

Identifying Causal Effects using Instrumental Time Series: Nuisance IV and Correcting for the Past

Nikolaj Thams

THAMS@MATH.KU.DK

Rikke Søndergaard

RSN@MATH.KU.DK

Sebastian Weichwald

SWEICHWALD@MATH.KU.DK

Jonas Peters

JONAS.PETERS@MATH.KU.DK

*Department of Mathematical Sciences
University of Copenhagen
Denmark*

Editor: Vanessa Didelez

Abstract

Instrumental variable (IV) regression relies on instruments to infer causal effects from observational data with unobserved confounding. We consider IV regression in time series models, such as vector auto-regressive (VAR) processes. Direct applications of i.i.d. techniques are generally inconsistent as they do not correctly adjust for dependencies in the past. In this paper, we outline the difficulties that arise due to time structure and propose methodology for constructing identifying equations that can be used for consistent parametric estimation of causal effects in time series data. One method uses extra nuisance covariates to obtain identifiability (an idea that can be of interest even in the i.i.d. case). We further propose a graph marginalization framework that allows us to apply nuisance IV and other IV methods in a principled way to time series. Our methods make use of a version of the global Markov property, which we prove holds for VAR(p) processes. For VAR(1) processes, we prove identifiability conditions that relate to Jordan forms and are different from the well-known rank conditions in the i.i.d. case (they do not require as many instruments as covariates, for example). We provide methods, prove their consistency, and show how the inferred causal effect can be used for distribution generalization. Simulation experiments corroborate our theoretical results. We provide ready-to-use Python code.

Keywords: causality, time series, instrumental variables, VAR processes, Markov property, distribution generalization

1. Introduction

Predicting a response variable Y from observations of covariates X may be insufficient to answer a scientific question at hand. Instead, we may wish to model how the response variable Y reacts to an intervention on X . Such modeling requires causal knowledge. For example, for i.i.d. data from a linear model $Y := \beta X + g(H, \varepsilon^Y)$, it is well-known that an ordinary least squares (OLS) regression of Y on X generally yields a biased estimator of the linear causal effect β from X on Y when an unobserved variable H confounds X and Y . Instead, we may obtain unbiased estimates of β by utilising instrumental variables (IVs) I that correlate with the covariates X , are independent of H , and affect Y only indirectly through X . IV regression, pioneered by Wright (1928) and Reiersøl (1945), is well-established in econometrics (Angrist et al., 1996; Staiger and Stock, 1997; Angrist and Krueger, 2001), statistics (Bowden and Turkington, 1985) and epidemiology (Hernán and Robins, 2006; Didelez et al., 2010). One approach for IV estimation in the linear i.i.d. model is the two-stage least squares (TSLS) estimator (Angrist and Imbens, 1995), which first estimates the effect from I to X (stage 1) and then regresses Y on the fitted values from the first regression (stage 2). Another

formulation, used by Hansen (1982), is the generalized method of moments (GMM), which uses the independence of the residual $Y - \beta X = g(H, \varepsilon^Y)$ from the instrument I : One can estimate β by selecting $\hat{\beta}$ such that the empirical correlation between $Y - \hat{\beta}X$ and I is minimized. If the dimension of I is greater than or equal to the dimension of X , these estimators are consistent (e.g., Hall, 2005).

In more recent approaches, causality and directed acyclic graph (DAG) representations have proved fruitful for studying instrumental variables for i.i.d. data (Pearl, 2009; Hernán and Robins, 2006; Didelez et al., 2010). Brito and Pearl (2002) proposed ‘generalized IV’, a graphical framework that enlarges the class of graphical models, in which IV methods can be used to identify causal effects. Similarly, ‘conditional IV’ (Pearl, 2009) relaxes the assumptions of IV by considering a conditional moment equation (see also Henckel, 2021). Moreover, Kang et al. (2016) demonstrate that identification and consistent estimation are possible when at least half of the instruments are valid, even without knowing which ones are invalid.

In many real-world applications (see Weigend (2018) for examples from various fields), the data are sampled not independently but rather as a time series that exhibits memory effects, with past values affecting present ones. For example, both price and demand on an electricity market are confounded by several factors, which makes estimation of price elasticity of demand (that is, the response in the demand to changes in price) difficult to identify. In this example, wind power penetration of the market may act as instrumental variable, since it affects the energy supply and in turn quantity purchased at that price (Neamtu, 2016; Hirth et al., 2022); it can be justified as a valid instrument since the daily amount of generated wind power depends on external weather conditions, but not other aspects of the market supply or demand. This is similar to how yield per acre acted as ‘curve shifter’ in the seminal work of Wright (1928) for estimating the price elasticity of demand for flaxseed.

Using IV methods in time series data poses a number of challenges that are not present when considering IV methods for i.i.d. data. For example, memory effects in the observed processes X , Y and I can obfuscate the assumption that I only affects present values of Y through the present value of X , because I and Y are confounded by common ancestors in the memory of the process. Later in this introduction, we show with a concrete example how these challenges arise in linear models. Additionally, memory effects, or serially correlated errors, in the confounder process H can make identification of the dependence on past states of the process difficult; for such settings, Fair (1970) proposes a search-based method for a subclass of first order vector auto-regressive (VAR) processes. If one is provided with identifying equations with serially correlated errors (such as the ones proposed in this paper), Newey and West (1987) construct confidence intervals by using heteroskedasticity and auto-correlation consistent (HAC) estimators to estimate long-run covariance matrices.

To overcome the challenges for IV methods in time series data, we establish a link between graphical models and IV methods for time series, which we then exploit to construct estimators and prove consistency. To help build this connection, we require that the global Markov property (Lauritzen, 1996) holds in VAR(p) processes. We prove this statement as Theorem 1. To the best of our knowledge this result is more general than existing results (e.g., Dahlhaus and Eichler, 2003; Lauritzen, 1996); its proof requires technical arguments taking into account that the graphs contain infinitely many nodes.

Causal inference on time series data has been considered before. E.g., causal inference can then be done using the principle of Granger causality (e.g., Wiener, 1956; Granger, 1969, 1980; Dahlhaus and Eichler, 2003; Eichler, 2007; Didelez et al., 2010) but such methods usually fail when some of the variables are unobserved (e.g., Peters et al., 2017, Chapter 10.3). It is also possible to extend independence-based methods (Spirtes et al., 2000; Pearl, 2009) to time series (Demiralp and Hoover, 2003; Chu and Glymour, 2008; Entner and Hoyer, 2010; Moneta et al., 2011) but as in the i.i.d. case, these methods cannot exploit identifiability that stems from IV conditions. Adjustment formulas (and its modifications such as the front-door criterion) have been extended to time series with known graph structure (Eichler and Didelez, 2010) but the specific form of confounding assumed in our setting does not allow for consistent estimation of the causal effect using these techniques – again,




Figure 1: Finite excerpt of a full time graph (formally defined in Section 2.2) of a VAR(1) process S satisfying (A2) below. We assume that each time point is observed only once. Our methodology estimates the causal effect β (highlighted in green) of X_{t-1} on Y_t , where $X = [X_t]_{t \in \mathbb{Z}}$ and $Y = [Y_t]_{t \in \mathbb{Z}}$ are subprocesses of S , that are confounded by a latent process $H = [H_t]_{t \in \mathbb{Z}}$ (so that, in general, standard regression yields an inconsistent estimator). Motivated by instrumental variables, one may aim to exploit the subprocess $I = [I_t]_{t \in \mathbb{Z}}$ that is independent of H and only acts on Y through X . However, simply using I_{t-2} as an instrument is generally inconsistent; the same holds when adding X_{t-2} and Y_{t-1} as a conditioning set, for example. This paper develops a graphical framework giving rise to several estimators which consistently estimate β .

this is similar to the i.i.d. case. Michael et al. (2024) use instrumental variables for longitudinal data; different from our work they consider binary instruments and primarily consider a setting where time series are observed multiple times (for example each corresponding to a patient), though the marginal structural mean model (Robins, 1997) that they employ also allows for estimation within a single observation of a time-series. Their target of inference differs to ours in that Michael et al. (2024) consider interventions on the whole treatment time series. The structural assumptions (such as independence of instrument and confounder) underlying their method are similar to those required by the estimators we develop below, though generally not identical.

Recently, Mogensen (2023) uses integrated covariances to conduct instrumental variable estimation of ‘normalized parameters’, and shows that these parameters have causal interpretations in both discrete time series and continuous time Hawkes processes.

Throughout this work we consider a joint process $S := [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ where H is latent. As a motivating example, let S be a linear VAR(1) process (though we generally consider VAR(p) processes), with dependencies represented in Figure 1, and assume that we observe a finite subsample S_1, \dots, S_T from a single instantiation of this process. Our goal is to estimate the coefficient β with which X_{t-1} linearly enters into Y_t . When S is fully observed, estimators that are consistent and asymptotically normal for the standard form parameters of the VAR(1) process S exist (e.g., Hamilton, 1994). Yet, in our setting, H , the confounding variable between X and Y , is unobserved, and such estimators are then generally not consistent. In this work, we seek to overcome this hidden confounding by using I_{t-2} as an instrument. Estimation in the time series setting is complicated by memory effects not present in the i.i.d. setting: For example, the instrument I_{t-2} is correlated to Y_t not only through the path $I_{t-2} \rightarrow X_{t-1} \rightarrow Y_t$ but also through an infinite number of paths in the past, due to common ancestors I_{t-j} , $j \geq 3$ in the instrument process I . This correlation violates the assumption that the instrument I_{t-2} only correlates with Y_t through X_{t-1} . Similarly, there are also observed confounders of X_{t-1} and Y_t , such as X_{t-2} which is a common ancestor of the two.

To obtain valid identifying IV methods for VAR(p) processes S , such as the one shown in Figure 1, we establish a general graph marginalization technique allowing us to read off the relevant separation statements from the reduced graphs. Based on these results, we outline two solutions that identify total causal effects in the considered time series. The first solution (which is based on ‘conditional

IV’ or ‘CIV’) identifies β using IV conditioned on some subset of past states of the time series. In the example above, we will see that, e.g., $I_{t-2} \perp\!\!\!\perp Y_t - \beta X_{t-1} | \{I_{t-3}, X_{t-2}, Y_{t-1}\}$ is an identifying equation.

The second solution is based on a modification of IV that can be used not only for time series but also for i.i.d. data. It adds nuisance treatment variables to the target causal effect and thereby allows for stronger identifiability results. This is a straightforward idea but we are not aware that this has been discussed explicitly, so we suggest to call it ‘nuisance IV’ (or ‘NIV’). Applied to the time series setting, nuisance IV yields a consistent estimator for the target of inference β by including nuisance regressors, such as Y_{t-1} , that also affect Y_t . In the example above, it yields the identifying equation $\{I_{t-2}, \dots, I_{t-m-1}\} \perp\!\!\!\perp Y_t - \beta X_{t-1} - \kappa Y_{t-1}$ for some m . As for the first solution, we will detail the reasoning behind these equations and conditions for identifiability.

Similar to the i.i.d. case, these two approaches induce identifying moment equations that are satisfied by the total causal effect β . Rank conditions guarantee that their solution is unique, allowing us to identify the causal effect. Unlike in the i.i.d. case, however, the standard conditions are not easily interpretable in the time series setting. We therefore develop sufficient and necessary conditions on the parameters of the data-generating process that provide insight on when identifiability holds. Our results imply that identifiability with nuisance IV depends on geometric multiplicities of eigenvalues of the parameter matrix in the VAR(1) process, and we show that if parameters are drawn at random from a continuous distribution, the causal effect β of X_{t-1} on Y_t is almost surely identifiable. In particular, it is possible to identify the causal effect even if the instrument I is univariate and the regressor X is multivariate. For both of the approaches (conditional IV and nuisance IV), we propose estimators and prove that, in case of identifiability, these estimators consistently estimate the causal effect.

