are also closed under R1.

Consider the left-hand-side of R2 in Figure 6, and suppose for a contradiction that  $A \rightarrow B \rightarrow C$  and  $A - C$  is an induced subgraph of  $\bar{\mathcal{G}}$  for some  $A, B, C \in V$ . Then  $A \rightarrow B$ ,  $B \rightarrow C$ , or both  $A \rightarrow B$  and  $B \rightarrow C$ , were added to  $\mathcal{G}$  in the construction of  $\bar{\mathcal{G}}$ . Because of  $A - C$ , suppose  $A$  and  $C$  are in the same bucket  $B_i$  for some  $i \in \{1, \dots, K\}$  in  $\mathcal{G}$ . Also, suppose  $B \in B_j$ . Because only directed edges between buckets are added,  $i \neq j$ .

Now,  $A \rightarrow B$  and  $B \rightarrow C$  cannot be both added to  $\mathcal{G}$  to construct  $\bar{\mathcal{G}}$ , because that would imply that  $i < j$  and  $j < i$ . By R1,  $B \rightarrow C - A$  cannot be an induced subgraph of MPDAG  $\mathcal{G}$ , so  $A \rightarrow B$  alone also could not be added to  $\mathcal{G}$ . Therefore,  $B \rightarrow C$  alone was added to  $\mathcal{G}$ . But  $C - A \rightarrow B$  is an induced subgraph of  $\mathcal{G}$ , so  $i < j$ , which contradicts the direction of  $B \rightarrow C$ . ■

## Appendix C. Graphical preliminaries

**Graphs, vertices, edges** A graph  $\mathcal{G} = (V, F)$  consists of a set of vertices (variables)  $V$  and a set of edges  $F$ . The graphs we consider are allowed to contain directed ( $\rightarrow$ ) and undirected ( $-$ ) edges and at most one edge between any two vertices. We can thus partition the set of edges  $F$  into a set of directed edges  $E$  and undirected edges  $U$  and denote graph  $\mathcal{G} = (V, F)$  as  $\mathcal{G} = (V, E, U)$ . The corresponding undirected graph is simply  $\mathcal{G}_U = (V, \emptyset, U)$ .

**Subgraphs and skeleton** An *induced subgraph*  $\mathcal{G}_{V'} = (V', F')$  of  $\mathcal{G} = (V, F)$  consists of  $V' \subseteq V$  and  $F' \subseteq F$  where  $F'$  are all edges in  $F$  between vertices in  $V'$ . A *skeleton* of a graph  $\mathcal{G} = (V, F)$  is an undirected graph  $\mathcal{G} = (V, F')$ , such that  $F'$  are undirected versions of all edges in  $F$ .

**Paths. Directed, undirected, causal, non-causal, proper paths** A *path*  $p$  from  $i$  to  $j$  in  $\mathcal{G}$  is a sequence of distinct vertices  $p = \langle i, \dots, j \rangle$  in which every pair of successive vertices are adjacent. A path consisting of undirected edges is an *undirected path*. A *directed path* from  $i$  to  $j$  is a path from  $i$  to  $j$  in which all edges are directed towards  $j$ , that is,  $i \rightarrow \dots \rightarrow j$ . We will use *causal path* instead of *directed path* when talking about causal graphs. Let  $p = \langle v_1, \dots, v_k \rangle$ ,  $k > 1$  be a path in  $\mathcal{G}$ ,  $p$  is a *possibly directed path* (*possibly causal path*) if no edge  $v_i \leftarrow v_j$ ,  $1 \leq i < j \leq k$  is in  $\mathcal{G}$ . Otherwise,  $p$  is a *non-causal path* in  $\mathcal{G}$  (see Definition 3.1 and Lemma 3.2 of Perković et al., 2017). A path from  $A$  to  $Y$  is *proper* (w.r.t.  $A$ ) if only its first vertex is in  $A$ .

