

---

# Confidence in Causal Discovery with Linear Causal Models

---

David Strieder<sup>1</sup>

Tobias Freidling<sup>1,2</sup>

Stefan Haffner<sup>1</sup>

Mathias Drton<sup>1</sup>

<sup>1</sup>Chair of Mathematical Statistics, Technical University of Munich, Germany

<sup>2</sup>DPMMS, University of Cambridge, UK

## Abstract

Structural causal models postulate noisy functional relations among a set of interacting variables. The causal structure underlying each such model is naturally represented by a directed graph whose edges indicate for each variable which other variables it causally depends upon. Under a number of different model assumptions, it has been shown that this causal graph and, thus also, causal effects are identifiable from mere observational data. For these models, practical algorithms have been devised to learn the graph. Moreover, when the graph is known, standard techniques may be used to give estimates and confidence intervals for causal effects. We argue, however, that a two-step method that first learns a graph and then treats the graph as known yields confidence intervals that are overly optimistic and can drastically fail to account for the uncertain causal structure. To address this issue we lay out a framework based on test inversion that allows us to give confidence regions for total causal effects that capture both sources of uncertainty: causal structure and numerical size of nonzero effects. Our ideas are developed in the context of bivariate linear causal models with homoscedastic errors, but as we exemplify they are generalizable to larger systems as well as other settings such as, in particular, linear non-Gaussian models.

## 1 INTRODUCTION

Anticipating the cause and effect of actions is a task the human brain is able to master every day, yet it is challenging to devise statistical methods that reliably infer cause-effect relations from available data. The field of causal discovery seeks to address this challenge by clarifying when it is theoretically possible to infer a causal effect and by developing

practical methods to estimate the effect from data [Pearl, 2009, Spirtes et al., 2000]. A widely studied approach adopts the paradigm of structural causal models, in which each variable is a function of a subset of other variables (its causes) and a stochastic error term; see also Peters et al. [2017] or Maathuis et al. [2019, Part IV]. The causal perspective results from viewing these functions as mechanisms that assign values based on the values of causes. Algorithms for causal discovery infer the structure of such a causal model, which is naturally represented by a directed graph whose edges point from causes to effects. Moreover, for a known causal graph, standard methods give point estimates and confidence intervals for a causal effect of interest; at least, this is the case in the causally sufficient setting, where all relevant variables have been observed.

It is straightforward to combine a method that learns a graph with a method to subsequently estimate and make confidence statements about causal effects. However, such a two-step approach tacitly conditions away the uncertainty that arises from the data-driven model choice and ignores the uncertainty that exists with respect to the causal structure. As a result, this approach is overly optimistic in its conclusions about existence and strength of causal effects. Despite the extensive literature that exists on causal discovery, we are not aware of prior work that accounts for uncertainty in the causal structure when providing confidence statements about the inferred causal direction or effect size. Here, the term ‘confidence’ is used in the technical sense of confidence sets with a given desired frequentist coverage probability. An entirely different approach that we do not consider here would be to form Bayesian credible sets. Indeed, a number of authors have pursued Bayesian approaches, but primarily with a focus on the uncertainty in the graphical structure as opposed to causal effects; see Hoyer and Hyttinen [2009], Claassen and Heskes [2012], and Cao et al. [2019] for three selected examples.

In this article we present a framework to construct confidence intervals for the total causal effect that account for the uncertainty inherent in the data-driven selection of a causal

model. The setting we focus on pertains to the case in which we only have access to observational data but consider restricted structural causal models for which causal structure and effects are nevertheless identifiable. Specifically, we consider the simplest such setting, namely, linear structural equation models with errors that are homoscedastic, i.e., of equal variance [Peters and Bühlmann, 2014]. Their causal ordering is identifiable as it corresponds to an ordering of conditional variances [Chen et al., 2019, Ghoshal and Honorio, 2018]. Our presentation focuses on this setting and also the two-variable case, but the framework we lay out in Section 3 is general and can be applied to larger systems as well as other model classes, in particular, linear models with non-Gaussian errors, as we also discuss in Sections 4-5.

## 2 BACKGROUND

This section reviews linear structural equation models and the total causal effect that is the central object of study. Moreover, we discuss the difficulties of applying resampling methods for statistical inference and review the test-inversion approach.

### 2.1 STRUCTURAL EQUATION MODELS

Consider observational data in the form of a sample of independent copies of a random vector  $X = (X_1, \dots, X_d)$  which, without loss of generality, is assumed to have zero mean. *Linear structural equation models* (LSEMs) assume that  $X$  solves the equation system

$$X_j = \sum_{i \neq j} \beta_{ji} X_i + \epsilon_j, \quad j = 1, \dots, d, \quad (1)$$

where  $B := [\beta_{ji}]_{j,i=1}^d$  are unknown parameters that constitute direct causal effects between the variables, and the  $\epsilon_j$  are independent error terms with mean zero. Following a line of work initiated by Peters and Bühlmann [2014], we further assume the errors to be homoscedastic, that is, for an unknown variance parameter  $\sigma^2 \in (0, \infty)$  we have

$$\text{Var}[\epsilon_1] = \dots = \text{Var}[\epsilon_d] = \sigma^2. \quad (2)$$

Each specific LSEM restricts a subset of the parameters  $\beta_{ji}$  to be zero. Put differently, each model is associated to a directed graph  $\mathcal{G}$  and constrains  $\beta_{ji} = 0$  whenever  $\mathcal{G}$  does not contain edge  $i \rightarrow j$ . As in related work, we assume  $\mathcal{G}$  to be a *directed acyclic graph* (DAG). Then the matrix  $B$  is permutation similar to a strictly lower triangular matrix, and system (1) admits the unique solution  $X = (I_d - B)^{-1} \epsilon$ , where  $I_d$  is the identity. Hence,  $X$  has covariance matrix

$$\mathbb{E}[XX^T] = \sigma^2 (I_d - B)^{-1} (I_d - B)^{-T}. \quad (3)$$

### 2.2 TOTAL CAUSAL EFFECT

In the causal interpretation of LSEMs the equations in (1) are viewed as making assignments, with the variable on the left-hand side being assigned the value specified on the right-hand side. In this framework, the effect of an experimental intervention that externally sets the value of  $X_i$  to  $x_i$  is then captured by replacing the  $i$ th equation in (1) by  $X_i = x_i$ . In probabilistic notation this is expressed as  $\text{do}(X_i = x_i)$ , see Pearl [2009].