Finally, we apply our findings to the task of distribution generalization (e.g., Christiansen et al., 2022; Jakobsen and Peters, 2022; Meinshausen, 2018; Rothenhäusler et al., 2021). In many systems, the causal effects are of value in themselves because they contribute to the understanding of the system but it also serves a purpose when predicting Y_{t+1} under an intervention on X_t . In a linear setting, the OLS estimator has the smallest expected mean squared error (MSE) among all linear predictors when predicting new test data from the observational distribution. However, as is known for the i.i.d. setting (e.g., Rojas-Carulla et al., 2018), causal estimators can have better worst-case predictive performance when there may be interventions on the covariates. Similarly, we show that in time series, under arbitrary interventions on X_t , our IV estimators are worst-case prediction optimal for Y_{t+1} .

Our work is structured as follows. Section 2 introduces the model and the assumptions considered in this paper; we review graphical representations of time series models and prove that the global Markov property holds for VAR(p) processes. In Section 3, we review theory on conditional instrumental variables and introduce the concept of nuisance IV. Our main results for instrumental variable regression for time series are presented in Section 4. We propose two approaches to overcome confounding from past values yielding identifying equations for the total causal effect: the first one is based on CIV and the second one uses NIV. For the latter, we characterize identifiability of the causal parameter in terms of parameters of the data-generating process. We also discuss how to use the causal effect to perform optimal prediction of Y_{t+1} under interventions on X_t . In Section 5 we empirically evaluate our method. All proofs are provided in Appendix E. Code can be found at <https://github.com/nikolajthams/its-time>.

2. Causal Time Series Models with Confounding

2.1 Definitions and Notation

We consider multivariate time series $X := [X_t]_{t \in \mathbb{Z}}$ and $I := [I_t]_{t \in \mathbb{Z}}$, a univariate process $Y := [Y_t]_{t \in \mathbb{Z}}$, and an unobserved multivariate process $H := [H_t]_{t \in \mathbb{Z}}$. Let d_X be the dimensionality of X_t , that is

$X_t \in \mathbb{R}^{d_X}$, and similarly for d_I , d_Y and d_H , with $d_Y = 1$. Let $S := [S_t]_{t \in \mathbb{Z}} = [I_t^\top, H_t^\top, X_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$, $S_t \in \mathbb{R}^d$, with $d := d_X + d_Y + d_I + d_H$.

In general, our results are presented for VAR(p) processes (e.g., Brockwell and Davis, 1991), though we apply our theory to the particular example of a VAR(1) process in Sections 4.2.1, 4.2.2, and 4.3. Many of the results hold more generally than in VAR(p) processes, which we discuss in Section 4.2.3. Given a (known) $p \in \mathbb{N}$, we say that a (weakly) stationary process S is a VAR(p) process if the following assumptions hold:

(A0) There are coefficient matrices $A_1, \dots, A_p \in \mathbb{R}^{d \times d}$ such that for all $t \in \mathbb{Z}$:

$$S_t = A_1 S_{t-1} + \dots + A_p S_{t-p} + \varepsilon_t, \quad (1)$$

where A_1, \dots, A_p are such that¹ $\det(I_d \lambda^p - A_1 \lambda^{p-1} - A_2 \lambda^{p-2} - \dots - A_p) = 0$ implies $|\lambda| < 1$ and the ε_t constitute an i.i.d. process with finite second moments.

(A1) (A0) is satisfied and, in addition $\varepsilon_t \sim \mathcal{N}(0, \Gamma)$, where Γ is a diagonal matrix.

(That is, unless stated otherwise, we assume that the VAR processes in this paper are Gaussian.) We use the notation $\alpha_{X,I}^1$ to refer to the submatrix of A_1 with rows corresponding to X and columns corresponding to I (see Figure 2a for an example), and similarly $\alpha_{I,I}^2$ etc. We use superscripts to denote individual components of ε , e.g., ε^Y . We consider Y_t as the *response variable*, and aim to estimate the total causal effect (TCE) of covariates on Y_t . For $1 \leq i, j \leq d$ and $l \in \mathbb{N}$, the TCE of S_{t-l}^i on S_t^j is defined as


$$\beta := \left(\sum_{\substack{1 \leq l_1, \dots, l_m \leq p \\ l_1 + \dots + l_m = l}} A_{l_1} \cdots A_{l_m} \right)_{j,i};$$

see Appendix A.4 for more details on defining the TCE. An important example is when we use X_{t-1} as covariates; in this case, the TCE equals $\beta = \alpha_{Y,X}^1$, which is also called the *direct causal effect* of X_{t-1} on Y_t . Unless specified otherwise, all causal effects in this paper refer to total causal effects.²

Both H and the noise ε are assumed to be unobserved; while the sequence of innovations ε_t^Y is assumed to be i.i.d. and independent of X_t , H can act as a confounder between X and Y and can have an autoregressive structure. Similar to the i.i.d. case (e.g., Hernán and Robins, 2006; Pearl, 2009; Peters et al., 2017), the existence of the confounder H implies that we cannot identify β by simply regressing Y_t on X_{t-1} . In Appendix A.2, we provide an example of two VAR(1) processes with two different parameter matrices that generate the same distribution over the observed time series.

We assume that the process has zero mean³, so for instance $\text{cov}(X_t, Y_t) = \mathbb{E}\{X_t Y_t^\top\} \in \mathbb{R}^{d_X \times d_Y}$. We assume that the data are sampled as follows: We obtain a sample over time points $t = 1, \dots, T$ of S such that for $t = 1$, S_1 follows the stationary distribution. We denote the sample with boldface $\mathbf{S} = [\mathbf{S}_t]_{t=1}^T$ such that $\mathbf{S} \in \mathbb{R}^{d \times T}$ and each column \mathbf{S}_t represents the process observation at time t . We let $\hat{\mathbb{E}}\mathbf{S}_t := \frac{1}{T} \sum_{t=1}^T \mathbf{S}_t$ denote the empirical mean of the process (here, the index t in $\hat{\mathbb{E}}\mathbf{S}_t$ does not refer to any specific time point). From (A1) it follows (Hamilton, 1994, Chapter 10) that for

-
1. A VAR(p) process satisfying this condition is sometimes referred to as a causal VAR(p) process (Brockwell and Davis, 1991); Peters et al. (2013) discuss a relation between this causality property and the independent noise assumption in SCMs.
 2. The notions of causal effect and total causal effect are motivated by interpreting the VAR(p) equations as a structural causal model (SCM), which we explain in detail in Appendix A.3. The interventional interpretation of an SCM is not required for any results of the paper, except for the ones presented in Section 4.3, where we discuss optimal predictions under interventions. Considering the model as an interventional model implies that H is the only unobserved process of relevance (see ‘interventional sufficiency’ Peters et al., 2017, Chapter 9).
 3. Since we can always subtract empirical means, the assumption of vanishing means does not come with any loss of generality.



(a) Matrix block structure of A_1 assumed in (A2). (b) Summary time graph of a VAR(1) process satisfying (A2).

Figure 2: The sparsity structure on the parameter matrix A_1 assumed in (A2), and a representation of the graphical structure induced by A_1 . Zeros in panel a) correspond to absent edges in panel b).

$T \rightarrow \infty$

$$\hat{\mathbb{E}}\mathbf{S}_t \xrightarrow{P} \mathbb{E}S_t \quad \text{and} \quad \text{cov}\{\mathbf{S}_{t-j}, \mathbf{S}_t\} := \frac{1}{T-j} \sum_{t=j+1}^T \mathbf{S}_{t-j} \mathbf{S}_t^\top \xrightarrow{P} \text{cov}\{S_{t-j}, S_t\}. \quad (2)$$

Even though we assume that only a finite time window \mathbf{S} is observed, we assume for simplicity that the data are generated (but not observed) over all time points $t \in \mathbb{Z}$, to ensure that subprocesses of S are correlated in the same way throughout time. The same assumption is common in the literature (e.g. Brockwell and Davis, 1991; Hamilton, 1994).

Finally, in the case where S is a VAR(1) process, we will sometimes assume additional structure on the coefficient matrix.

(A2) Assume that S satisfies (A1) for $p = 1$ and that A_1 has the sparsity structure displayed in Figure 2a.

Under (A2), none of the other time series components enters the assignment for I and I itself neither appears in the assignments for Y nor in the one of H ; the assignment for Y reduces to

$$Y_t = \alpha_{Y,Y}Y_{t-1} + \beta X_{t-1} + \eta_t$$

with $\eta_t := \alpha_{Y,H}H_{t-1} + \varepsilon_t^Y$. We refer to I as an *instrumental time series*. As in the i.i.d. case, (A2) cannot in general be tested from data and must therefore rest on background knowledge.

2.2 Graph Representations of VAR(p) Processes

Equation (1) can be represented by a directed graph. We choose a representation with infinitely many nodes that follows directly from the structural coefficients and is similar or identical to other representations that have been discussed in the literature (see, e.g., Dahlhaus and Eichler, 2003; Peters et al., 2013). The graph will prove helpful when establishing identifying equations for causal effects and constructing consistent estimators. The *full time graph* (e.g., Peters et al., 2013) is defined as an infinite directed graph with nodes I_t, H_t, X_t, Y_t , for any $t \in \mathbb{Z}$. For $k \in \mathbb{N}$, it contains a directed edge from (j, t) to $(i, t+k)$, $i, j \in \{I, H, X, Y\}$, if $\alpha_{ij}^k \neq 0$. For a VAR(1) process satisfying (A2) an extract of this graph is shown in Figure 1. We define the full time graph for higher-order VAR(p) processes accordingly. The *summary time graph* only has a single node per time series component. It contains a directed edge from i to j , for some $i, j \in \{I, H, X, Y\}$ if and only if the full time graph contains an edge from (j, t) to $(i, t+k)$ for some $k \in \mathbb{N}$. For a VAR(1) process satisfying (A2), such a graph is visualized in Figure 2b. (The difference between the full time graph and a

summary time graph is also sometimes referred to as ‘unrolling’ or ‘unfolding’.) Neither of these two graphical representations requires Gaussianity (that is, the constructions only require (A0)). We now introduce some standard graph terminology (e.g., Lauritzen, 1996; Koller and Friedman, 2009; Pearl, 2009). A *path* p , $p = (v_1, e_1, v_2, \dots, e_{n-1}, v_n)$, is an alternating sequence of distinct vertices v_i and edges e_i such that v_i and v_{i+1} are connected by e_i . We say that p is a *directed path* from v_1 to v_n if for every i , e_i points from v_i to v_{i+1} . For two nodes v and u , we say that u is a *descendant* of v if there exists a directed path from v to u and otherwise u is a *non-descendant* of v . We write $\text{ND}(v)$ and $\text{DE}(v)$ for the sets of non-descendants and descendants of v , respectively, using the convention that neither of them contain v itself. For a path p and any $i \in \{2, \dots, n-1\}$, we say that v_i is a *collider on p* if $(v_{i-1}, e_{i-1}, v_i, e_i, v_{i+1})$ is of the form $v_{i-1} \rightarrow v_i \leftarrow v_{i+1}$ and else v_i is a *non-collider on p* . We say that the path p is *unblocked, given the set B* if for every non-collider v_i in p , $v_i \notin B$ and for every collider v_i on p , $(v_i \cup \text{DE}(v_i)) \cap B \neq \emptyset$. Otherwise, we say that p is *blocked* by B . If all paths between distinct vertices v and u are blocked by a set B neither containing v nor u , we say that v and u are *d -separated* by B . Similarly, we say that disjoint sets V and U are *d -separated* by B if all nodes $v \in V$ and $u \in U$ are *d -separated* by B .

In Section 4.2, we also consider marginalized graphs, which are acyclic directed mixed graphs (ADMGs), containing both directed (\rightarrow) and bidirected (\leftrightarrow) edges. If we define v to be a collider on a path whenever two surrounding edges have arrowheads at v (e.g. $u_1 \leftrightarrow v \leftarrow u_2$), and define descendants only with respect to directed edges, *d -separation* also extends to ADMGs. See Richardson (2003) for details.