**Directed cycles** A directed path from  $i$  to  $j$  and the edge  $j \rightarrow i$  form a *directed cycle*.

**Colliders, shields and definite status paths** If a path  $p$  contains  $i \rightarrow j \leftarrow k$  as a subpath, then  $j$  is a *collider* on  $p$ . A path  $\langle i, j, k \rangle$  is an (*un*)*shielded triple* if  $i$  and  $k$  are (not) adjacent. A path is *unshielded* if all successive triples on the path are unshielded. A node  $v_j$  is a *definite non-collider* on a path  $p$  if there is at least one edge out of  $v_j$  on  $p$ , or if  $v_{j-1} - v_j - v_{j+1}$  is a subpath of  $p$  and  $\langle v_{j-1}, v_j, v_{j+1} \rangle$  is an unshielded triple. A node is of *definite status* on a path if it is a collider, a definite non-collider or an endpoint on the path. A path  $p$  is of definite status if every node on  $p$  is of definite status.

**Subsequences and subpaths** A *subsequence* of a path  $p$  is obtained by deleting some nodes from  $p$  without changing the order of the remaining nodes. A subsequence of a



path is not necessarily a path. For a path  $p = \langle v_1, v_2, \dots, v_m \rangle$ , the *subpath* from  $v_i$  to  $v_k$  ( $1 \leq i \leq k \leq m$ ) is the path  $p(v_i, v_k) = \langle v_i, v_{i+1}, \dots, v_k \rangle$ .

**Ancestral relations** If  $i \rightarrow j$ , then  $i$  is a *parent* of  $j$ , and  $j$  is a *child* of  $i$ . If there is a causal path from  $k$  to  $l$ , then  $k$  is an *ancestor* of  $l$ , and  $l$  is a *descendant* of  $k$ . If there is a possibly causal path from  $k$  to  $l$ , then  $k$  is a *possible ancestor* of  $l$ , and  $l$  is a *possible descendant* of  $k$ . We use the convention that every vertex is a descendant, ancestor, possible ancestor and possible descendant of itself. The sets of parents, ancestors, descendants and possible descendants of  $i$  in  $\mathcal{G}$  are denoted by  $\text{Pa}(i, \mathcal{G})$ ,  $\text{An}(i, \mathcal{G})$ ,  $\text{De}(i, \mathcal{G})$  and  $\text{PossDe}(i, \mathcal{G})$  respectively. For a set of vertices  $A$ , we let  $\text{Pa}(A, \mathcal{G}) = (\cup_{i \in A} \text{Pa}(i, \mathcal{G})) \setminus A$ , whereas,  $\text{An}(A, \mathcal{G}) = \cup_{i \in A} \text{An}(i, \mathcal{G})$ ,  $\text{De}(A, \mathcal{G}) = \cup_{i \in A} \text{De}(i, \mathcal{G})$  and  $\text{PossDe}(A, \mathcal{G}) = \cup_{i \in A} \text{PossDe}(i, \mathcal{G})$ .

**DAGs, PDAGs** A *directed graph* contains only directed edges. A *partially directed graph* may contain both directed and undirected edges. A directed graph without directed cycles is a *directed acyclic graph* (DAG). A *partially directed acyclic graph* (PDAG) is a partially directed graph without directed cycles.

**Blocking and d-separation** (See Definition 1.2.3 of Pearl (2009) and Lemma C.1 of Henckel et al. (2022)). Let  $Z$  be a set of vertices in an PDAG  $\mathcal{G} = (V, E, U)$ . A definite status path  $p$  is *blocked* by  $Z$  if (i)  $p$  contains a non-collider that is in  $Z$ , or (ii)  $p$  contains a collider  $C$  such that no descendant of  $C$  is in  $Z$ . A definite status path that is not blocked by  $Z$  is *open* given  $Z$ . If  $A, B$  and  $Z$  are three pairwise disjoint sets of nodes in a PDAG  $\mathcal{G} = (V, E, U)$ , then  $Z$  *d-separates*  $A$  from  $B$  in  $\mathcal{G}$  if  $Z$  blocks every definite status path between any node in  $A$  and any node in  $B$  in  $\mathcal{G}$ .

**CPDAGs, MPDAGs** Several DAGs can encode the same d-separation relationships. Such DAGs form a *Markov equivalence class* which is uniquely represented by a *completed partially directed acyclic graph* (CPDAG) (Meek, 1995; Andersson et al., 1997). A PDAG  $\mathcal{G} = (V, E, U)$  is a *maximally oriented* PDAG (MPDAG) if it is closed under orientation rules R1-R4 of (Meek, 1995), presented in Figure 6. The MPDAG can then be alternatively defined as any PDAG that does not contain graphs on the left-hand side of each orientation rule as induced subgraphs. Both DAGs and CPDAGs are types of MPDAGs (Meek, 1995).