Our interest is in the *total causal effect* that an intervention on variable  $X_i$  has on another variable  $X_j$ . In linear models we may quantify this effect by considering a unit change in the intervention value  $x_i$ . The total effect of  $X_i$  on  $X_j$  is then

$$\mathcal{C}(i \rightarrow j) := \frac{d}{dx_i} \mathbb{E}[X_j | \text{do}(X_i = x_i)] = (I_d - B)^{-1}_{ji}.$$

Note that  $\mathcal{C}(i \rightarrow j) = 0$  if there does not exist a directed path from  $i$  to  $j$  in the underlying DAG  $\mathcal{G}$ .

Our goal in this article is to construct confidence intervals for  $\mathcal{C}(i \rightarrow j)$  when the underlying causal structure is unknown and has to be learned. We emphasize that this is a well-defined problem when the causal structure is identifiable, as in the homoscedastic setting laid out in Section 2.1. Indeed, in such a setting, every feasible distribution (or here simply, covariance matrix) uniquely determines a minimal causal graph which entails a unique value for  $\mathcal{C}(i \rightarrow j)$ .

Lacking alternative methods, a commonly used “naive” approach is to form confidence intervals by splitting the two involved tasks: First, the model, i.e., pattern of zero-entries of  $B$ , is estimated and, second, confidence intervals for the causal effects within the model are derived. This procedure, however, tacitly conditions away the uncertainty that arises from the data-driven model choice. For this reason, such naive confidence intervals often have poor coverage probabilities, especially under high uncertainty with respect to the model. For example, consider a bivariate example with a small true causal effect in one of the two possible causal directions. If the wrong causal ordering is learned, which happens with probability close to 0.5, one concludes with certainty yet wrongly that the causal effect is zero or formally one concludes, no matter the significance level, the confidence interval is the singleton set 0. Hence, almost half the time the confidence interval will not cover the true non-zero parameter.

In the remainder of the paper we propose methods to address this issue.

### 2.3 RESAMPLING

Bootstrapping, subsampling and other resampling procedures are often applied to construct confidence intervals

based on estimators whose sampling distribution is difficult to derive. Resampling also offers a seemingly straightforward solution to the problem considered here. This attempted solution proceeds by computing for each resampled data set a causal effect estimate that is obtained by concatenating a consistent model selection method that learns the graph and a consistent estimator of the causal effect in the learned model. This concatenation is well defined by the identifiability of homoscedastic LSEMs [Peters and Bühlmann, 2014, Chen et al., 2019], and it yields a function  $T_{ij}$  such that  $T_{ij}(\Sigma) = \mathcal{C}(i \rightarrow j)$  for all covariance matrices  $\Sigma = \text{Var}[X_1, \dots, X_d]$  that come from a homoscedastic LSEM specified by (1) and (2). However, bootstrap procedures that evaluate  $T_{ij}$  on resampled data are not valid when there is non negligible uncertainty about the model. Indeed, bootstrap/subsampling procedures may fail drastically when the mapping  $T_{ij}$  lacks smoothness [Andrews and Guggenberger, 2010, Drton and Williams, 2011]. Our simulations in Section 3.4 validate those problems as can be seen in Table 1, even in our simplified two-variable setting with generous sample size bootstrap methods do not achieve the required coverage probability. We emphasize that Drton and Williams [2011] demonstrate that for complex composite hypotheses settings (as in our case), the asymptotic behavior of bootstrap tests and confidence intervals is difficult to predict, even in low dimensions.

The subtleties for homoscedastic LSEMs stem from the fact that the set of covariance matrices that are associated to at least one possible DAG is a union of smooth manifolds and singular at the intersections of the manifolds. Figure 1 depicts this for  $d = 2$  variables, where the two largest graphs  $1 \rightarrow 2$  and  $1 \leftarrow 2$  each define a 2-dimensional subset of the 3-dimensional cone of covariance matrices. The singular points are where the two models meet, and their existence invalidates bootstrapping as a method for correctly capturing model uncertainty. This fact is again underlined by our simulations in Section 3.4, see Table 1. Coverage failure occurs for small true causal effects and therefore further discredits bootstrapping as a method for calculating reliable confidence intervals for causal effects in practice.

Let the cone of positive definite matrices be denoted by

$$\mathcal{M} := \{\Sigma \in \mathbb{R}^{d \times d} : \Sigma^T = \Sigma, \Sigma \text{ positive definite}\}.$$

It is possible to define continuous extensions  $\tilde{T} : \mathcal{M} \rightarrow \mathbb{R}$  for  $T_{ij}$ . However, in our exploration of such continuous extensions for  $d = 2$  variables, we were unable to give extensions for which there exists a unique scaling sequence  $\tau_n$  such that  $\tau_n(\tilde{T}(\hat{\Sigma}) - \tilde{T}(\Sigma))$  always converges to a nondegenerate limit – here,  $\hat{\Sigma}$  is the empirical covariance matrix. Convergence to nondegenerate limits is required for subsampling methods [Politis et al., 1999].




Figure 1: The set of  $2 \times 2$  covariance matrices coming from LSEMs with homoscedastic errors.

## 2.4 INVERSION OF TESTS

In order to circumvent the difficulties posed by the non-smooth nature of the causal effect of interest, we develop in the following sections an approach that leverages the duality between statistical hypothesis tests and confidence regions. Let  $\alpha \in (0, 1)$  be a fixed significance level. Suppose that for each attainable value  $\psi$  of the causal effect  $\mathcal{C}(i \rightarrow j)$  we have a level  $\alpha$  test of the hypothesis that the effect is indeed  $\psi$ . Let  $A(\psi)$  be the acceptance region, i.e., the set of all data sets for which the test does not reject  $\psi$  as a hypothesized causal effect. Then a  $(1 - \alpha)$ -confidence region for  $\mathcal{C}(i \rightarrow j)$  given the data set  $x$  is obtained as

$$C(x) := \{\psi : x \in A(\psi)\}.$$

This approach shifts the burden to the construction of suitable tests of a hypothesized causal effect. Without knowledge of the precise model, this remains challenging. In the next section we present three concrete solutions based on likelihood ratio tests of order constraints [Silvapulle and Sen, 2005], and the recent theory of universal inference [Wasserman et al., 2020].