2.2.1 MARKOV PROPERTIES OF VAR(p) PROCESSES

The representation described above satisfies a Markov property, which enables us to read off conditional independences from the full time graph. This will be an important tool in our theory (and will be used in many of the proofs), because it enables the use of graphical models to develop IV methodology in time series. Markov properties of time series have been discussed before (e.g., Dahlhaus and Eichler, 2003; Eichler, 2012) but to the best of our knowledge, the statement of Theorem 1 has not been proved before, so we add a proof in this paper. For example, some of the results in the above works consider finite graphs while the statement of Theorem 1 is about graphs with infinitely many nodes. For this reason, the result does not formally follow from standard results in, for example, Lauritzen (1996), either.

Theorem 1. *Consider $p \in \mathbb{N}$, a (weakly) stationary time series S satisfying (A0) for that p , and consider disjoint, finite sets A, B, C of nodes in the full time graph $\mathcal{G}_{\text{full}}$. If A and C are d -separated given B in $\mathcal{G}_{\text{full}}$, then $A \perp\!\!\!\perp C \mid B$.*

The proof of Theorem 1 is inspired by Lauritzen et al. (1990, Sec. 6). As all other proofs in this paper, it can be found in Appendix E.

3. Nuisance Effects in Instrumental Variable Regression

In this work, we establish two identifying equations for causal effect estimation in time series: The first one is based on conditional instrumental variables (CIV) and the second one on a generalization that we term nuisance instrumental variables (NIV). We regard the idea of NIV as interesting in its own right, as it can be applied in the i.i.d. setting, too. In this section, we therefore first review CIV regression for i.i.d. data, and then introduce NIV regression; instrumental variable (IV) regression is a special case of CIV regression, where the conditioning set is empty. In Section 4, we extend the CIV and NIV estimators to VAR(p) processes via a reduction of the full time graph to a marginalized graph.

3.1 Instrumental Variables and Conditional Instrumental Variables

Consider a linear SCM (see Appendix A.3) over variables V , and let $\mathcal{I}, \mathcal{X}, \mathcal{B}, \{Y\} \subseteq V$ be disjoint collections of variables⁴ from V , and let \mathcal{G} be the corresponding DAG. Assume that \mathcal{I}, \mathcal{X} and Y have zero mean and finite second moment and let β be the causal coefficient with which \mathcal{X} enters the structural equation for Y , that is,

$$Y = \beta\mathcal{X} + \gamma W + \varepsilon^Y,$$

for some variables $W \subseteq V \setminus \mathcal{X}$; (some of the entries of β can be zero, so not all variables in \mathcal{X} have to be parents of Y). In this setup, β is both the total and the direct causal effect of \mathcal{X} on Y . We consider the following three requirements on $\mathcal{I}, \mathcal{X}, \mathcal{B}$ and Y :

(CIV1) \mathcal{I} and Y are d -separated given \mathcal{B} in the graph $\mathcal{G}_{\mathcal{X} \not\rightarrow Y}$, that is the graph \mathcal{G} where all direct edges from \mathcal{X} to Y are removed,

(CIV2) \mathcal{B} does not contain a descendant of $\mathcal{X} \cup \{Y\}$ in \mathcal{G} , and

(CIV3) the matrix $\mathbb{E}[\text{cov}(\mathcal{X}, \mathcal{I}|\mathcal{B})]$ has rank $d_{\mathcal{X}}$, that is, full row rank.

If requirements (CIV1) and (CIV2) are met, $Y - \beta\mathcal{X} \perp\!\!\!\perp \mathcal{I}|\mathcal{B}$, and in particular β satisfies the *CIV moment equation*⁵

$$\mathbb{E}[\text{cov}(Y - \beta\mathcal{X}, \mathcal{I}|\mathcal{B})] = 0. \quad (3)$$

If, additionally, requirement (CIV3) is met, β is the unique solution to Equation (3),

$$\mathbb{E}[\text{cov}(Y - b\mathcal{X}, \mathcal{I}|\mathcal{B})] = 0 \implies b = \beta.$$

(Conditional IV with univariate \mathcal{X} has been discussed in the literature (Pearl, 2009; Henckel, 2021; Brito and Pearl, 2002). Since we allow $d_{\mathcal{X}} > 1$, we add a short proof in Appendix E.1.) In this case, we say that β is *identified by CIV* or, more precisely, *identified by CIV $_{\mathcal{X} \rightarrow Y}(\mathcal{I}|\mathcal{B})$* . If requirements (CIV1) and (CIV3) are satisfied for $\mathcal{B} = \emptyset$ (requirement (CIV2) is trivially satisfied for $\mathcal{B} = \emptyset$), $\text{CIV}_{\mathcal{X} \rightarrow Y}(\mathcal{I}|\emptyset)$ reduces to instrumental variables (IV) regression (Reiersøl, 1945; Anderson and Rubin, 1949; Bowden and Turkington, 1985; Angrist et al., 1996), which we refer to as $\text{IV}_{\mathcal{X} \rightarrow Y}(\mathcal{I})$. We use the term *proper CIV* when $\mathcal{B} \neq \emptyset$.

For a finite sample $\mathbf{X}, \mathbf{Y}, \mathbf{I}$, and \mathbf{B} , we consider an empirical counterpart of Equation (3) which, however, may not have a solution in the overidentified setting, that is when $d_{\mathcal{I}} > d_{\mathcal{X}}$; to overcome this, for any positive definite weight matrix W , we define the estimator $\hat{b}(W)$ as

$$\hat{b}(W) := \arg \min_b \|\text{côv}(\mathbf{Y} - b\mathbf{X}, \mathbf{I}|\mathbf{B})\|_W^2, \quad (4)$$

where $\|x\|_W^2 := x^\top W x$ and $\text{côv}(\mathbf{Y} - b\mathbf{X}, \mathbf{I}|\mathbf{B})$ is the empirical covariance of the residuals after regressing out \mathbf{B} . We refer to this estimator as $\text{CIV}_{\mathbf{X} \rightarrow \mathbf{Y}}(\mathbf{I}|\mathbf{B})$. If $\mathcal{I}, \mathcal{X}, Y$ and \mathcal{B} are zero mean random vectors, the minimizer of Equation (4) is given by

$$\hat{b}(W) = \hat{\mathbb{E}}[r_{\mathbf{Y}} r_{\mathbf{I}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}} r_{\mathbf{X}}^\top] \left(\hat{\mathbb{E}}[r_{\mathbf{X}} r_{\mathbf{I}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}} r_{\mathbf{X}}^\top] \right)^{-1}, \quad (5)$$

where $r_{\mathbf{Y}} := \mathbf{Y} - \hat{\mathbb{E}}[\mathbf{Y}|\mathbf{B}]$ are the residuals after regressing \mathbf{Y} on \mathbf{B} , and similarly $r_{\mathbf{X}} := \mathbf{X} - \hat{\mathbb{E}}[\mathbf{X}|\mathbf{B}]$ and $r_{\mathbf{I}} := \mathbf{I} - \hat{\mathbb{E}}[\mathbf{I}|\mathbf{B}]$.

4. Below, the different variables will take different roles (such as instruments or regressors). We use the calligraphic notation \mathcal{I}, \mathcal{X} , and \mathcal{B} to denote collections of observed variables, being used as instruments, regressors, and conditioning sets, respectively. Individual variables are denoted by non-calligraphic letters, such as I .

5. Here we use the definition $\text{cov}(A, C|B) := \mathbb{E}[AC^\top|B] - \mathbb{E}[A|B]\mathbb{E}[C^\top|B] = \text{cov}(A - \mathbb{E}[A|B], C - \mathbb{E}[C|B])$, which even accommodates for nonlinear relationships between the variables and \mathcal{B} .




Figure 3: CIV and (proper) NIV are complementary for identifying causal effects in that they can be used in different settings. (*Left*) A graph where the effect $X \rightarrow Y$ can be identified by $\text{CIV}_{X \rightarrow Y}(I|B)$ (provided that requirement (CIV3) holds) but not by any IV or proper NIV method. (*Middle*) A graph where $X \rightarrow Y$ can be identified by $\text{NIV}_{X \rightarrow Y}(I, Z)$, but not by any proper CIV method. (*Right*) A graph where $X \rightarrow Y$ can be identified by both $\text{CIV}_{X \rightarrow Y}(I|B)$ and $\text{NIV}_{X \rightarrow Y}(\mathcal{I}, Z)$ with $\mathcal{I} = \{I, B\}$.

Choosing the two-stage least squares (TSLS) weight matrix $W_{\text{TSLS}} := \mathbb{E}[r_{\mathbf{I}} r_{\mathbf{I}}^\top]^{-1}$ corresponds to the procedure where one regresses $r_{\mathbf{X}}$ on $r_{\mathbf{I}}$ and then returns the regression coefficient of $r_{\mathbf{Y}}$ on the fitted values $\hat{r}_{\mathbf{X}}$. In a linear Gaussian model, the IV estimator $\hat{b}(W_{\text{TSLS}})$ has the lowest asymptotic variance among all positive definite weight matrices W (Hall, 2005) (see Appendix B.1, where we review the asymptotic theory for IV estimators in i.i.d. data).

3.2 Nuisance instrumental variables

IV estimation is a special case of CIV with the empty set as conditioning set. The example in Figure 3 (left) shows a graph where an effect between X and Y cannot be identified using the IV estimator: no variable is d -separated from H , and hence no valid instruments for (unconditional) IV exist. Yet, the causal effect is identified by⁶ $\text{CIV}_{X \rightarrow Y}(I|B)$ because B satisfies requirements (CIV1) to (CIV3).

In the case shown in Figure 3 (middle), the effect from X to Y cannot be identified by $\text{IV}_{X \rightarrow Y}(I)$ because of the unblocked path $I \rightarrow Z \rightarrow Y$. We cannot use proper CIV, either, because the path $I \rightarrow Z \leftarrow H \rightarrow Y$ is unblocked given Z , violating requirement (CIV1). Nevertheless, we can identify the effect from X to Y by adding an additional regressor variable. If $d_I \geq d_X + d_Z$, then $\text{IV}_{\{X, Z\} \rightarrow Y}(I)$ satisfies the assumption for identifying the effect $\{X, Z\} \rightarrow Y$. In particular, from this we can extract the effect of interest, $X \rightarrow Y$.

We refer to this approach as *nuisance instrumental variables (NIV)*. More formally, consider collections of variables $\mathcal{I}, \mathcal{X}, \mathcal{Z}$ and a response variable Y . We say that β satisfies the *NIV moment equation* if there exist $\alpha \in \mathbb{R}^{d_Y \times d_Z}$ such that

$$\text{cov}(Y - \beta\mathcal{X} - \alpha\mathcal{Z}, \mathcal{I}) = 0. \quad (6)$$

We say that β is *identified by NIV* or, more formally *identified by* $\text{NIV}_{\mathcal{X} \rightarrow \mathcal{Y}}(I, Z)$ if additionally β is the only solution to the moment equation; that is for all $a \in \mathbb{R}^{d_Y \times d_Z}$ and $b \in \mathbb{R}^{d_Y \times d_X}$

$$\text{cov}(Y - b\mathcal{X} - a\mathcal{Z}, \mathcal{I}) = 0 \implies b = \beta. \quad (7)$$

We refer to \mathcal{Z} as a *nuisance regressor*. If we use both a nuisance regressor \mathcal{Z} and condition on a variable \mathcal{B} , the conditions become

there exists α s.t. $\mathbb{E}[\text{cov}(Y - \beta\mathcal{X} - \alpha\mathcal{Z}, \mathcal{I}|\mathcal{B})] = 0$ and $\mathbb{E}[\text{cov}(Y - b\mathcal{X} - a\mathcal{Z}, \mathcal{I}|\mathcal{B})] = 0 \implies b = \beta$,

6. In a slight abuse of notation, we sometimes omit parentheses indicating sets and write $\text{CIV}_{\mathcal{X} \rightarrow \mathcal{Y}}(I|\mathcal{B})$ instead of $\text{CIV}_{\mathcal{X} \rightarrow \mathcal{Y}}(\mathcal{I}|\mathcal{B})$ if $\mathcal{I} = \{I\}$, for example.

and we write $\text{NIV}_{\mathcal{X} \rightarrow Y}(\mathcal{I}, \mathcal{Z}|\mathcal{B})$; this corresponds to extracting the entries relevant for \mathcal{X} from the output of $\text{CIV}_{\mathcal{X} \cup \mathcal{Z} \rightarrow Y}(\mathcal{I}|\mathcal{B})$.⁷ By choosing $\mathcal{Z} = \emptyset$, NIV extends CIV. When $\mathcal{Z} \neq \emptyset$, we use the term *proper NIV*. The following theorem proves that requirements (CIV1) to (CIV3) are sufficient to establish identifiability of NIV.