**Background knowledge and constructing MPDAGs** A PDAG  $\mathcal{G}'$  is *represented* by another PDAG  $\mathcal{G}$  (equivalently  $\mathcal{G}$  represents  $\mathcal{G}'$ ) if  $\mathcal{G}'$  and  $\mathcal{G}$  have the same adjacencies and unshielded colliders and every directed edge  $i \rightarrow j$  in  $\mathcal{G}$  is also in  $\mathcal{G}'$ . Let  $R$  be a set of directed edges representing background knowledge. Algorithm 1 of Meek (1995) describes how to incorporate background knowledge  $R$  in an MPDAG  $\mathcal{G}$ . If Algorithm 1 does not return a FAIL, then it returns a new MPDAG  $\mathcal{G}'$  that is represented by  $\mathcal{G}$ . Background knowledge  $R$  is *consistent* with MPDAG  $\mathcal{G}$  if and only if Algorithm 1 does not return a FAIL (Meek, 1995).

**Remark 32** *The MPDAG output by  $\text{ConstructMPDAG}(\mathcal{G}, R)$  is the same independent of the ordering of edges in  $R$ . This stems from the fact that the orientation rules of Meek (1995) are necessary and sufficient for the construction of an MPDAG given a set of adjacencies and unshielded colliders.*

**$\mathcal{G}$  and  $[\mathcal{G}]$**  If  $\mathcal{G}$  is a MPDAG, then  $[\mathcal{G}]$  denotes every DAG represented by  $\mathcal{G}$ .




Figure 6: The orientation rules from Meek (1995). If the graph on the left-hand side of a rule is an induced subgraph of a PDAG  $\mathcal{G}$ , then *orient* the blue undirected edge (–) as shown on the right-hand side of the rule. Hence, the graphs on the left-hand side of each rule are not allowed to be induced subgraphs of an MPDAG.

---

**Algorithm 1: ConstructMPDAG**, (Meek, 1995; Perković et al., 2017)

---

**Data:** MPDAG  $\mathcal{G}$ , background knowledge  $R$

**Result:** MPDAG  $\mathcal{G}'$  or FAIL

```

1 Let  $\mathcal{G}' = \mathcal{G}$ ;
2 while  $R \neq \emptyset$  do
3   Choose an edge  $\{X \rightarrow Y\}$  in  $R$ ;
4    $R = R \setminus \{X \rightarrow Y\}$ ;
5   if  $\{X - Y\}$  or  $\{X \rightarrow Y\}$  is in  $\mathcal{G}'$  then
6     Orient  $\{X \rightarrow Y\}$  in  $\mathcal{G}'$ ;
7     Close the edge orientations under the rules in Figure 6 in  $\mathcal{G}'$ ;
8   else
9     FAIL;
10  end
11 end
    
```

---

**Causal and partial causal ordering of vertices** A total ordering,  $<$ , of vertices  $V' \subseteq V$  is *consistent* with a DAG  $\mathcal{D} = (V, E, \emptyset)$  and called a *causal ordering* of  $V'$  if for every  $i, j \in V'$ , such that  $i < j$  and such that  $i$  and  $j$  are adjacent in  $\mathcal{D}$ ,  $i \rightarrow j$  is in  $\mathcal{D}$ . There can be more than one causal ordering of  $V'$  in a DAG  $\mathcal{D} = (V, E, \emptyset)$ . For example, in DAG  $i \leftarrow j \rightarrow k$  both orderings  $j < i < k$  and  $j < k < i$  are consistent.

Since an MPDAG may contain undirected edges, there is generally no unique causal ordering of vertices in an MPDAG. Instead, we define a *partial causal ordering*,  $<$ , of a vertex set  $V'$ ,  $V' \subseteq V$  in an MPDAG  $\mathcal{G} = (V, E, U)$  as a total ordering of pairwise disjoint vertex sets  $A_1, \dots, A_k$ ,  $k \geq 1$ ,  $\cup_{i=1}^k A_i = V'$ , that satisfy the following: if  $A_i < A_j$  and there is an edge between  $i \in A_i$  and  $j \in A_j$  in  $\mathcal{G}$ , then  $i \rightarrow j$  is in  $\mathcal{G}$ .