## 2.5 GAUSSIAN LIKELIHOOD

In our construction of test statistics, we will assume the data to consist of random vectors  $X^{(1)}, \dots, X^{(n)}$  drawn independently from a (centered) multivariate normal distribution  $\mathcal{N}(0, \Sigma)$  with density  $p(x|\Sigma)$ . Let  $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n X^{(i)} X^{(i)T}$ , and let  $\mathcal{M}_0 \subset \mathcal{M}$  be a subset of covariance matrices. In order to test

$$H_0 : \Sigma \in \mathcal{M}_0 \quad \text{against} \quad H_1 : \Sigma \in \mathcal{M} \setminus \mathcal{M}_0,$$

we use the likelihood-ratio statistic

$$\lambda_n = 2 \left( \sup_{\Sigma \in \mathcal{M}} \ell_n(\Sigma) - \sup_{\Sigma \in \mathcal{M}_0} \ell_n(\Sigma) \right)$$

based on the log-likelihood  $\ell_n(\Sigma) = \sum_{i=1}^n \log p(X_{(i)}|\Sigma)$  with

$$\frac{2}{n} \ell_n(\Sigma) = -\log \det(2\pi\Sigma) - \text{tr}(\Sigma^{-1}\hat{\Sigma}).$$

### 3 BIVARIATE CASE

This section develops the details of our approach in the two-dimensional setting, with two variables  $X_1$  and  $X_2$ .

#### 3.1 REPRESENTATION OF CAUSAL EFFECT

For  $d = 2$ , the model uncertainty boils down to uncertainty about the direction of the single edge,  $1 \rightarrow 2$  versus  $1 \leftarrow 2$ , and we maintain the two possible LSEMs

$$\begin{aligned} \text{(M1)} \quad X_1 &= \epsilon_1, & X_2 &= \beta_{21}X_1 + \epsilon_2, \\ \text{(M2)} \quad X_1 &= \beta_{12}X_2 + \epsilon_1, & X_2 &= \epsilon_2. \end{aligned}$$

The mere assumption of a homoscedastic LSEM imposes structure on the covariance matrix  $\Sigma$  of  $X = (X_1, X_2)^T$ . Under model (M1), we obtain

$$\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{pmatrix} = \sigma^2 \begin{pmatrix} 1 & \beta_{21} \\ \beta_{21} & \beta_{21}^2 + 1 \end{pmatrix},$$

which leads to the relations

$$\beta_{21} = \frac{\Sigma_{12}}{\Sigma_{11}}, \quad \beta_{21}^2 + 1 = \frac{\Sigma_{22}}{\Sigma_{11}}, \quad \Sigma_{11}^2 = \det(\Sigma), \quad (4)$$

with  $\det(\Sigma) = \Sigma_{11}\Sigma_{22} - \Sigma_{12}^2$ . Analogously, for model (M2),

$$\beta_{12} = \frac{\Sigma_{12}}{\Sigma_{22}}, \quad \beta_{12}^2 + 1 = \frac{\Sigma_{11}}{\Sigma_{22}}, \quad \Sigma_{22}^2 = \det(\Sigma). \quad (5)$$

Hence, the set of  $2 \times 2$  covariance matrices  $\Sigma$  that are possible under homoscedastic LSEMs is  $\mathcal{M}_r = \mathcal{M}_{r1} \cup \mathcal{M}_{r2}$ , where

$$\mathcal{M}_{ra} := \{\Sigma \in \mathcal{M} : \Sigma_{aa}^2 = \det(\Sigma)\}, \quad a = 1, 2. \quad (6)$$

By symmetry, when considering total causal effects we may focus on the effect of  $X_1$  on  $X_2$ , which is  $\beta_{21}$  under model (M1) and zero under model (M2). So, by (4),

$$\mathcal{C}(1 \rightarrow 2) = \frac{\Sigma_{12}}{\Sigma_{11}} \mathbb{1}\{\Sigma \in \mathcal{M}_{r1}\}. \quad (7)$$

#### 3.2 CONSTRAINED LIKELIHOOD-RATIO TESTS

Our construction of confidence sets inverts tests of hypotheses that specify  $\mathcal{C}(1 \rightarrow 2) = \psi$ . We now present two approaches to perform likelihood ratio tests. Likelihood ratio statistics are easily defined but their probability distributions are generally difficult to determine at singularities, as encountered here where the alternative  $\mathcal{M}_r$  is a union, and thus non-smooth. To simplify distribution theory, we will relax the alternative to be the entire positive definite cone  $\mathcal{M}$ .

#### 3.2.1 Testing Inequality Constraints

Our first approach exploits that model selection for homoscedastic LSEMs can be achieved by ordering variances [Chen et al., 2019]. Indeed, if  $\Sigma \in \mathcal{M}_r$ , then  $\Sigma \in \mathcal{M}_{r1}$  precisely when  $\Sigma_{11} \leq \Sigma_{22}$ . We will use this fact to set up hypotheses that encode  $\mathcal{C}(1 \rightarrow 2) = \psi$  for given  $\psi$ . Three cases arise:  $\psi = 0$ ,  $0 < |\psi| < 1$  and  $|\psi| \geq 1$ .

*Case  $\psi = 0$ .* Given  $\Sigma \in \mathcal{M}_r$ , we have  $\mathcal{C}(1 \rightarrow 2) = 0$  if and only if  $\Sigma_{11} \geq \Sigma_{22}$ . Hence, we conduct the test of

$$H_0 : \Sigma_{11} \geq \Sigma_{22} \quad \text{against} \quad H_1 : \Sigma \in \mathcal{M}.$$

In this scenario, the asymptotic null distribution of  $\lambda_n$  depends on the unknown value of  $\Sigma$ , but it is easy to see that the stochastically largest asymptotic distribution arises when  $\Sigma_{11} = \Sigma_{22}$ , in which case

$$\lambda_n \xrightarrow{\mathcal{D}} 0.5\chi_0^2 + 0.5\chi_1^2, \quad \text{as } n \rightarrow \infty,$$

where  $\chi_d^2$  denotes a chi-square distribution with  $d$  degrees of freedom and  $\chi_0^2 \equiv 0$ ; see Silvapulle and Sen [2005].