Theorem 2 (Nuisance IV). *Consider a linear SCM (see Appendix A.3) over variables V , and let $\mathcal{I}, \mathcal{X}, \mathcal{Z}, \mathcal{B}, \{Y\} \subseteq V$ be disjoint collections of variables from V , and let \mathcal{G} be the corresponding DAG. Assume that $\mathcal{I}, \mathcal{X}, \mathcal{Z}$ and Y have zero mean and finite second moment and let β and α be the causal coefficients with which \mathcal{X} and \mathcal{Z} enter the structural equation for Y , respectively (some of the entries of β and α can be zero, so not all variables in \mathcal{X} and \mathcal{Z} have to be parents of Y). Let $\tilde{\mathcal{X}} := \mathcal{X} \cup \mathcal{Z}$. If requirements (CIV1) to (CIV3) are satisfied in \mathcal{G} for $\mathcal{I}, \tilde{\mathcal{X}}, \mathcal{B}$ and Y , the causal effect β of \mathcal{X} on Y is identified by $\text{NIV}_{\mathcal{X} \rightarrow Y}(\mathcal{I}, \mathcal{Z}|\mathcal{B})$.*

Even though this is a straight-forward extension of IV regression, we are not aware of any work describing the idea of NIV. It will prove useful in the time series setting and even in the i.i.d. setting it is a strict generalization of CIV: there are graphs, such as the one in Figure 3 (middle), where the causal effect $X \rightarrow Y$ is neither identified by IV nor by CIV.

For some graphs the effect $X \rightarrow Y$ can be identified by (proper) NIV and (proper) CIV. For example, in the graph in Figure 3 (right) the effect $X \rightarrow Y$ can be identified both by $\text{CIV}_{X \rightarrow Y}(I|B)$ and by $\text{NIV}_{X \rightarrow Y}(\{I, B\}, Z)$. When estimated from a finite sample, the two resulting estimators are not identical, and, as the following proposition establishes, the two approaches cannot in general be sorted in terms of asymptotic variance.⁸

Proposition 3. *If an effect can be identified by CIV and by NIV, then the estimators cannot be strictly sorted in terms of asymptotic variance. More specifically, there exist data generating processes, for which CIV has strictly smaller asymptotic variance and others, for which NIV has strictly smaller asymptotic variance.*

The idea of NIV can be naturally applied in time series settings, too. In Section 4.2 we show that causal effects in VAR(1) processes as described in Section 2 can be estimated both by CIV and NIV. To establish this result, we first develop a marginalization technique of time series graphs, which allows us to apply the above results.

4. Instrumental Time Series Regression

4.1 Time Series Reduction

In Section 2.2, a VAR(p) process is represented by its full time graph (see, e.g., Figure 1). We now show that instruments and conditioning sets for VAR(p) processes can be found by considering marginalized time graphs, which are obtained by marginalization of the full time graph to a finite set of nodes; they resemble latent projections (e.g., Verma, 1991) but are projections of graphs that are not finite.

Definition 4. *Consider a process $S = [S_t]_{t \in \mathbb{Z}}$ satisfying (A1) and let \mathcal{G}_{full} be the full time graph of S as defined in Section 2.2. Let $M = \{S_{t_1}^{i_1}, \dots, S_{t_m}^{i_m}\}$ be a finite collection of nodes in \mathcal{G}_{full} . The marginalized time graph, \mathcal{G}_M , is the graph over nodes M where for all $i, j \in M$ there is:*

- i) a directed edge $i \rightarrow j$ if and only if $i \rightarrow j$ in \mathcal{G}_{full} or there exists $m_1 \in \mathbb{N}$, $v_1, \dots, v_{m_1} \notin M$ and a directed path $i \rightarrow v_1 \rightarrow \dots \rightarrow v_{m_1} \rightarrow j$ in \mathcal{G}_{full} , and*

7. As such, one could also write $\text{CIV}_{\underline{\mathcal{X}} \cup \mathcal{Z} \rightarrow Y}(\mathcal{I}|\mathcal{B})$ instead of $\text{NIV}_{\mathcal{X} \rightarrow Y}(\mathcal{I}, \mathcal{Z}|\mathcal{B})$, where the underline indicates which causal effect one is interested in.

8. In the i.i.d. setting, the asymptotic variances of both NIV and CIV estimators can be described by closed form expressions, see Appendix B.1.

ii) a bidirected edge $i \leftrightarrow j$ if and only if there exists $m_1, m_2 \in \mathbb{N}$, $v_1, \dots, v_{m_1}, w_1, \dots, w_{m_2}, U \notin M$ in \mathcal{G}_{full} such that there exists directed paths $U \rightarrow v_1 \rightarrow \dots \rightarrow v_{m_1} \rightarrow i$ and $U \rightarrow w_1 \rightarrow \dots \rightarrow w_{m_2} \rightarrow j$.

The following theorem establishes that the CIV conditions being satisfied in a marginalized time graph implies a moment condition that can be used for identifying the causal effect. In the previous section, we stated identifiability results in terms of the causal effect β . The following theorem is stated in terms of the total causal effect (see Section 2.1); this generalizes the causal effect and may for instance be interesting if some predictors X are unobserved or are observed but cannot be intervened on. To do so, we state a slight modification of requirement (CIV1).

(CIV1') \mathcal{I} and Y are d -separated given \mathcal{B} in the graph where we remove all edges from a node $i \in \mathcal{X}$ to a node $j \notin \mathcal{X}$ if the edge lies on a directed path from \mathcal{X} to Y that contains only a single node in \mathcal{X} (namely i).

Theorem 5 (Time series IV by marginalization). *Consider a process $S = [S_t]_{t \in \mathbb{Z}}$ satisfying (A1) with full time graph \mathcal{G}_{full} . Let Y be some node in \mathcal{G}_{full} and let $\mathcal{X}, \mathcal{I}, \mathcal{Z}$, and \mathcal{B} be disjoint collections of nodes from \mathcal{G}_{full} . Let $\tilde{\mathcal{X}} := \mathcal{X} \cup \mathcal{Z}$ and define $M := \{Y\} \cup \mathcal{X} \cup \mathcal{I} \cup \mathcal{Z} \cup \mathcal{B}$. Assume that requirements (CIV1') and (CIV2) are satisfied for $\mathcal{I}, \tilde{\mathcal{X}}, \mathcal{B}$ and Y in \mathcal{G}_M (see Definition 4). Then, the following three statements hold. (i) The total causal effect $[\beta, \alpha]$ of $[\mathcal{X}^\top, \mathcal{Z}^\top]^\top$ on Y satisfies the NIV moment equation*

$$\mathbb{E}[\text{cov}(Y - b\mathcal{X} - a\mathcal{Z}, \mathcal{I}|\mathcal{B})] = 0. \quad (8)$$

(ii) Further, if requirement (CIV3) is satisfied for $\mathcal{I}, \tilde{\mathcal{X}}, \mathcal{B}$, then $[\beta, \alpha]$ is the unique solution to Equation (8). (iii) If, additionally, $\mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{Z}$ and \mathbf{B} are observations of $\mathcal{X}, Y, \mathcal{I}, \mathcal{Z}$ and \mathcal{B} at T time points, W is a positive definite matrix, and

$$[\hat{b}, \hat{a}] := \arg \min_{b, a} \|\text{côv}(\mathbf{Y} - b\mathbf{X} - a\mathbf{Z}, \mathbf{I}|\mathbf{B})\|_W^2, \quad (9)$$

then \hat{b} is a consistent estimator for β .

We now apply the above result to VAR(1) processes satisfying (A2). In this case, the total causal effect coincides with the β defined in Section 2.1 (and (CIV1') and (CIV1) become equivalent). The generality of Theorem 5, however, can be used to develop similar results for VAR(p) processes with $p \neq 1$ and for total causal effects between arbitrary variables in the process.

4.2 Instrumental VAR(1) Processes

We now consider estimating the causal effect β in VAR(1)-processes, such as the one displayed in Figure 2a. A first attempt might be to estimate the effect β from X_{t-1} to Y_t by directly adapting the i.i.d. case and using $\text{IV}_{X_{t-1} \rightarrow Y_t}(I_{t-2})$. However, in general, this estimator is not consistent: In Proposition 17 in Appendix C.1, we show that the resulting estimator converges to $(1 - \alpha_{I, I} \alpha_{Y, Y})^{-1} \beta$. This naive IV approach fails due to the memory of the processes: The instrument, I_{t-2} , correlates with the response, Y_t , not only through the directed path $I_{t-2} \rightarrow X_{t-1} \rightarrow Y_t$, but also through paths reaching into the past, such as $I_{t-2} \leftarrow I_{t-3} \rightarrow X_{t-2} \rightarrow Y_{t-1} \rightarrow Y_t$. Instead, the concepts of NIV and time series reductions introduced above provide us with a principled approach to selecting which variables to include into the regression and allow us to construct estimators that adjust for the past.

4.2.1 BLOCKING THE PAST USING CONDITIONAL IV

Consider a VAR(1) process S satisfying (A2). By Theorem 5, any set satisfying requirements (CIV1) to (CIV3) in the corresponding marginalised time graph yields a consistent estimator for the causal




Figure 4: Three different marginalizations of the full time graph $\mathcal{G}_{\text{full}}$ of a VAR(1) process satisfying (A2) (see Figure 1). Using Theorem 5, we use these marginalizations for identification using CIV (left) and NIV (middle and right). (*Left*) Marginalization to nodes I_{t-2}, X_{t-1} and Y_t and their lagged values. (*Middle*) Marginalization to nodes I_{t-2}, X_{t-1} and Y_t and the lagged value of Y_t . (*Right*) Marginalization to m instrument nodes $I_{t-2}, \dots, I_{t-m-1}$, and X_{t-1}, Y_t , and Y_{t-1} .

effect of X_{t-1} on Y_t . Thus, define $M := \{I_{t-3}, I_{t-2}, X_{t-2}, X_{t-1}, Y_{t-1}, Y_t\}$ and consider the marginalization \mathcal{G}_M of the full time graph with respect to M , see Figure 4 (left). In this graph, every path from I_{t-2} to Y_t either goes through X_{t-1} or I_{t-3} . Indeed, the assumptions for Theorem 5 are satisfied when choosing $\mathcal{I}_t := \{I_{t-2}\}$ and either $\mathcal{B}_t = \{I_{t-3}\}$ or $\mathcal{B}_t := \{I_{t-3}, X_{t-2}, Y_{t-1}\}$. Formally, we have the following theorem.

Theorem 6 (Identification with conditioning set). *Consider a VAR(1) process $S = [S_t]_{t \in \mathbb{Z}}$ with $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2). Let either $\mathcal{B}_t := \{I_{t-3}\}$ or $\mathcal{B}_t := \{I_{t-3}, X_{t-2}, Y_{t-1}\}$. Then, the following three statements hold. (i) The causal effect β of X_{t-1} on Y_t satisfies the CIV moment condition $\mathbb{E}[\text{cov}(Y_t - \beta X_{t-1}, I_{t-2} | \mathcal{B}_t)] = 0$. (ii) Furthermore, if $\mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)]$ has rank d_X , then β is identified by $\text{CIV}_{X_{t-1} \rightarrow Y_t}(I_{t-2} | \mathcal{B}_t)$. (iii) If, additionally, $\mathbf{X}_t, \mathbf{Y}_t, \mathbf{I}_t$, and \mathbf{B}_t are observations of X, Y, I and \mathcal{B} at T time points, then β can be consistently estimated as $T \rightarrow \infty$ by $\text{CIV}_{\mathbf{X}_{t-1} \rightarrow \mathbf{Y}_t}(\mathbf{I}_{t-2} | \mathbf{B}_t)$, that is, the output of Algorithm 1.*

Theorem 6 establishes that the bias due to confounding from the past (see beginning of Section 4.2) can be overcome by choosing either $\mathcal{B}_t = \{I_{t-3}\}$ or $\mathcal{B}_t = \{I_{t-3}, X_{t-2}, Y_{t-1}\}$ as conditioning set. In Section 5, we compare these two choices empirically. The assumption that $\alpha_{X,I}$ has full rank ensures the relevance condition, requirement (CIV3). In Appendix B.2, we state the asymptotic distribution of an estimator $\hat{\beta}$ constructed using Theorem 6.