**Buckets and bucket decomposition** Algorithm 2 describes how to obtain an ordered bucket decomposition for a set of vertices  $V$  in an MPDAG  $\mathcal{G} = (V, E, U)$ . By Perković

(2020, Lemma 1), the ordered list of buckets output by Algorithm 2 is a partial causal ordering of  $V$  in  $\mathcal{G}$ .

---

**Algorithm 2:** Partial causal ordering (Perković, 2020)

---

**input** : vertex set  $V$  in MPDAG  $\mathcal{G}=(V, E, U)$  and MPDAG  $\mathcal{G}$ .  
**output:** An ordered list  $B=(B_1, \dots, B_k), k \geq 1$ , of the bucket decomposition of  $V$  in  $\mathcal{G}$ .

- 1 Let  $\mathcal{G}_U$  denote the undirected subgraph of  $\mathcal{G}$ ;
- 2 Let  $ConComp$  be the bucket decomposition (i.e., maximal connected components) of  $V$  in  $\mathcal{G}_U$ ;
- 3 Let  $B$  be an empty list;
- 4 **while**  $ConComp \neq \emptyset$  **do**
- 5     Let  $C$  be any element from  $ConComp$ ;
- 6     Let  $\bar{C}$  be the set of vertices in  $ConComp$  that are not in  $C$ ;
- 7     **if** all edges between  $C$  and  $\bar{C}$  are into  $C$  in  $\mathcal{G}$  **then**
- 8         Remove  $C$  from  $ConComp$ ;
- 9         Add  $C$  to the beginning of  $B$ ;
- 10    **end**
- 11 **end**
- 12 **return**  $B$ ;

---

**Lemma 33** (see Lemma D.1 (i) of Perković, 2020) *Let  $A$  and  $Y$  be disjoint node sets in MPDAG  $\mathcal{G} = (V, E, U)$ . Suppose that there is no proper possibly causal path from  $A$  to  $Y$  that starts with an undirected edge in  $\mathcal{G}$ , that is, suppose that the criterion in Theorem 2 is satisfied. Further, let  $D = An(Y, \mathcal{G}_{V \setminus A})$  and  $D = \bigcup_{i=1}^K D_i$  for  $D_i = D \cap B_i, i = 1, \dots, K$ , where  $B_1, \dots, B_K$  is the bucket decomposition of  $V$ . Then for all  $i \in \{1, \dots, K\}$ , there is no proper possibly causal path from  $A$  to  $B_i$  that starts with an undirected edge in  $\mathcal{G}$ .*

## Appendix D. Additional simulation results

In this section, we report additional simulation results.

### D.1 With estimated graphs

The setup is the same as Section 7.1 of main text, but we replace the true CPDAG with the CPDAG estimated with the greedy equivalence search algorithm (Chickering, 2002) based on the same sample. The relative squared errors of the contending estimators are shown in Fig. 7 and are summarized in Table 4. Compared to the results with the true CPDAG, the performance improvement of  $\mathcal{G}$ -regression is more modest but still matters in practice. The reduced improvement is due to the error in estimating the graph, which diminishes as  $n$  increases.

Table 4: Geometric average (brackets: median) of relative squared errors compared to  $\mathcal{G}$ -regression when CPDAGs are estimated