*Case  $0 < |\psi| < 1$ .* When  $\psi \neq 0$ , it must be that  $X_2$  is the causally dependent variable which corresponds to  $\Sigma_{11} \leq \Sigma_{22}$  and  $\Sigma_{12}/\Sigma_{11} = \psi$ , according to (7). Hence, we test

$$H_0 : \Sigma_{12} = \psi\Sigma_{11}, \quad \Sigma_{11} \leq \Sigma_{22} \quad \text{against} \quad H_1 : \Sigma \in \mathcal{M}.$$

For the least favorable covariance matrix in  $H_0$ , it holds that

$$\lambda_n \xrightarrow{\mathcal{D}} 0.5\chi_1^2 + 0.5\chi_2^2, \quad \text{as } n \rightarrow \infty.$$

*Case  $|\psi| \geq 1$ .* Again the two constraints  $\Sigma_{11} \leq \Sigma_{22}$  and  $\Sigma_{12} = \psi\Sigma_{11}$  have to be satisfied. However, the Cauchy-Schwarz inequality yields  $\Sigma_{22} \geq \Sigma_{12}^2/\Sigma_{11} = \psi^2\Sigma_{11} \geq \Sigma_{11}$ , since  $|\psi| \geq 1$ . Consequently, the inequality condition is automatically fulfilled and it suffices to test

$$H_0 : \Sigma_{12} = \psi\Sigma_{11} \quad \text{against} \quad H_1 : \Sigma \in \mathcal{M}.$$

The likelihood ratio statistic satisfies

$$\lambda_n \xrightarrow{\mathcal{D}} \chi_1^2, \quad \text{as } n \rightarrow \infty.$$

We remark that in the first and second case one may also follow a two-step procedure that uses a suitably calibrated pretest to decide which asymptotic distribution to employ [Silvapulle, 1996]. We report no details on this approach here as we found the power gains to be only very slight.

Based on the above suite of tests, we may form a confidence interval from the accepted values of  $\psi$ , which we determine in practice by inspecting a fine grid of choices. In our simulations in Section 3.4 we refer to this method as `LRT1`. We will also consider a heuristic variant in which we compute the likelihood ratio statistics by restricting the null and the alternative to the union of the two LSEMs, i.e., to  $\mathcal{M}_r$ , but still set critical values based on the asymptotic distributions given above. We refer to this method as `LRT1b`.

### 3.2.2 Testing Polynomial Constraints

The previous method encodes membership in model (M1) via the key inequality  $\Sigma_{11} \leq \Sigma_{22}$ . As an alternative we may directly work with the set of covariance matrices  $\mathcal{M}_{r,1}$  given in (6) when specifying null hypotheses. However, to retain simple distributional approximations we continue to relax the alternative to be the entire p.d.-cone  $\mathcal{M}$ .

Recall that the causal effect  $\mathcal{C}(1 \rightarrow 2)$  is non-zero only under model (M1), i.e., if  $\Sigma \in \mathcal{M}_{r,1}$ . In this case the effect is  $\mathcal{C}(1 \rightarrow 2) = \Sigma_{12}/\Sigma_{11}$ . All matrices with  $\mathcal{C}(1 \rightarrow 2) = 0$  belong to  $\mathcal{M}_{r,2}$ . Thus, we test the null hypotheses

$$H_0 : \begin{cases} \Sigma_{12} = \psi \Sigma_{11} \text{ and } \Sigma \in \mathcal{M}_{r,1}, & \text{if } \psi \neq 0, \\ \Sigma \in \mathcal{M}_{r,2}, & \text{if } \psi = 0. \end{cases}$$

We write  $\Theta_0^{(\psi)}$  for the respective sets of covariance matrices.

*Case  $\psi \neq 0$ .* If  $\psi \neq 0$ , the set  $\Theta_0^{(\psi)}$  is a one-dimensional submanifold of the three-dimensional p.d.-cone  $\mathcal{M}$ . The likelihood ratio statistic in this case is found to be

$$\lambda_n^{(\psi)} = 2n \log \left\{ \frac{1}{2 \det(\hat{\Sigma})^{1/2}} \text{tr} \left[ \begin{pmatrix} 1 + \psi^2 & -\psi \\ -\psi & 1 \end{pmatrix} \hat{\Sigma} \right] \right\},$$

and for  $\Sigma \in \Theta_0^{(\psi)}$  we have

$$\lambda_n^{(\psi)} \xrightarrow{\mathcal{D}} \chi_2^2, \quad \text{as } n \rightarrow \infty.$$

*Case  $\psi = 0$ .* The set  $\Theta_0^{(0)}$  is a two-dimensional submanifold of  $\mathcal{M}$  and yields the likelihood ratio statistic

$$\lambda_n^{(0)} = 2n \log \left\{ \frac{1}{2 \det(\hat{\Sigma})^{1/2}} \left( \hat{\Sigma}_{11} - \frac{\hat{\Sigma}_{12}^2}{\hat{\Sigma}_{22}} + \hat{\Sigma}_{22} \right) \right\}.$$

Here, for  $\Sigma \in \Theta_0^{(0)}$ ,

$$\lambda_n^{(0)} \xrightarrow{\mathcal{D}} \chi_1^2, \quad \text{as } n \rightarrow \infty.$$

Given the explicit form of the likelihood ratio statistics we can explicitly determine an asymptotic confidence set for the total causal effect.

**Theorem 1.** *Let  $\alpha \in (0, 1)$ . Then an asymptotic  $(1 - \alpha)$  confidence set for the causal effect  $\mathcal{C}(1 \rightarrow 2)$  is given by*

$$C = \{\psi \neq 0 : \lambda_n^{(\psi)} \leq \chi_{2,1-\alpha}^2\} \cup \{0 : \lambda_n^{(0)} \leq \chi_{1,1-\alpha}^2\}.$$

Furthermore, if we define

$$K := 2 \hat{\Sigma}_{11} \det(\hat{\Sigma})^{1/2} \exp\left(\frac{1}{2n} \chi_{2,1-\alpha}^2\right) - \hat{\Sigma}_{11}^2 - \det(\hat{\Sigma})$$

and  $K \geq 0$ , then

$$\{\psi \neq 0 : \lambda_n^{(\psi)} \leq \chi_{2,1-\alpha}^2\} = \left[ \frac{\hat{\Sigma}_{12} - \sqrt{K}}{\hat{\Sigma}_{11}}, \frac{\hat{\Sigma}_{12} + \sqrt{K}}{\hat{\Sigma}_{11}} \right].$$

*Remark.* Since the dimension of  $\Theta_0^{(0)}$  exceeds that of  $\Theta_0^{(\psi)}$  for  $\psi \neq 0$ , we are led to consider two different degrees of freedom for chi-square limits. As a result there exist data for which we reject a zero effect but accept positive and negative effects that are arbitrarily small in magnitude. However, this case arises very rarely. If it does, it may be preferable to simply include zero in the confidence interval.