Theorem 6 is robust to misspecification in the number of lags: Though we assume (A2), which includes a VAR(1) assumption, the proof of Theorem 6 is still valid if higher-order lags (where each coefficient matrix has the sparsity pattern of Figure 2a) are present, as long as the I process is marginally VAR(1), i.e. $\alpha_{I,I}^j = 0$ for $j = 2, \dots, p$. In this case, the conditioning set \mathcal{B}_t still yields a valid instrument setting. Other choices of the marginalization set M are possible, too, and yield alternative ways of estimating the causal effect $X_{t-1} \rightarrow Y_t$: We now illustrate an alternative strategy for identification using nuisance IV.

4.2.2 BLOCKING THE PAST USING NUISANCE IV

The graph in Figure 4 (*middle*) shows the marginalization of the full time graph in Figure 2a to nodes $M := \{I_{t-2}, X_{t-1}, Y_{t-1}, Y_t\}$. The effect $X_{t-1} \rightarrow Y_t$ cannot be consistently estimated using only these variables in a CIV; if we condition on Y_{t-1} , for example, the path $I_{t-2} \leftrightarrow Y_{t-1} \leftrightarrow Y_t$ is

Algorithm 1 Estimating the causal effect β of X_{t-1} on Y_t in a VAR(1) process satisfying (A2).

Input: Sample $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_T] \in \mathbb{R}^{d_X \times T}$, $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_T] \in \mathbb{R}^{d_Y \times T}$, $\mathbf{I} = [\mathbf{I}_1, \dots, \mathbf{I}_T] \in \mathbb{R}^{d_I \times T}$

1: Align observations, by setting $\mathbf{Y} = [\mathbf{Y}_s, \dots, \mathbf{Y}_T]$ and either

$$(CIV) \quad \mathbf{I} := [\mathbf{I}_{s-2}, \dots, \mathbf{I}_{T-2}] \quad \mathbf{X} := [\mathbf{X}_{s-1}, \dots, \mathbf{X}_{T-1}] \quad \text{and} \quad \mathbf{B} := [\mathbf{I}_{s-3}, \dots, \mathbf{I}_{T-3}]$$

$$(NIV) \quad \mathbf{I} := \begin{bmatrix} \mathbf{I}_{s-2} & & \mathbf{I}_{T-2} \\ \vdots & \dots & \vdots \\ \mathbf{I}_{s-m+1} & & \mathbf{I}_{T-m+1} \end{bmatrix} \quad \text{and} \quad \mathbf{X} := \begin{bmatrix} \mathbf{X}_{s-1} & \dots & \mathbf{X}_{T-1} \\ \mathbf{Y}_{s-1} & \dots & \mathbf{Y}_{T-1} \end{bmatrix},$$

where s is chosen such that all indices are positive.

2: Compute regression estimates $\hat{E}[\mathbf{X}|\mathbf{B}]$, $\hat{E}[\mathbf{Y}|\mathbf{B}]$, and $\hat{E}[\mathbf{I}|\mathbf{B}]$.

3: Compute residual processes $r_{\mathbf{X}} := \mathbf{X} - \hat{E}[\mathbf{X}|\mathbf{B}]$, $r_{\mathbf{Y}} := \mathbf{Y} - \hat{E}[\mathbf{Y}|\mathbf{B}]$, and $r_{\mathbf{I}} := \mathbf{I} - \hat{E}[\mathbf{I}|\mathbf{B}]$.

$$4: W := \left(\frac{1}{T-s+1} r_{\mathbf{I}} r_{\mathbf{I}}^\top \right)^{-1}$$

$$5: \hat{\beta} := r_{\mathbf{Y}} r_{\mathbf{I}}^\top W r_{\mathbf{I}} r_{\mathbf{X}} \left(r_{\mathbf{X}} r_{\mathbf{I}}^\top W r_{\mathbf{I}} r_{\mathbf{X}}^\top \right)^{-1}$$

6: If \mathbf{I} and \mathbf{X} are chosen according to NIV, set $\hat{\beta} := \hat{\beta}_{1:d_X}$.

Output: Estimate of causal effect $\hat{\beta}$

unblocked, violating requirement (CIV1). Instead, we can include Y_{t-1} as a nuisance regressor: We identify the effect of X_{t-1} on Y_t using the instrument $\mathcal{I}'_t := \{I_{t-2}\}$ by NIV $_{X_{t-1} \rightarrow Y_t}(\mathcal{I}'_t, Y_{t-1})$, defined in Section 3.2. This model satisfies the (nuisance) IV requirement (CIV1), because the only open paths between I_{t-2} and Y_t include either the edge $X_{t-1} \rightarrow Y_t$ or the edge $Y_{t-1} \rightarrow Y_t$.

Because the resulting d_X -dimensional estimate is extracted from the $d_X + d_Y$ -dimensional solution IV $_{\{X_{t-1}, Y_{t-1}\} \rightarrow Y_t}(\mathcal{I}'_t)$, we require that $\text{rank } \mathbb{E} \{ [X_{t-1}^\top, Y_{t-1}^\top]^\top (\mathcal{I}'_t)^\top \} = d_X + d_Y$. If $d_I < d_X + d_Y$, this rank condition is not met for $\mathcal{I}'_t := \{I_{t-2}\}$. To overcome this, one can increase the instrument set to $\mathcal{I}_t = \{I_{t-2}, I_{t-3}, \dots, I_{t-m-1}\}$ for a user-specified m : Figure 4 (right) shows the marginalization of the full time graph to $M := \mathcal{I}_t \cup \{X_{t-1}, Y_{t-1}, Y_t\}$. Again, requirements (CIV1) and (CIV2) are satisfied in \mathcal{G}_M , but the instrument set now has dimension $|\mathcal{I}| = md_I$ (implying that we should choose $m \geq (d_X + d_Y)/d_I$) and, provided the rank condition now holds, β is identified by NIV $_{X_{t-1} \rightarrow Y_t}(\mathcal{I}_t, Y_{t-1})$. The following theorem formalizes this discussion.

Theorem 7 (Identification with nuisance regressor). *Consider a VAR(1) process $S = [S_t]_{t \in \mathbb{Z}}$ with $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2). Let $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$ for an $m \geq 1$ and $\mathcal{Z}_t := \{Y_{t-1}\}$. Then, the following three statements hold. (i) There exists $\alpha \in \mathbb{R}$ such that the causal effect β of X_{t-1} on Y_t satisfies the NIV moment condition $\text{cov}(Y_t - \beta X_{t-1} - \alpha \mathcal{Z}_t, \mathcal{I}_t) = 0$. (ii) Further, if $\mathbb{E}[[X_{t-1}^\top, \mathcal{Z}_t^\top]^\top \mathcal{I}_t^\top]$ has rank $d_X + d_Y$, β is identified by NIV $_{X_{t-1} \rightarrow Y_t}(\mathcal{I}_t, \mathcal{Z}_t)$. (iii) If, additionally, $\mathbf{X}_t, \mathbf{Y}_t, \mathbf{I}_t$, and \mathbf{Z}_t are observations of X, Y, \mathcal{I} and \mathcal{Z} at T time points, then β can be consistently estimated as $T \rightarrow \infty$ by NIV $_{\mathbf{X}_{t-1} \rightarrow \mathbf{Y}_t}(\mathbf{I}_t, \mathbf{Z}_t)$, that is, the output of Algorithm 1.*

In Appendix B.2, we state the asymptotic distribution of an estimator $\hat{\beta}$ constructed using Theorem 7. Similar to the discussion in Section 4.2.1, Theorem 7 is robust to misspecification of the number of lags. While Theorem 6 is still valid in VAR(p) processes if the I process is a VAR(1) process, Theorem 7 is valid in VAR(p) processes if the Y process is VAR(1), since in that case Y_{t-1} yields a valid nuisance regressor.

Theorem 7 shows that identification is possible if $\text{rank } \mathbb{E} \{ [X_{t-1}^\top, Y_{t-1}^\top]^\top \mathcal{I}_t^\top \} = d_X + d_Y$ is met, where we use the lagged instrument set $\mathcal{I}_t = \{I_{t-2}, \dots, I_{t-m-1}\}$ (it is easy to see that for this choice

of \mathcal{I}_t , identifiability also holds for $\beta = 0$ and the rank being d_X). Satisfying this relevance criterion implies requirement (CIV3). For the CIV approach in Section 4.2.1, this is directly related to the rank of the parameter α_{XI} . For NIV, the relevance criterion depends on the parameter matrix in an intricate way. We now provide necessary and sufficient conditions for when this rank condition is satisfied. Moreover, we show in Corollary 9 below that for almost all parameter matrices one can obtain sufficiently high rank to identify the effect $X_{t-1} \rightarrow Y_t$ by increasing the number of lags used.

Define the following submatrices of the parameter matrix A from (A2).

$$A_I := \begin{pmatrix} \alpha_{X,I} \\ 0 \end{pmatrix} \in \mathbb{R}^{d_X+1} \quad \text{and} \quad A_{XY} := \begin{pmatrix} \alpha_{X,X} & \alpha_{X,Y} \\ \beta & \alpha_{Y,Y} \end{pmatrix} \in \mathbb{R}^{(d_X+1) \times (d_X+1)}. \quad (10)$$

The following theorem outlines conditions for $\mathbb{E}[[X_{t-1}^\top, Y_{t-1}^\top]^\top \mathcal{I}_t^\top]$ to have full rank when $d_I = 1$ and we use $d_X + d_Y = d_X + 1$ lags as instruments:

Theorem 8. *Consider a process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}$ satisfying (A2). Assume that $d_I = d_Y = 1$ and let $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$, where $m = d_X + d_Y$. Let A_{XY} and A_I be defined as in Equation (10). The following three statements are equivalent:*

- i) $\text{rank } \mathbb{E}[[X_{t-1}^\top, Y_{t-1}^\top]^\top \mathcal{I}_t^\top] = d_X + d_Y$.
- ii) The matrix $[A_{XY}^0 A_I, A_{XY}^1 A_I, \dots, A_{XY}^{d_X} A_I]$ is invertible, where A_{XY}^0 is the identity matrix of size $(d_X + d_Y) \times (d_X + d_Y)$.
- iii) Different Jordan blocks of J have different eigenvalues and for all $q \in \{1, \dots, k\}$, the coefficient $w_{\sum_{i=1}^q m_i}$ is non-zero; here, $J = Q^{-1} A_{XY} Q$ is the Jordan normal form⁹ of A_{XY} , with k Jordan blocks $J = \text{diag}(J_{m_1}(\lambda_1), \dots, J_{m_k}(\lambda_k))$, each with size m_i and eigenvalue λ_i , and w are the coefficients of A_I in the basis of the generalized eigenvectors Q , that is, $w = Q^{-1} A_I$.

Intuitively, Theorem 8 states that identification of β is possible if the information is passed on in a sufficiently diverse way from I_{t-k} to (X_t, Y_t) . For every $k \in \{0, \dots, d_X\}$, $A_{XY}^k A_I$ corresponds to the path $I_{t-(k+1)} \xrightarrow{A_I} (X_{t-k}, Y_{t-k}) \xrightarrow{A_{XY}} \dots \xrightarrow{A_{XY}} (X_t, Y_t)$ in Figure 1. Theorem 8 ii) requires these to be sufficiently different (that is, linearly independent), for the matrix $\mathbb{E}[[X_{t-1}^\top, Y_{t-1}^\top]^\top \mathcal{I}_t^\top]$ to have full rank. While there are parameter matrices that do not satisfy Theorem 8 (see Appendix C.2 for examples), the following corollary shows that, when chosen randomly, almost all parameter matrices A allow for using multiple lags I_{t-2-j} as instruments.