$ A $	$ V  = 20$				$ V  = 50$				$ V  = 100$			
	$n = 100$		$n = 1000$		$n = 100$		$n = 1000$		$n = 100$		$n = 1000$	
<b>adj.O</b>												
1	1.0	(1.0)	1.0	(1.0)	1.2	(1.0)	1.3	(1.0)	1.8	(1.1)	1.6	(1.0)
2	2.0	(1.1)	3.1	(1.2)	2.4	(1.3)	3.1	(1.4)	3.2	(1.9)	3.7	(2.0)
3	3.3	(1.7)	5.2	(2.7)	4.0	(2.4)	5.9	(2.8)	4.7	(2.5)	5.5	(2.8)
4	4.6	(2.2)	7.9	(4.2)	5.0	(2.1)	9.0	(5.7)	10	(5.9)	8.9	(5.6)
<b>IDA.M</b>												
1	2.9	(1.4)	4.1	(1.4)	4.5	(2.7)	10	(5.7)	7.3	(4.5)	18	(11)
2	4.2	(2.0)	6.6	(2.1)	7.3	(4.8)	14	(7.2)	13	(7.9)	22	(14)
3	6.2	(3.1)	6.8	(2.5)	12	(7.1)	16	(8.3)	15	(10)	28	(18)
4	9.5	(5.6)	9.0	(3.1)	13	(10)	20	(12)	19	(14)	37	(26)
<b>IDA.R</b>												
1	2.9	(1.4)	4.1	(1.4)	4.5	(2.7)	10	(5.7)	7.3	(4.5)	18	(11)
2	2.7	(1.3)	4.6	(1.2)	4.5	(2.3)	9.6	(4.0)	8.5	(5.9)	15	(9.5)
3	3.1	(1.5)	4.1	(1.2)	5.8	(3.0)	7.8	(2.5)	7.6	(5.2)	14	(8.9)
4	3.6	(1.6)	4.2	(1.3)	4.9	(2.8)	8.2	(3.6)	8.1	(5.4)	15	(10)




Figure 7: Violin plots for the relative squared errors of contending estimators (‘ $\cdot$ ’: geometric mean, ‘+’: median). The estimated CPDAGs are provided to the estimators.

### D.2 With rescaled coefficients

It can be argued that the data generating mechanism described in Section 7.1 is less realistic because the variance of a vertex grows with its topological ordering. To remedy this, we rescale the coefficients  $\{\gamma_{ij}\}$  properly such that the variances are roughly equalized along the topological order. The results (based on true graphs) are reported in Table 5.

Table 5: Geometric average (brackets: median) of relative squared errors compared to  $\mathcal{G}$ -regression (coefficients rescaled)

A	V  = 20				V  = 50				V  = 100			
	n = 100		n = 1000		n = 100		n = 1000		n = 100		n = 1000	
<b>adj.0</b>												
1	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)
2	1.1	(1.0)	1.1	(1.0)	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)
3	1.2	(1.0)	1.1	(1.0)	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)
4	1.2	(1.0)	1.1	(1.0)	1.1	(1.0)	1.1	(1.0)	1.1	(1.0)	1.0	(1.0)
<b>IDA.M</b>												
1	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)
2	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)	1.2	(1.0)	1.0	(1.0)
3	1.2	(1.0)	1.1	(1.0)	1.2	(1.0)	1.0	(1.0)	1.2	(1.0)	1.0	(1.0)
4	1.2	(1.0)	1.0	(1.0)	1.2	(1.0)	1.0	(1.0)	1.2	(1.0)	1.0	(1.0)
<b>IDA.R</b>												
1	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)	1.0	(1.0)
2	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)	1.2	(1.0)	1.0	(1.0)
3	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)
4	1.2	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)	1.1	(1.0)	1.0	(1.0)

## References

- Takeshi Amemiya. *Advanced Econometrics*. Harvard University Press, 1985.
- Theodore Wilbur Anderson and Ingram Olkin. Maximum-likelihood estimation of the parameters of a multivariate normal distribution. *Linear algebra and its applications*, 70: 147–171, 1985.
- Steen A. Andersson, David Madigan, and Michael D. Perlman. A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25:505–541, 1997.
- Peter J. Bickel, Chris A. J. Klaassen, Ya’acov Ritov, and Jon A. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*, volume 4. Johns Hopkins University Press, Baltimore, 1993.
- Kenneth A. Bollen. *Structural Equations with Latent Variables*. Wiley, New York, 1989.
- Wenyu Chen, Mathias Drton, and Y. Samuel Wang. On causal discovery with an equal-variance assumption. *Biometrika*, 106(4):973–980, 2019.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov):507–554, 2002.
- A. Philip Dawid. Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika*, 68(1):265–274, 1981.
- Mathias Drton. Computing all roots of the likelihood equations of seemingly unrelated regressions. *Journal of Symbolic Computation*, 41(2):245–254, 2006.
- Mathias Drton. Algebraic problems in structural equation modeling. In *The 50th Anniversary of Gröbner Bases*, pages 35–86. Mathematical Society of Japan, 2018.