*Proof.* If  $\psi \neq 0$ , it is easy to see that  $\lambda_n^{(\psi)} \leq \chi_{2,1-\alpha}^2$  if and only if

$$\begin{aligned} & \psi^2 \hat{\Sigma}_{11} - 2\psi \hat{\Sigma}_{12} + \hat{\Sigma}_{11} + \hat{\Sigma}_{22} \\ & - 2 \det(\hat{\Sigma})^{1/2} \exp\left(\frac{1}{2n} \chi_{2,1-\alpha}^2\right) \leq 0. \end{aligned}$$

The inequality features a strictly convex quadratic polynomial in  $\psi$ . The confidence interval is nonempty if the quadratic has real roots, which occurs for  $K \geq 0$ . The roots are  $(\hat{\Sigma}_{12} \pm \sqrt{K})/\hat{\Sigma}_{11}$  and give the claimed explicit lower and upper end of the confidence interval. The confidence set is then completed by checking whether we accept (and thus include)  $\psi = 0$ .  $\square$

In our simulations, we refer to this method as LRT2.

### 3.3 SPLIT LIKELIHOOD RATIO TESTS

Wasserman et al. [2020] introduced the framework of universal inference, a general method for constructing hypothesis tests and confidence regions that are conservative but valid in finite samples. Universal inference employs a modification of the classical likelihood-ratio statistic termed the *split likelihood ratio* (SLR), which is especially appealing for irregular composite hypotheses where asymptotic distributions are intractable. As its name indicates, the SLR statistic is based on a data splitting approach. Type-I error control is guaranteed by an application of Markov's inequality.

The method proceeds by splitting the data into two subsets  $D_0 = \{X^{(1)}, \dots, X^{(k)}\}$  and  $D_1 = \{X^{(k+1)}, \dots, X^{(n)}\}$ . Let  $\ell_0(\Sigma)$  denote the log-likelihood function based on  $D_0$ , that is,  $\ell_0(\Sigma) = \sum_{i=1}^k \log p(X^{(i)} | \Sigma)$ . We then calculate the *profile log-likelihood function*

$$\ell_{\dagger}(\psi) = \max\{\ell_0(\Sigma) : \Sigma \in \mathcal{M}_r, \mathcal{C}(1 \rightarrow 2) = \psi\},$$

and choose any estimator  $\tilde{\Sigma}^1$  based on  $D_1$ . Then

$$C = \{\psi : \ell_0(\tilde{\Sigma}^1) - \ell_{\dagger}(\psi) \leq \log(1/\alpha)\} \quad (8)$$

is a (conservative) confidence set for the total causal effect  $\mathcal{C}(1 \rightarrow 2)$  with confidence level  $1 - \alpha$ .

Let  $\hat{\Sigma}^0 = \frac{1}{k} \sum_{i=1}^k X^{(i)} X^{(i)T}$  be the empirical covariance matrix for  $D_0$ . In each of our two possible LSEMs, (M1) and (M2), maximizing  $\ell_0$  with respect to the variance parameter

$\sigma^2$  is straightforward; recall (3). We find that for fixed  $B$ , the maximum of  $\ell_0$  over  $\sigma^2 > 0$  is

$$-k \log \left\{ \pi \operatorname{tr} \left[ (I_d - B)^T (I_d - B) \hat{\Sigma}^0 \right] \right\} - k. \quad (9)$$

Now we assume that the causal effect  $\mathcal{C}(1 \rightarrow 2)$  equals a fixed value  $\psi \in \mathbb{R}$  and we maximize  $\ell_0$  further over any remaining parameters.

*Case  $\psi \neq 0$ .* The covariance matrix  $\Sigma$  is in  $\mathcal{M}_{r_1}$  and in its parametrization  $\beta_{21} = \psi$ . From (9), we find that for  $\psi \neq 0$ ,

$$\ell_{\dagger}(\psi) = -k \log \left\{ \pi \operatorname{tr} \left[ \begin{pmatrix} 1 + \psi^2 & -\psi \\ -\psi & 1 \end{pmatrix} \hat{\Sigma}^0 \right] \right\} - k. \quad (10)$$

*Case  $\psi = 0$ .* Now,  $\Sigma$  may be any matrix in  $\mathcal{M}_{r_2}$ . Straightforward calculations show that

$$\max_{\beta_{12}} \operatorname{tr} \left[ \begin{pmatrix} 1 & -\beta_{12} \\ -\beta_{12} & 1 + \beta_{12}^2 \end{pmatrix} \hat{\Sigma}^0 \right] = \hat{\Sigma}_{11}^0 - \frac{(\hat{\Sigma}_{12}^0)^2}{\hat{\Sigma}_{22}^0} + \hat{\Sigma}_{22}^0.$$

Inserting this expression in (9) yields the profile  $\ell_{\dagger}(0)$ .

With these preparations, we can now explicitly calculate the boundaries of the confidence set for the total causal effect given in (8).

**Theorem 2.** *Let  $\alpha \in (0, 1)$  and define*

$$G_a := 2\hat{\Sigma}_{aa}^0 \alpha^{-1/k} \det(\tilde{\Sigma}^1)^{1/2} \exp \left( \frac{1}{2} \operatorname{tr} [(\tilde{\Sigma}^1)^{-1} \hat{\Sigma}^0] - 1 \right) - (\hat{\Sigma}_{aa}^0)^2 - \det(\hat{\Sigma}^0), \quad a = 1, 2.$$

(i) *If  $G_1 \geq 0$ , then the nonzero elements of the confidence set  $C$  from (8) are the nonzero elements of the interval  $[L, U]$  with*

$$L := \frac{\hat{\Sigma}_{12}^0 - \sqrt{G_1}}{\hat{\Sigma}_{11}^0}, \quad U := \frac{\hat{\Sigma}_{12}^0 + \sqrt{G_1}}{\hat{\Sigma}_{11}^0}.$$

(ii) *The set  $C$  from (8) contains zero if and only if  $G_2 \geq 0$ .*

*Proof.* Expanding the formula from (10), we obtain that a value  $\psi \neq 0$  satisfies the inequality for the confidence set in (8) if and only if

$$\begin{aligned} & \psi^2 \hat{\Sigma}_{11}^0 - 2\psi \hat{\Sigma}_{12}^0 + \hat{\Sigma}_{11}^0 + \hat{\Sigma}_{22}^0 \\ & - 2\alpha^{-1/k} (\det \tilde{\Sigma}^1)^{1/2} \exp \left( \frac{1}{2} \operatorname{tr} [(\tilde{\Sigma}^1)^{-1} \hat{\Sigma}^0] - 1 \right) \leq 0. \end{aligned}$$

If  $G_1 \geq 0$ , the involved convex quadratic function has real roots at the claimed values of  $L$  and  $U$ . If  $G_1 < 0$ , the inequality has no solutions and the nonzero part of the confidence set remains empty.