Corollary 9. *Consider a VAR(1) process S with $d_I = 1$ and parameter matrix A , and assume that sparsity pattern of A is given by (A2) and that the non-zero entries of A are drawn from any distribution which has density with respect to Lebesgue measure. Then β is identifiable with probability 1.*

The following corollary provides a sufficient condition for identifiability of β , when instruments are multivariate.

Corollary 10. *Consider a VAR(1) process S satisfying (A2) with $d_I > 1$ instrument processes $I^{(1)}, \dots, I^{(d_I)}$. Assume that there is at least one instrument process $I^{(j)}$ such that both of the following conditions hold.*

- i) $I_t^{(j)}$ is independent of $I_s^{(i)}$ for all t, s and $i \neq j$, and
- ii) the requirements of Theorem 8 are satisfied for the reduced process $(I^{(j)}, X, Y)$.

Then β is identifiable.

9. See Appendix E.9 for the definition of Jordan normal forms and the notation that we use.

Using a single instrument at $d_X + 1$ lags allows for a simple condition for identifiability. But in finite samples, using instruments with high time lags may come at a loss of efficiency as the estimation procedure may suffer from weak dependencies between the residual and the instrument due to the mixing of the time series, see Section 5 for an empirical investigation.

4.2.3 EXTENSION TO NON-VAR NONLINEAR PROCESSES

The results in Sections 4.2.1 and 4.2.2 assume that the entire process S is a VAR(1) process, see (A2), which was done mostly for presentation purposes. The key arguments and statements hold more generally. We have argued that similar arguments hold for VAR(p) processes and we now consider a setting, where Y_t satisfies a linear structural equation (as in a VAR(1) process) but we do not assume any VAR(p) structure or linearity on the remaining subprocesses. We outline the assumptions needed to obtain the same identification results as in Sections 4.2.1 and 4.2.2 (a similar relaxation can be applied when Y_t behaves like a VAR(p) process). Let $[S_t]_{t \in \mathbb{Z}} = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ and assume that for all $t \in \mathbb{Z}$

$$Y_t := \beta X_{t-1} + \alpha_{Y,Y} Y_{t-1} + g(\varepsilon_t^Y, H_{t-1}), \quad (11)$$

where ε_t^Y is a sequence of i.i.d. random variables that, for all t , are independent of $[Y_s]_{s < t}$ and $[X_s, I_s, H_s]_{s \leq t}$ and g is a measurable function. Without assuming a VAR(p) process we make the following assumptions on S .

- (A1') $[S_t]_{t \in \mathbb{Z}}$ is covariance stationary.
- (A2') S satisfies Equation (2), that is, empirical first and second moments converge to their population version.
- (A3') There exists a $p \in \mathbb{N}$ such that for all $i \in \{1, \dots, d\}$, there exists a function f^i such that for all $t \in \mathbb{Z}$

$$S_t^i = f^i(S_{t-1}, \dots, S_{t-p}) + \varepsilon_t^i.$$

This induces a full time graph $\mathcal{G}_{\text{full}}$ (Peters et al., 2013); we assume that for all finite disjoint collections of nodes A, B, C from $\mathcal{G}_{\text{full}}$ such that A and C are d -separated given B in $\mathcal{G}_{\text{full}}$, we have $A \perp\!\!\!\perp C | B$. Furthermore, for all $t \in \mathbb{Z}$, $H_{t-1} \in \text{ND}(I_{t-2})$, $Y_{t-1} \in \text{ND}(I_{t-2})$ and ε_Y^t is independent of any finite set $A \subseteq \text{ND}(Y_t)$.

- (A4') For all $t \in \mathbb{Z}$ and $m \in \mathbb{N}$, we have $(\varepsilon_t^Y, H_{t-1}) \perp\!\!\!\perp (I_{t-2}, \dots, I_{t-2-m})$.

Using these assumptions, we can restate Theorem 6 without the VAR(1) assumption.

Proposition 11 (Identification with conditioning set relaxing the VAR(p) assumption). *Consider a process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying Equation (11) and (A1')–(A3'). Let \mathcal{B}_t be a set of variables satisfying $\text{PA}(I_{t-2}) \subseteq \mathcal{B}_t \subseteq \text{ND}(Y_t) \cap \text{ND}(I_{t-2})$ in $\mathcal{G}_{\text{full}}$. Then, (i), (ii) and (iii) from Theorem 6 hold.*

Similarly, we can extend Theorem 7 to more general time series models, too.

Proposition 12 (Identification with nuisance regressor relaxing the VAR(p) assumption). *Consider a process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying Equation (11) and (A1'), (A2'), and (A4'). Let $\mathcal{Z}_t := \{Y_{t-1}\}$ and $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$ for an $m \geq 1$. Then, (i), (ii), and (iii) from Theorem 7 hold.*

Equation (11) allows for binary covariates or instruments. The assumed model in Equation (11) does not allow for the response Y_t to be binary, though Propositions 11 and 12 readily extend to other parameterizations of the parameter of interest. Another generalization that is possible, is to

let β_t change over time. We show in Appendix C.4 that one can also develop estimating equations for this setting.

So far, we have focused on estimating causal effects. Knowledge of such causal effects can be of interest in itself. In the following section, we discuss that they can also be used for prediction and forecasting under interventions.

4.3 Optimal Prediction under Interventions

Causal estimates may facilitate improved prediction under intervention. Suppose that we have inferred the causal effect $X_{t-1} \rightarrow Y_t$ for instance using the methods presented in Section 4.2 above. How do we best predict Y_{t+1} given that we perform the intervention $\text{do}(X_t := x)$ (see Appendix A.3 for an introduction to do-interventions) and that we observe the past values of the time series (but the value of X_t had it not been intervened on is not observed)?

Due to the hidden confounding, the conditional mean of Y_{t+1} given its past (which could be consistently estimated by OLS regression, for example) is in general not optimal in terms of mean squared prediction error (MSPE). Intuitively, this is because the conditional mean also encompasses the effects of the latent process H onto Y_{t+1} . In the i.i.d. setting, it has been observed that using the causal parameter yields a smaller MSPE and can be worst-case optimal under arbitrarily large interventions (e.g., Rojas-Carulla et al., 2018; Christiansen et al., 2022).

In VAR(p) processes, the intervention $\text{do}(X_t := x)$ partially breaks the confounding of X_t and Y_{t+1} from the past. If an unobserved process H confounds X and Y , the marginalized process (X, Y) is not a Markov process (neither with or without intervention), and therefore the lagged observations $\{X_{t-k}, k = 1, \dots, m\}$ and $\{Y_{t-j}, j = 1, \dots, l\}$ can further improve prediction of Y_{t+1} , due to being informative of H , and therefore of Y_{t+1} . Fixing the number of lags, the following proposition shows that, under the intervention $\text{do}(X_t := x)$, the optimal linear prediction consist of a mix of (population) regression parameters for non-intervened variables and causal parameters for the intervened variable X_t .

Proposition 13. *Consider a VAR(1) process $S = [S_t]_{t \in \mathbb{Z}}$ satisfying (A2). Let β be the causal effect from X_t to Y_{t+1} , and let, for an arbitrary $m, \ell \in \mathbb{N}$, $(\alpha_{Y,X}, \alpha_{Y,Y})$ be the population vector of coefficients when regressing $Y_{s+1} - \beta X_s$ on $\{X_{s-k}, k = 1, \dots, m\} \cup \{Y_{s-j}, j = 0, \dots, l\}$. Then*

$$(\alpha_{Y,Y}, \beta, \alpha_{Y,X}) = \arg \min_{a,b,c} \mathbb{E}_{\text{do}(X_t := x)} \left\{ Y_{t+1} - \sum_{j=0}^{\ell} a_j Y_{t-j} - b X_t - \sum_{k=1}^m c_k X_{t-k} \right\}^2.$$

That is, under the intervention $\text{do}(X_t := x)$, the causal coefficient can be used to optimally predict Y_{t+1} . We state the corresponding finite sample algorithm in Appendix C.3.

5. Simulation Experiments

We test the empirical performance of our proposed estimators in simulation experiments. (Appendices D.2 and D.3 discuss additional experiments on settings with weak instruments.)

Data generating process. We generate data by first simulating a matrix A with the sparsity structure of Figure 2a and all non-zero entries being drawn independently uniformly at random from $(-0.9, -0.1) \cup (0.1, 0.9)$. By Corollary 9, the causal effect under this sampling scheme is almost surely identifiable by NIV. Unless specified differently, we use $d_I = 3, d_X = 2$ and $d_Y = 1$. Any such randomly generated matrix is used to generate data only if it satisfies the eigenvalue condition in (A1) with a margin of 0.1. Furthermore, all noise variables ε_t^i are independently randomly drawn from a normal distribution with mean 0 and standard deviation 1.

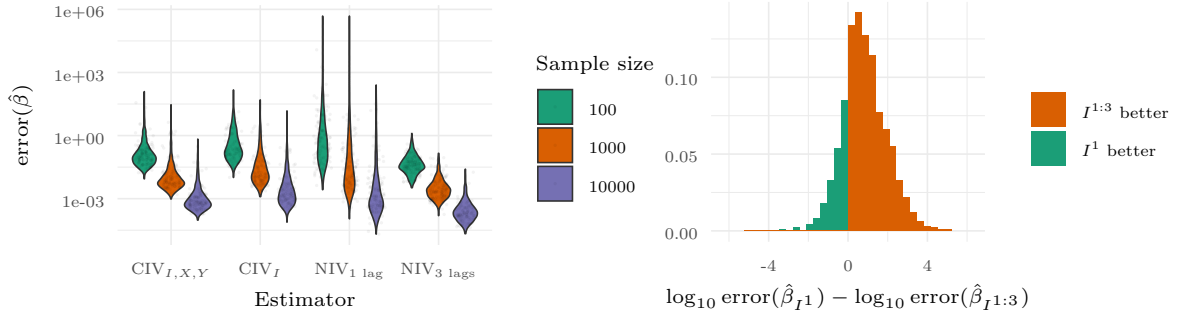


Figure 5: (Left) Distributions of the average error (in log scale) of different consistent CIV and NIV estimators for various sample sizes T , see Section 5.1 (‘Consistency’). Each point corresponds to the average over repeated draws from the same parameter matrix, and different points correspond to different parameter matrices. The $\text{NIV}_{1 \text{ lag}}$ estimators have heavier tails, corresponding to the fact that this is a just-identified case, where we use a 3-dimensional instrument to estimate a 3-dimensional causal effect. (Right) Histogram of log error ratios for two different NIV estimators in a model with a 3-dimensional instrument. I^1 uses 6 lags from a 1-dimensional instrument process, while $I^{1:3}$ uses 2 lags of a 3-dimensional instrument process, see Section 5.1 (‘Using more lags or additional instruments’). A value larger than zero indicates that I^1 yields a larger error than $I^{1:3}$. This is the case for the majority of the considered settings.

Evaluation of the estimators. We simulate $m = 1,000$ random matrices and for each we simulate $s = 10$ data sets. We fit estimators $\hat{\beta}_j$ for each $j \in \{1, \dots, s\}$, and for a given matrix, we compute the average error, $\text{error}(\hat{\beta}) := \text{mean}(\|\hat{\beta}_1 - \beta\|_2^2, \dots, \|\hat{\beta}_s - \beta\|_2^2)$.