- Mathias Drton and Michael Eichler. Maximum likelihood estimation in Gaussian chain graph models under the alternative Markov property. *Scandinavian Journal of Statistics*, 33(2):247–257, 2006.
- Mathias Drton and Marloes H. Maathuis. Structure learning in graphical modeling. *Annual Review of Statistics and Its Application*, 4:365–393, 2017.
- Mathias Drton and Thomas S. Richardson. Multimodality of the likelihood in the bivariate seemingly unrelated regressions model. *Biometrika*, 91(2):383–392, 2004.
- Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on algebraic statistics*, volume 39. Springer Science & Business Media, 2008.
- Mathias Drton, Michael Eichler, and Thomas S. Richardson. Computing maximum likelihood estimates in recursive linear models with correlated errors. *Journal of Machine Learning Research*, 10(81):2329–2348, 2009.
- Mathias Drton, Rina Foygel, and Seth Sullivant. Global identifiability of linear structural equation models. *The Annals of Statistics*, 39(2):865–886, 2011.
- Marco Eigenmann, Preetam Nandy, and Marloes H. Maathuis. Structure learning of linear Gaussian structural equation models with weak edges. In *Proceedings of the 33rd Annual Conference on Uncertainty in Artificial Intelligence (UAI-17)*, 2017.
- Zhuangyan Fang and Yangbo He. IDA with background knowledge. In *Proceedings of the 36th Annual Conference on Uncertainty in Artificial Intelligence (UAI-20)*, 2020.
- F. Richard Guo and Emilija Perkovic. Minimal enumeration of all possible total effects in a Markov equivalence class. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS-21)*, volume 130, pages 2395–2403, 2021.
- Shantanu Gupta, Zachary C. Lipton, and David Childers. Estimating treatment effects with observed confounders and mediators. *arXiv preprint arXiv:2003.11991*, 2020.
- Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054, 1982.
- Alan Hauser and Peter Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464, 2012.
- Leonard Henckel, Emilija Perković, and Marloes H. Maathuis. Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, to appear, 2022.
- Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition, 2012.