The inequality that is equivalent to inclusion of  $\psi = 0$  in the confidence set is similar.  $\square$

So far we have not made a specific choice for  $\tilde{\Sigma}^1$ . In (8), we observe that for small confidence sets it is desirable to form estimates  $\tilde{\Sigma}^1$  that achieve large values of  $\ell_0(\tilde{\Sigma}^1)$ . Since the universal inference approach poses no problems due to irregular geometry of hypothesis/alternative, it is natural to form an estimate that exploits the assumed validity of at least one of the two LSEMs. In other words, we choose  $\tilde{\Sigma}^1$  to be the maximum likelihood estimator under the restriction that  $\Sigma \in \mathcal{M}_r$ . We refer to this method as `SLRT`.

For problems with more than two variables determining the profile log-likelihood is computationally involved in the sense that higher degree polynomial equations need to be solved. As an alternative we also report results for a heuristic in which we take  $\tilde{\Sigma}^1$  to be the unrestricted sample covariance matrix and replace the profile log-likelihood  $\ell_{\dagger}(\psi)$  by  $\ell_0$  evaluated at a consistent moment estimator. To form this estimator under the constraint of an assumed causal effect  $\mathcal{C}(i \rightarrow j) = \psi$ , we consider every causal ordering that permits the effect, use the sample variance to estimate the variance of the first variable in the ordering, use sample covariances for all covariances of pairs of variables other than  $(i, j)$  and fill the rest of the matrix so that the constraint  $\mathcal{C}(i \rightarrow j) = \psi$  holds and all error variances are equal. We then choose the ordering that leads to the maximum value of  $\ell_0$ . We refer to this method as `estSLRT`.

### 3.4 SIMULATIONS

In this section we present the results of a simulation study with the aim to compare the empirical coverage probabilities and widths of the different proposed confidence intervals for the causal effect. Our simulation experiment was designed as follows. We generated pseudo random numbers according to model (M1) ( $1 \rightarrow 2$ ) and model (M2) ( $1 \leftarrow 2$ ), respectively, with standard normal errors. For a selection of different values of  $\beta \equiv \beta_{12}$  ( $\beta \equiv \beta_{21}$ ) and different sample sizes  $n$ , we simulated 10000 independent data sets, for which we then determined the confidence sets for  $\alpha = 0.05$ .

The resulting empirical coverage probabilities are reported in Table 1, where all five proposed methods achieve the desired coverage frequency of 0.95. Furthermore, all proposed methods seem to be particularly conservative if the causal effect is zero. In general the split likelihood ratio methods are, as expected, the most conservative. For the purpose of comparison we included the empirical coverage probabilities of confidence intervals for the causal effect calculated with two different bootstrapping methods, as explained in Section 2.3. For each resampled data set `Bootstrap1` simply uses the sample covariance to select the direction (via (M1) if  $\hat{\Sigma}_{11} \leq \hat{\Sigma}_{22}$ , (M2) if  $\hat{\Sigma}_{22} < \hat{\Sigma}_{11}$ ) and subsequently calculates the causal effect based on the selected model. `Bootstrap2` employs an established causal discovery algorithm for the equal variance case `GDS`, proposed by Peters and Bühlmann [2014], to estimate the causal effect

Table 1: Empirical Coverage of 95%-Confidence Intervals for the Total Causal Effect of  $X_1$  on  $X_2$  (10000 Replications).

method	$n \setminus \beta$	$X_1 \rightarrow X_2$					$X_2 \rightarrow X_1$				
		0	0.05	0.1	0.2	0.5	0	0.05	0.1	0.2	0.5
LRT1	100	1.00	0.95	0.95	0.96	0.98	1.00	1.00	0.99	0.98	1.00
	500	1.00	0.95	0.95	0.96	0.97	1.00	0.99	0.98	0.98	1.00
	1000	1.00	0.96	0.96	0.97	0.98	1.00	0.98	0.97	0.99	1.00
LRT1b	100	1.00	0.97	0.97	0.96	0.95	1.00	1.00	0.99	0.96	0.96
	500	1.00	0.97	0.97	0.96	0.98	1.00	1.00	0.99	0.94	1.00
	1000	1.00	0.97	0.97	0.95	0.97	1.00	1.00	0.99	0.94	1.00
LRT2	100	0.97	0.96	0.96	0.96	0.96	0.96	0.97	0.97	0.97	0.97
	500	0.97	0.97	0.96	0.97	0.97	0.97	0.97	0.97	0.97	1.00
	1000	0.97	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.97	1.00
SLRT	100	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	500	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	1000	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
estSLRT	100	0.99	0.98	0.98	0.98	0.98	0.99	0.99	0.99	0.99	0.99
	500	0.99	0.98	0.98	0.98	0.98	0.99	0.99	0.99	0.99	0.99
	1000	0.99	0.98	0.98	0.98	0.98	0.99	0.99	0.99	0.99	0.99
Bootstrap1	100	0.99	0.89	0.89	0.91	0.94	1.00	1.00	0.99	0.99	1.00
	500	0.99	0.89	0.90	0.93	0.95	1.00	0.99	0.99	0.99	1.00
	1000	0.99	0.89	0.90	0.94	0.94	1.00	0.99	0.98	0.99	1.00
Bootstrap2	100	1.00	0.56	0.74	0.93	0.95	1.00	1.00	1.00	1.00	1.00
	500	1.00	0.63	0.84	0.94	0.95	1.00	1.00	1.00	0.99	1.00
	1000	1.00	0.67	0.92	0.95	0.96	1.00	1.00	1.00	0.99	1.00

for each resampled data set. In principle the GDS method is a greedy search algorithm that maximizes the likelihood. As expected, and theoretically explained in Section 2.3, bootstrapping methods do not work in practice and do not achieve the the required coverage frequency.