5.1 Identification of the causal effect

Consistency The estimators CIV and NIV are consistent (Theorems 6 and 7) and we perform an experiment to compare their finite sample properties for different sample sizes. We simulate data from the scheme described above and fit two $\text{CIV}_{X_{t-1} \rightarrow Y_t}(I_{t-2} | \mathcal{B}_t)$ estimators, where $\mathcal{B}_t = \{I_{t-3}\}$ (CIV_I) or $\mathcal{B}_t = \{I_{t-3}, X_{t-2}, Y_{t-1}\}$ ($\text{CIV}_{I,X,Y}$) and two $\text{NIV}_{X_{t-1} \rightarrow Y_t}(\mathcal{I}_t, Y_{t-1})$ estimators, where $\mathcal{I}_t = \{I_{t-2}\}$ ($\text{NIV}_{1 \text{ lag}}$) or $\mathcal{I}_t = \{I_{t-2}, I_{t-3}, I_{t-4}\}$ ($\text{NIV}_{3 \text{ lag}}$). All of these estimators are consistent (see Section 4.2). We plot the errors obtained for different sample sizes T in Figure 5 (left). For all estimators, the errors decrease with increasing sample size supporting the consistency results in Theorem 5. In general, there is no empirical support for either of the NIV or CIV estimators to be strictly better than the others in terms of speed of convergence in terms of sample size. As discussed in Section 4.2.1, both $\mathcal{B}_t = \{I_{t-3}\}$ and $\mathcal{B}_t = \{I_{t-3}, X_{t-2}, Y_{t-1}\}$ block confounding from past values. We observe that removing X_{t-2} and Y_{t-1} from \mathcal{B}_t increases the upper tail of the error distribution, supporting the intuition that while the conditioning variables X_{t-2} and Y_{t-1} are not necessary for identification, they reduce finite sample variance. Similarly, for the NIV estimators, using 3 lags instead of a single lag shrinks the upper tail of the error distribution. In Appendix D.1 we print the correlation matrices between the error of the estimators, and in Appendix D.2 we conduct a similar experiment, but with fixed data generating mechanisms.

Using more lags or additional instruments. We compare using multiple lags of a single instrument to using multiple instrument processes. We consider the model described above with $d_I = 3$ independent instrument processes, and use either the first instrument process (I^1) or all three instrument processes ($I^{1:3}$) for estimation in NIV. For $I^{1:3}$ we use 2 lags of each of the three




Figure 6: Median error (log scale) for the NIV estimator as we vary Δ , the difference between the eigenvalues of $\alpha_{X,X}$, see Section 5.1 (‘Estimation close to non-identifiability’). The causal effect is identifiable if and only if $\Delta \neq 0$, see Theorem 8. Indeed, the error does not decrease for $\Delta = 0$ and decreases the faster, the further Δ is away from 0.

processes (‘recent instruments’), while for I^1 we use 6 lags (‘distant instruments’), such that both models use in total 6 instruments. In this way, we can inspect the benefit of using more recent instruments if available.

Figure 5 (right) shows a histogram of the log error ratio of the two different estimators: A large value indicates that model I^1 , which uses many lags of only a single instrument process, incurs a higher error. In the majority of parameter settings, using recent instruments yields a lower error, in some cases several orders of magnitude, when compared to the more distant instruments. Although adding lags of a univariate instrument process can yield identifiability of the causal effect of a regressor process that is not univariate (see Theorem 8), this simulation experiment supports the notion that using more instruments (if available) is preferred over using more distant lags.

Estimation close to non-identifiability. Example 1 in Appendix C.2 shows a setting, in which β is not identifiable by NIV (in the setting of Corollary 9, this happens with probability zero). We examine the behaviour of the NIV estimator in scenarios that are close to this non-identifiable setting. We consider $d_I = 1$, $\alpha_{X,X} = \text{diag}(-0.6, -0.6 + \Delta)$ and draw the remaining parameters uniformly from $(-0.9, -0.1) \cup (0.1, 0.9)$. As per Corollary 9, the causal effect is identifiable, except for $\Delta = 0$, in which case $\alpha_{X,X}$ has two Jordan blocks with the same eigenvalue. In Figure 6 we plot the median error for changing Δ and sample size T .¹⁰ The further Δ is from 0, the faster (in terms of sample size) the estimator converges to β . In the non-identified setting $\Delta = 0$, we do not observe the error to decrease with increasing sample size. This observation is in line with Theorem 8.

5.2 Using the causal parameter for prediction under intervention on X

We support empirically that a linear prediction using OLS parameters for non-intervened variables and causal parameters for the intervened variables achieves minimal square loss for prediction under intervention (see Proposition 13). We consider the model with $d_X = d_I = 1$ and, ensuring strong hidden confounding, draw the non-zero entries $\alpha_{i,H}$, $i \in \{X, H, Y\}$ uniformly at random from $(-0.9, -0.5) \cup (0.5, 0.9)$ (instead of $(-0.9, -0.1) \cup (0.1, 0.9)$ as for the other non-zero entries of A). The prediction task follows Section 4.3: Given observations $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_{T-1}] \in \mathbb{R}^{d_X \times T-1}$ and $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_T] \in \mathbb{R}^{d_Y \times T}$, with $T = 3,000$, the goal is to predict $\mathbf{Y}_{T+1} \in \mathbb{R}^{d_Y \times 1}$ under an intervention $\text{do}(X_T := n \cdot \sigma)$ where σ is the standard deviation of the process $[X_t]_{t=1}^{T-1}$ and $n \in \{1, 5\}$. In Section 4.3 we discuss that IV is prediction optimal under arbitrary interventions $\text{do}(X_T := x)$.

10. Here, we report the median, since the non-identifiability when $\Delta = 0$ implies that the mean is ill-behaved; for those lines where $\Delta \neq 0$, plotting the mean instead of the median yields a similar plot (not shown).




Figure 7: The figure plots the loss $\text{MSPE}(\hat{Y}_{T+1})$ under an intervention $\text{do}(X_T := x)$ when Y_{T+1} is predicted using OLS against the loss when Y_{T+1} is predicted using one of the IV methods we develop for time series. The OLS prediction is based on the regression $Y_{t+1} \sim X_t + X_{t-1} + X_{t-2} + Y_t + Y_{t-1}$ while the IV predictions of Y_{T+1} are based on the procedure discussed in Section 4.3 (see Algorithm 2 in Appendix C.3). We plot this both for $\text{do}(X_T := \sigma)$ and $\text{do}(X_T := 5\sigma)$, where σ is the standard deviation of X_t in the unintervened distribution. The results show that as the intervention strength increases, the OLS prediction error increases at a faster rate than the IV prediction error. 17 outliers were removed from the right-hand side plot, of which 10 had a larger OLS MSPE than IV MSPE.

In particular, one would expect that OLS becomes increasingly inferior to IV methods when one increases n and thereby the intervention strength $\text{do}(X_T := n \cdot \sigma)$.

We compare prediction for Y_{T+1} via Algorithm 2 in Appendix C.3 (with $m = 2$ and $l = 1$) with prediction based on the OLS regression $Y_{t+1} \sim X_t + X_{t-1} + X_{t-2} + Y_t + Y_{t-1}$ (‘OLS’). For each repetition and matrix A , we obtain CIV and NIV estimates of β on a separate sample first, and then obtain predictions for Y_{T+1} following Algorithm 2 using either $\hat{\beta}_{\text{CIV}}$ (‘CIV’) or $\hat{\beta}_{\text{NIV}}$ (‘NIV’). For each of the $m = 100$ random matrices A we compute $\text{error}(\hat{Y}_{T+1})$ as the mean of the squared prediction error $(\hat{Y}_{T+1} - Y_{T+1})^2$ over the $s = 100$ repetitions.

In Figure 7, we plot the error of the OLS estimate against the error of the IV estimate (either CIV or NIV). When $(X_T := \sigma)$, the intervention is within the normal range of X_t , and the errors of OLS and IV estimates are similar. As we perform a stronger intervention $(X_T := 5\sigma)$, the OLS error exceeds the IV error in most simulations. This robustness in prediction under intervention is in line with the result in Proposition 13.

6. Conclusion and Future Work

In this work, we have developed IV methods for time series data that allow us to identify the causal effect of a process X on a response process Y , based on an instrument process I that exhibits memory effects. Simple adaptations of ordinary IV estimators generally fail to identify the causal effect due to confounding from the past, as we show in Proposition 17. We have developed the concept of nuisance IV (NIV), see Theorem 2, a marginalization framework for time series graphs, see Theorem 1, and have proved a version of the global Markov property for $\text{VAR}(p)$ processes, see Theorem 5. Based on these principles, we propose two classes of estimation methods that properly adjust for confounding from the past: one based on choosing the correct conditioning set (CIV), see Theorem 6, and another

one based on nuisance regressors (NIV), see Theorem 7. The procedures find solutions to moment conditions that, in their population version, are satisfied for the true causal parameters. Unlike in the i.i.d. case, the identifiability conditions (which are usually rank conditions) do not have a simple interpretation. Theorem 8 provides necessary and sufficient conditions on the parameters of the underlying data-generating process for the causal parameter being the unique solution to the corresponding moment equation.

The experiments show that while several different choices of instruments, conditioning sets and nuisance regressors allow us to consistently identify the causal effect, the modelling choices impact the finite sample performance. For example, for identifiability in the case of NIV, we only need the number of lags used as instruments, m , to be so large that $d_I \cdot m = d_X + d_Y$, but using more lags may help in finite samples in that this shrinks the upper tail of the error distribution, see Figure 5 (left).

We have further argued that identifying the causal effect may be of interest not only for causal inference, but also for prediction of Y under the intervention $\text{do}(X_t := x)$, where the minimal expected squared error can be obtained by a mix of causal parameters and regression coefficients, see Proposition 13 and Section 5.2.

For future work, it may be fruitful to develop principled techniques for deciding which estimator yields the best finite sample performance (see, e.g., Henckel, 2021, Chapter 4) and to construct confidence statements, either based on Appendix B.2 or other techniques (e.g., Newey and West, 1987; Shah and Peters, 2020). In settings where both NIV and CIV estimators are valid, it would be interesting to combine the estimating equations to obtain a single estimator and to analyze both the statistical and identifiability properties of such an approach. Our results assume that the response variable Y_t is continuous. However, we think one can generalize the NIV and CIV procedures, for example to a setting where Y_t follows a generalized linear model. Further, it would be interesting to understand whether the identifiability obtained by using lagged instruments in Theorem 8 can be achieved in processes that do not satisfy the linearity assumption. In addition, it is known that overidentification can help to detect or even correct for certain types of model violations (see, e.g., Hansen, 1982; Chao et al., 2014). It may be worthwhile to explore to which extent such methods can also be used in the time series setting – in particular, as the NIV approach comes with a natural way of obtaining a certain degree of overidentification by including more of the previous time points (see also the discussion in Section 5.1). It could further be interesting to extend the results to VARMA processes (e.g., Scherrer and Deistler, 2019). Finally, as for the i.i.d. case (Imbens and Newey, 2009; Chesher, 2003; Saengkyongam et al., 2022), considering independence, rather than vanishing covariances may yield stronger identifiability results but may come with computational and statistical challenges.

Acknowledgments

We thank Aurélien Bibaut for pointing out typos in an earlier version of Figure 4 and Ignacio Gonzalez Pérez for discovering and discussing potential fixes for errors in the proofs of Theorems 1 and 5. During parts of this project, NT, RS and JP were supported by a research grant (18968) from VILLUM FONDEN and JP and SW by the Carlsberg Foundation.