- Patrik O. Hoyer, Aapo Hyvarinen, Richard Scheines, Peter L. Spirtes, Joseph Ramsey, Gustavo Lacerda, and Shohei Shimizu. Causal discovery of linear acyclic models with arbitrary distributions. In *Proceedings of the 24th Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 282–289, 2008.
- 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.
- Tjalling C. Koopmans and Olav Reiersøl. The identification of structural characteristics. *The Annals of Mathematical Statistics*, 21(2):165–181, 1950.
- Manabu Kuroki and Zhihong Cai. Selection of identifiability criteria for total effects by using path diagrams. In *Proceedings of the 20th 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.
- Manabu Kuroki and Hisayoshi Nanmo. Variance formulas for estimated mean response and predicted response with external intervention based on the back-door criterion in linear structural equation models. *AStA Advances in Statistical Analysis*, pages 1–19, 2020.
- Steffen L. Lauritzen. *Graphical Models*. Oxford University Press, New York, 1996.
- Jianzhou Liu. Some Löwner partial orders of Schur complements and Kronecker products of matrices. *Linear algebra and its applications*, 291(1-3):143–149, 1999.
- Marloes H. Maathuis and Diego Colombo. A generalized 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.
- Daniel Marbach, Thomas Schaffter, Dario Floreano, Robert J. Prill, and Gustavo Stolovitzky. The DREAM4 in-silico network challenge. Draft, version 0.3 <http://gnw.sourceforge.net/resources/DREAM4%20in%20silico%20challenge.pdf>, 2009a.
- Daniel Marbach, Thomas Schaffter, Claudio Mattiussi, and Dario Floreano. Generating realistic in silico gene networks for performance assessment of reverse engineering methods. *Journal of Computational Biology*, 16(2):229–239, 2009b.
- Christopher Meek. Causal inference and causal explanation with background knowledge. In *Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 403–410, 1995.
- 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. Comment: graphical models, causality and intervention. *Statistical Science*, 8(3):266–269, 1993.
- Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- Judea Pearl. *Causality*. Cambridge University Press, Cambridge, 2nd edition, 2009.
- Judea Pearl and Thomas S. Verma. A theory of inferred causation. In *Studies in Logic and the Foundations of Mathematics*, volume 134, pages 789–811. Elsevier, 1995.
- Emilija Perković. Identifying causal effects in maximally oriented partially directed acyclic graphs. In *Proceedings of the 36th Annual Conference on Uncertainty in Artificial Intelligence (UAI-20)*, 2020.
- Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. A complete generalized adjustment criterion. In *Proceedings of the 31st Annual Conference on Uncertainty in Artificial Intelligence (UAI-15)*, 2015.
- Emilija Perković, Markus Kalisch, and Marloes H. Maathuis. Interpreting and using CPDAGs with background knowledge. In *Proceedings of the 33rd Annual Conference on Uncertainty in Artificial Intelligence (UAI-17)*, 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.
- Jonas Peters and Peter Bühlmann. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2020. URL <https://www.R-project.org/>.
- James M. 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*, 7:1393–1512, 1986.
- Dominik Rothenhäusler, Jan Ernest, and Peter Bühlmann. Causal inference in partially linear structural equation models. *The Annals of Statistics*, 46(6A):2904 – 2938, 2018.
- Andrea Rotnitzky and Ezequiel Smucler. Efficient adjustment sets for population average causal treatment effect estimation in graphical models. *Journal of Machine Learning Research*, 21(188):1–86, 2020.
- John D. Sargan. The estimation of economic relationships using instrumental variables. *Econometrica: Journal of the Econometric Society*, pages 393–415, 1958.
- Richard Scheines, Peter Spirtes, Clark Glymour, Christopher Meek, and Thomas Richardson. The TETRAD project: constraint based aids to causal model specification. *Multivariate Behavioral Research*, 33(1):65–117, 1998.



- Shohei Shimizu. LiNGAM: Non-Gaussian methods for estimating causal structures. *Behaviormetrika*, 41(1):65–98, 2014.
- 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(72):2003–2030, 2006.
- Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvärinen, Yoshinobu Kawahara, Takashi Washio, Patrik O. Hoyer, and Kenneth Bollen. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *Journal of Machine Learning Research*, 12:1225–1248, 2011.
- Galen R. Shorack. *Probability for Statisticians*. Springer, 2000.
- Ilya Shpitser, Tyler VanderWeele, and James M. Robins. On the validity of covariate adjustment for estimating causal effects. In *Proceedings of the 26th Annual Conference on Uncertainty in Artificial Intelligence (UAI-10)*, pages 527–536, 2010.
- Ezequiel Smucler, Facundo Sapienza, and Andrea Rotnitzky. Efficient adjustment sets in causal graphical models with hidden variables. *arXiv preprint arXiv:2004.10521*, 2020.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition, 2000.
- Robert H. Strotz and H. O. A. Wold. Recursive vs. nonrecursive systems: An attempt at synthesis (part I of a triptych on causal chain systems). *Econometrica*, 28(2):417–427, 1960.
- Seth Sullivant, Kelli Talaska, and Jan Draisma. Trek separation for Gaussian graphical models. *The Annals of Statistics*, 38(3):1665–1685, 2010.
- Anastasios Tsiatis. *Semiparametric Theory and Missing Data*. Springer, New York, 2006.
- Mark J. van der Laan and Sherri Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.
- Aad W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.
- Yuhao Wang, Liam Solus, Karren Dai Yang, and Caroline Uhler. Permutation-based causal inference algorithms with interventions. In *Advances in Neural Information Processing Systems 30*, pages 5822–5831. 2017.
- Janine Witte, Leonard Henckel, Marloes H Maathuis, and Vanessa Didelez. On efficient adjustment in causal graphs. *Journal of Machine Learning Research*, 21(246):1–45, 2020.
- Sewall Wright. Correlation and causation. *Journal of Agricultural Research*, 20:557–585, 1921.
- Sewall Wright. The method of path coefficients. *The Annals of Mathematical Statistics*, 5(3):161–215, 1934.