Figure 2 displays the mean width of the smallest interval containing the constructed confidence set. The widths are plotted against the sample size  $n$  for a true causal effect of size 0.5 (in different directions). We note that while the confidence sets predominantly are intervals, it is possible that they are "torn" with  $\{0\}$  as a disconnected component, reflecting the larger null hypothesis that is associated to a zero effect. The more conservative split likelihood ratio methods yield wider confidence intervals. The `estSLRT` heuristic outperforms the standard `SLRT`. As it does not fully optimize the profile log-likelihood function, `estSLRT` produces smaller sets, yet the desired empirical coverage is (easily) achieved. In the case of no causal effect the confidence intervals converge to zero for all proposed methods.

Figure 3 shows the percentage of times zero is in the calculated confidence sets. The percentages are plotted for a total causal effect  $\mathcal{C}(1 \rightarrow 2) = 0.5$  against the sample size and for sample size  $n = 500$  against the size of the causal effect. All proposed methods are consistent and exclude the possibility of no causal effect with increasing sample size. Therefore, all proposed methods not only yield correct confidence sets for the total causal effect but also successfully help decide whether the effect is nonzero or not. Figure 3 also shows that we exclude the possibility of no causal effect more frequently the higher the actual causal effect is.

Even though it seems that in Figure 2 and 3 the `LRT1b` method seems to perform best, we should stress that we




Figure 2: Average maximum width of 95%-confidence intervals for the causal effect of  $X_1$  on  $X_2$  (10000 replications).




Figure 3: Percentage of times zero contained in the 95%-confidence intervals for the total causal effect of  $X_1$  on  $X_2$  (10000 replications). (Left)  $\mathcal{C}(1 \rightarrow 2) = 0.5$  against sample size. (Right)  $n = 500$  against size of the total causal effect.

only have theoretical guarantees for the methods `LRT1`, `LRT2` and `SLRT`. Out of those three methods the `LRT2` method seems to perform best, but this method also has to be handled with care as we will see in the following real data example.

### 3.5 DATA EXAMPLES

For a real world data benchmark we used the cause effect pairs data set presented in Mooij et al. [2016]. It consists of different cause effect data pairs from various fields, for which the true causal direction is determined by domain knowledge. For the application of homoscedastic LSEMs, we selected the following pairs: `pair66` and `pair67` containing daily stock returns, `pair76` containing the average annual rate of change of population and total dietary consumption, and `pair89` and `pair90` which describe the degree of root decomposition in forests and grasslands, respectively. The first three data sets exhibit a causal effect

from  $X_1$  to  $X_2$  while the last two feature a causal effect from  $X_2$  to  $X_1$ . Before calculating the confidence intervals for the causal effect of  $X_1$  on  $X_2$ , we centered the data.

The results in Figure 4 show that method LRT2 which performed very well in the simulations produced an empty confidence set for all five data pairs. In brief, the method is not able to cope with model misspecification (where we may still wish to obtain a confidence interval for a well-defined parameter).<sup>1</sup> The method estSLRT similarly suffers from this problem. Indeed, both LRT2 and estSLRT contrasts a linear and homoscedastic null hypothesis against a general Gaussian alternative. Even for normal data, departures from homoscedasticity may favor the alternative in all testing problems that correspond to the confidence set, which then remains empty.

Although it was not the most statistically efficient method in our simulations, the method LRT1 performs best for the real world data. As the union of all its tested null hypotheses coincides with the alternative (the p.d. cone  $\mathcal{M}$ ), it always produces a nonempty confidence set. Under misspecification the interval targets the parameter

$$\tilde{\mathcal{C}}(1 \rightarrow 2) = \frac{\Sigma_{12}}{\Sigma_{11}} \mathbf{1}\{\Sigma_{11} \leq \Sigma_{22}\};$$

an extension of the parameter defined in (7). In this sense the method is less sensitive to departures from homoscedasticity, or even linearity if the true covariance matrix  $\Sigma$  is defined as furnishing the KL-best normal approximation to the data-generating distribution. The interval width for LRT1 is here similar or even smaller than for the heuristic LRT1b.

The standard split likelihood ratio method SLRT also performs well, but produces considerably wider intervals than LRT. The width of the estimated confidence intervals with the split likelihood ratio methods slightly vary depending on how the real data set is (randomly) split.

Finally, we note that all proposed methods recognize that there is no causal effect from  $X_1$  to  $X_2$  in the last two data pairs. For the first data pair, some uncertainty remains at the 95%-level about whether an effect is indeed present (LRT1 and SLRT). For comparison, we also calculated the confidence intervals for the causal effect of  $X_2$  on  $X_1$ , see Figure 5, and the results lead to the same conclusions.

## 4 HIGHER DIMENSIONS

In this section we give an outlook on how one may extend the proposed methods to higher dimensional cases. The asymptotic distribution of the first method LRT1 via testing inequalities for the conditional variances is a mixture of chi-square distributions. The calculation of those

<sup>1</sup>E.g., one could consider the causal effect associated with the matrix  $\Sigma \in \mathcal{M}_r$  such that the normal distribution  $N(0, \Sigma)$  is closest in KL divergence to the data-generating distribution.




Figure 4: 95%-Confidence intervals for the total causal effect of  $X_1$  on  $X_2$  for different real world data pairs.

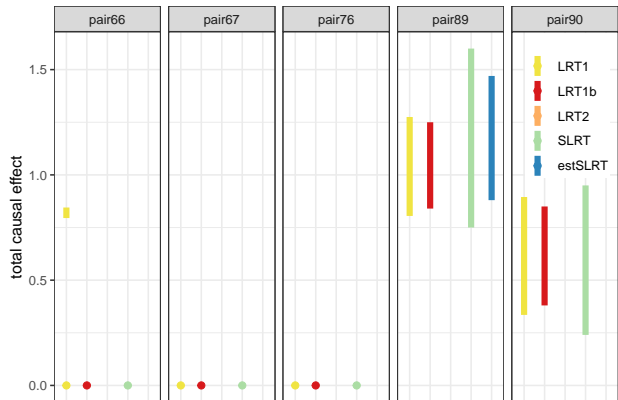


Figure 5: 95%-Confidence intervals for the total causal effect of  $X_2$  on  $X_1$  for different real world data pairs.






Figure 6: Empirical coverage and average maximum width of the 95%-confidence interval for the total causal effect of  $X_1$  on  $X_2$  (10000 replications).

mixture weights is difficult in higher dimensions. Using data-dependent critical values as used in Al Mohamad et al. [2020] may be an option to push the methodology to moderately small dimension, but at this point we have not yet explored this option.

Calculating asymptotic distributions for the second method LRT2 remains feasible in any individual LSEM that allows for a given effect  $\mathcal{C}(i \rightarrow j)$  to be nonzero. However, one has to then address the issue that several LSEM allow for  $\mathcal{C}(i \rightarrow j) \neq 0$  and the relevant null hypothesis becomes a union of smooth manifolds. One way to address this problem would be to form an intersection union test, but this is again a topic for future work.

The approach that is the simplest for extension to higher dimensional cases are the split likelihood ratio tests. We illustrate this in the following simulations for the case of  $d = 3$  variables. The three dimensional case allows for six possible models based on the ordering of the three involved variables  $X_1, X_2$  and  $X_3$ . Figure 6 displays the average maximum width and empirical coverage of the estimated confidence intervals for the total causal effect of  $X_1$  on  $X_2$  in the six different models. In the upper three cases there is a true causal effect of size 0.5, while in the other cases there is no causal effect of  $X_1$  on  $X_2$ . The constructed confidence intervals have a high coverage, exceeding the desired coverage with the method SLRT in all cases and with the heuristic method estSLRT in five of the six possible cases.

## 5 CONCLUSION

We proposed new methods to construct confidence intervals for the total causal effect in problems in which causal




Figure 7: Empirical coverage and average maximum width of the 95%-confidence interval for the total causal effect of  $X_1$  on  $X_2$  in a LiNGAM setting (100 replications).

structure is unknown but identifiable. We cope with this uncertainty in a test inversion approach that accounts for both types of uncertainty: causal structure and numerical size of nonzero effects. For two-variable problems the empirical results for the LRT1 method that tests inequalities among variances are very promising, but it may prove difficult to extend this method to higher dimensional cases and settings other than homoscedastic LSEMs. The second proposed LRT2 method tests the polynomial constraints imposed by the LSEM assumption and can be extended to higher dimensional cases. However, it can be sensitive to departures from the modeling assumptions, as seen in our real data example. This is due to the fact that the precise model assumptions were incorporated in the null hypotheses, but for simplicity in distribution theory not in the alternative. An interesting problem for further research would be to improve our understanding of possible asymptotic approximations for the LRT2 statistics when the alternative is not relaxed but kept as the union of all homoscedastic LSEMs.

The last proposed SLRT methods based on the theory of universal inference are the most conservative but also easiest to apply methods. They can be extended rather directly not only to higher dimensional cases but also to other modeling frameworks. To illustrate the latter point, we briefly consider the usage of the split likelihood ratio methods for LSEMs with non-Gaussian errors (LiNGAM). Shimizu et al. [2006] showed that under these assumptions unique identification is possible and the causal structure imposes constraints on the (conditional) moments, see Wang and Drton [2020]. We can thus use empirical likelihood methods [Wang and Drton, 2017] to form a split likelihood ratio and construct confidence intervals for the causal effect. Figure 7 shows the empirical coverage and average maximum width of these confidence intervals for the causal effect of  $X_1$  on  $X_2$  in the bivariate LiNGAM setting. We simulated the data with a causal effect of size 0.5 (in different directions) and uniform  $(-1, 1)$  distributed error terms. We observe that the method is conservative yet informative.

## Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 883818).

## References

- Diaa Al Mohamad, Erik W. van Zwet, Eric Cator, and Jelle J. Goeman. Adaptive critical value for constrained likelihood ratio testing. *Biometrika*, 107(3):677–688, 2020.
- Donald W. K. Andrews and Patrik Guggenberger. Asymptotic size and a problem with subsampling and with the  $m$  out of  $n$  bootstrap. *Econometric Theory*, 26(2):426–468, 2010.
- Xuan Cao, Kshitij Khare, and Malay Ghosh. Posterior graph selection and estimation consistency for high-dimensional Bayesian DAG models. *Ann. Statist.*, 47(1):319–348, 2019.
- Wenyu Chen, Mathias Drton, and Y. Samuel Wang. On causal discovery with an equal-variance assumption. *Biometrika*, 106(4):973–980, 2019.
- Tom Claassen and Tom Heskes. A Bayesian approach to constraint based causal inference. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, UAI’12, page 207–216, Arlington, Virginia, USA, 2012. AUAI Press.
- Mathias Drton and Benjamin Williams. Quantifying the failure of bootstrap likelihood ratio tests. *Biometrika*, 98(4):919–934, 2011.
- Asish Ghoshal and Jean Honorio. Learning linear structural equation models in polynomial time and sample complexity. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84, pages 1466–1475. PMLR, 09–11 Apr 2018.
- Patrik O. Hoyer and Antti Hyttinen. Bayesian discovery of linear acyclic causal models. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI ’09, page 240–248, Arlington, Virginia, USA, 2009. AUAI Press.
- Marloes Maathuis, Mathias Drton, Steffen Lauritzen, and Martin Wainwright, editors. *Handbook of graphical models*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, 2019.
- Joris M. Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *J. Mach. Learn. Res.*, 17:Paper No. 32, 102, 2016.
- Judea Pearl. *Causality*. Cambridge University Press, Cambridge, second edition, 2009. Models, reasoning, and inference.
- Jonas Peters and Peter Bühlmann. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2017. Foundations and learning algorithms.
- Dimitris N. Politis, Joseph P. Romano, and Michael Wolf. *Subsampling*. Springer Series in Statistics. Springer-Verlag, New York, 1999.
- Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-Gaussian acyclic model for causal discovery. *J. Mach. Learn. Res.*, 7:2003–2030, 2006.
- Mervyn J. Silvapulle. A test in the presence of nuisance parameters. *J. Amer. Statist. Assoc.*, 91(436):1690–1693, 1996.
- Mervyn J. Silvapulle and Pranab K. Sen. *Constrained statistical inference*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, 2005.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2000.
- Y. Samuel Wang and Mathias Drton. Empirical likelihood for linear structural equation models with dependent errors. *Stat*, 6:434–447, 2017.
- Y. Samuel Wang and Mathias Drton. High-dimensional causal discovery under non-Gaussianity. *Biometrika*, 107(1):41–59, 2020.
- Larry Wasserman, Aaditya Ramdas, and Sivaraman Balakrishnan. Universal inference. *Proceedings of the National Academy of Sciences*, 117(29):16880–16890, 2020.