References

- T. W. Anderson and H. Rubin. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics*, 20(1):46–63, 1949.
- J. D. Angrist and G. W. Imbens. Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association*, 90

- (430):431–442, 1995.
- J. D. Angrist and A. B. Krueger. Instrumental variables and the search for identification: From supply and demand to natural experiments. *Journal of Economic Perspectives*, 15(4):69–85, 2001.
- J. D. Angrist, G. W. Imbens, and D. B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- S. Bongers, P. Forré, J. Peters, and J. M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885 – 2915, 2021.
- R. J. Bowden and D. A. Turkington. *Instrumental Variables*, volume 8 of *Econometric Society Monographs*. Cambridge University Press, 1985.
- C. Brito and J. Pearl. Generalized instrumental variables. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 85–93, 2002.
- P. J. Brockwell and R. A. Davis. *Time Series: Theory and Methods*. Springer, 2nd edition, 1991.
- J. C. Chao, J. A. Hausman, W. K. Newey, N. R. Swanson, and T. Woutersen. Testing overidentifying restrictions with many instruments and heteroskedasticity. *Journal of econometrics*, 178:15–21, 2014.
- A. Chesher. Identification in nonseparable models. *Econometrica*, 71(5):1405–1441, 2003.
- R. Christiansen, N. Pfister, M. E. Jakobsen, N. Gnecco, and J. Peters. A causal framework for distribution generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6614–6630, 2022.
- T. Chu and C. Glymour. Search for additive nonlinear time series causal models. *Journal of Machine Learning Research*, 9:967–991, 2008.
- R. Dahlhaus and M. Eichler. Causality and graphical models in time series analysis. In P. Green, N. Hjort, and S. Richardson, editors, *Highly structured stochastic systems*, pages 115–137. Oxford University Press, 2003.
- S. Demiralp and K. Hoover. Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and Statistics*, 65:745–767, 2003.
- V. Didelez, S. Meng, and N. A. Sheehan. Assumptions of IV methods for observational epidemiology. *Statistical Science*, 25(1):22–40, 2010.
- M. Eichler. Granger causality and path diagrams for multivariate time series. *Journal of Econometrics*, 137(2):334–353, 2007.
- M. Eichler. Graphical modelling of multivariate time series. *Probability Theory and Related Fields*, 153:233–268, 2012.
- M. Eichler and V. Didelez. On Granger causality and the effect of interventions in time series. *Lifetime Data Analysis*, 16(1):3–32, 2010.
- D. Entner and P. O. Hoyer. On causal discovery from time series data using fci. In *Proceedings of the 5th European Workshop on Probabilistic Graphical Models*, 2010.
- R. C. Fair. The estimation of simultaneous equation models with lagged endogenous variables and first order serially correlated errors. *Econometrica: Journal of the Econometric Society*, pages 507–516, 1970.

- C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–38, 1969.
- C. W. J. Granger. Testing for causality: A personal viewpoint. *Journal of Economic Dynamics and Control*, 2(1):329–352, 1980.
- A. R. Hall. *Generalized method of moments*. Oxford University Press, 2005.
- J. D. Hamilton. *Time Series Analysis*. Princeton University Press, 1st edition, 1994.
- L. P. Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054, 1982.
- L. Henckel. *Graphical Tools for Efficient Causal Effect Estimation*. PhD thesis, ETH Zurich, Zurich, 2021.
- M. A. Hernán and J. M. Robins. Instruments for causal inference: an epidemiologist’s dream? *Epidemiology*, pages 360–372, 2006.
- L. Hirth, T. Khanna, and O. Ruhnau. The (very) short-term price elasticity of German electricity demand, 2022. URL <http://hdl.handle.net/10419/249570>.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- G. W. Imbens and W. K. Newey. Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica*, 77(5):1481–1512, 2009.
- M. E. Jakobsen and J. Peters. Distributional robustness of K-class estimators and the PULSE. *The Econometrics Journal*, 25(2):404–432, 2022.
- H. Kang, A. Zhang, T. T. Cai, and D. S. Small. Instrumental variables estimation with some invalid instruments and its application to mendelian randomization. *Journal of the American Statistical Association*, 111(513):132–144, 2016.
- D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT Press, 2009.
- S. L. Lauritzen. *Graphical models*. Clarendon Press, 1996.
- S. L. Lauritzen, A. P. Dawid, B. N. Larsen, and H.-G. Leimer. Independence properties of directed Markov fields. *Networks*, 20(5):491–505, 1990.
- H. Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer Berlin Heidelberg, 2005.
- N. Meinshausen. Causality from a distributional robustness point of view. In *2018 IEEE Data Science Workshop (DSW)*, pages 6–10. IEEE, 2018.
- H. Michael, Y. Cui, S. A. Lorch, and E. J. Tchetgen Tchetgen. Instrumental variable estimation of marginal structural mean models for time-varying treatment. *Journal of the American Statistical Association*, 119(546):1240–1251, 2024.
- S. W. Mogensen. Instrumental processes using integrated covariances. In *Proceedings of the 2nd Conference on Causal Learning and Reasoning*, volume 213. Proceedings of Machine Learning Research, 2023.
- A. Moneta, N. Chlass, D. Entner, and P. Hoyer. Causal search in structural vector autoregressive models. In F. Popescu and I. Guyon, editors, *Proceedings of the Neural Information Processing Systems Mini-Symposium on Causality in Time Series*, volume 12 of *Proceedings of Machine Learning Research*, pages 95–114. PMLR, 2011.

- I. Neamtu. *Wind power effects and price elasticity of demand for the Nordic Electricity Markets*. PhD thesis, PhD dissertation, Aarhus BSS, Aarhus University, 2016.
- W. K. Newey and K. D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation. *Econometrica*, 55(3):703–708, 1987.
- J. Pearl. *Causality*. Cambridge University Press, 2009.
- J. Peters, D. Janzing, and B. Schölkopf. Causal inference on time series using restricted structural equation models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 26, pages 154–162, 2013.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT Press, 2017.
- J. Peters, S. Bauer, and N. Pfister. Causal models for dynamical systems. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 671–690. ACM, 2022.
- O. Reiersøl. *Confluence analysis by means of instrumental sets of variables*. PhD thesis, Almqvist & Wiksell, 1945.
- T. Richardson. Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1):145–157, 2003.
- J. M. Robins. Causal inference from complex longitudinal data. In *Latent variable modeling and applications to causality*, pages 69–117. Springer, 1997.
- M. Rojas-Carulla, B. Schölkopf, R. Turner, and J. Peters. Causal transfer in machine learning. *Journal of Machine Learning Research*, 19(36):1–34, 2018.
- T. J. Rothenberg. Identification in parametric models. *Econometrica: Journal of the Econometric Society*, pages 577–591, 1971.
- D. Rothenhäusler, N. Meinshausen, P. Bühlmann, and J. Peters. Anchor regression: Heterogeneous data meet causality. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(2):215–246, 2021.
- S. Saengkyongam, L. Henckel, N. Pfister, and J. Peters. Exploiting independent instruments: Identification and distribution generalization. In *International Conference on Machine Learning*, pages 18935–18958. PMLR, 2022.
- W. Scherrer and M. Deistler. Chapter 6 - vector autoregressive moving average models. In H. D. Vinod and C. Rao, editors, *Conceptual Econometrics Using R*, volume 41 of *Handbook of Statistics*, pages 145–191. Elsevier, 2019.
- R. D. Shah and J. Peters. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3):1514–1538, 2020.
- P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman. *Causation, prediction, and search*. MIT press, 2000.
- D. Staiger and J. H. Stock. Instrumental variables regression with weak instruments. *Econometrica: journal of the Econometric Society*, pages 557–586, 1997.
- J. H. Stock, J. H. Wright, and M. Yogo. A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*, 20(4):518–529, 2002.

- T. S. Verma. Invariant properties of causal models. Technical report, UCLA Cognitive Systems Laboratory, 1991.
- A. S. Weigend. *Time series prediction: Forecasting the future and understanding the past*. Routledge, 2018.
- N. Wiener. The theory of prediction. In *Modern Mathematics for Engineers*. McGraw-Hill, 1956.
- P. G. Wright. *Tariff on animal and vegetable oils*. Macmillan Company, 1928.
- S. Wright. The method of path coefficients. *The Annals of Mathematical Statistics*, 5(3):161–215, 1934.




Figure 8: Illustration of two different causal mechanisms ($S^{(1)}$, left, and $S^{(2)}$, right) which are observationally equivalent for any $b, c \in \mathbb{R}$ and $a \in [-1, 1]$ and Gaussian noise distribution $\mathcal{N}(0, \text{Id})$.

Appendix A. Additional details for Section 2

A.1 Relation to VARMA Processes

In this section, we discuss that the partially observed VAR(1) process can also be viewed as a VARMA(p, q) process. In this perspective, the difficulty of identifying β when H is unobserved is linked to the non-uniqueness of vector autoregressive moving average (VARMA) process representations. The observed process $[I^\top, X^\top, Y^\top]_{t \in \mathbb{Z}}^\top$ can be obtained as a linear transformation of the VAR(1) process S and as such it has a VARMA(p, q) process representation where $p \leq d$ and $q \leq (d-1)$ (Lütkepohl, 2005, Corollary 11.1.1). Intuitively, the dependencies between I, X, Y induced by the unobserved H process can instead be modelled by serially correlated errors and higher-order memory. In contrast to a VAR(p) process, however, the parameters of a VARMA process in standard form are not identified and different parameter settings may induce the same distribution over $[I^\top, X^\top, Y^\top]_{t \in \mathbb{Z}}^\top$ (Lütkepohl, 2005, Chapter 12.1). As such, it is not straight-forward to obtain β from a VARMA representation of the observed process, even when choosing a canonical representation such as the echelon form or the final equations form. In this work, we propose another approach and describe how to exploit the instrumental variables idea to identify β when H is unobserved, without needing to estimate all of A .

A.2 Observational Equivalence

Without an instrument, the causal effect β is, due to the hidden confounding, not identifiable in general. In a fully observed Gaussian VAR(1) process, the parameter matrix (which contains the causal effect) and the covariance matrix of the noises are uniquely determined by the distribution, and can be identified by least squares regression on the previous time step, for example (Hamilton, 1994, Chapter 11). This is not the case if parts of the system are unobserved. We consider a VAR(1) process over $H = [H^1, H^2]^\top, X$, and Y , where H is latent, and provide two different sets of parameters which entail the same observational distribution, that is, the same joint distribution of the observed process $[S_{XY,t}]_{t \in \mathbb{Z}} = [X_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$. Nevertheless, the causal effects in the two cases are different (one is 0, the other is $b \neq 0$), and so are the induced interventional distributions when intervening on X , see Appendix A.3 below. Consider the two coefficient matrices

$$A_1 := \begin{matrix} & \begin{matrix} H^1 \\ H^2 \\ X \\ Y \end{matrix} \\ \begin{matrix} H^1 \\ H^2 \\ X \\ Y \end{matrix} & \begin{pmatrix} a & 0 & 0 & 0 \\ c & 0 & 0 & 0 \\ c & 0 & 0 & 0 \\ 0 & b & 0 & 0 \end{pmatrix} \end{matrix} \quad \text{and} \quad A_2 := \begin{matrix} & \begin{matrix} H^1 \\ H^2 \\ X \\ Y \end{matrix} \\ \begin{matrix} H^1 \\ H^2 \\ X \\ Y \end{matrix} & \begin{pmatrix} a & 0 & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & c & 0 & 0 \\ 0 & 0 & b & 0 \end{pmatrix} \end{matrix}, \quad (12)$$

with coefficients $a \in (-1, 1)$ and $b, c \in \mathbb{R} \setminus \{0\}$. Consider the processes $S^{(1)}$ and $S^{(2)}$ satisfying

$$S_{t+1}^{(i)} = A_i S_t^{(i)} + \varepsilon_{t+1}^i, \quad (13)$$

with $\varepsilon_t^i \sim \mathcal{N}(0, \text{Id})$. Figure 8 depicts parts of the corresponding full time graphs. Let $S_{XY}^{(1)}$ and $S_{XY}^{(2)}$ denote the subprocesses where only X and Y are observed. The following result shows, that $S_{XY}^{(1)}$ and $S_{XY}^{(2)}$ are identically distributed, that is the two models arising from A_1 and A_2 are *observationally equivalent* (e.g., Rothenberg, 1971).

Proposition 14. *Let $S^{(1)}, S^{(2)}$ be the processes defined from Equation (13) with respective parameter matrices A_1 and A_2 from Equation (12) and $\mathcal{N}(0, \text{Id})$ -distributed noise. Then the observed subprocesses $S_{XY}^{(1)}$ and $S_{XY}^{(2)}$ are identically distributed.*

Proof Since each of the processes are jointly Gaussian with zero mean, their distributions are uniquely determined by the autocovariance matrices. This means that the observed processes are identically distributed if and only if $\mathbb{E}(S_{XY,t}^{(1)} S_{XY,t-s}^{(1)\top}) = \mathbb{E}(S_{XY,t}^{(2)} S_{XY,t-s}^{(2)\top})$ for all $s \geq 0$. For $i \in \{1, 2\}$, the s -th autocovariance of $S^{(i)}$ is given by:

$$\mathbb{E}(S_t^{(i)} S_{t-s}^{(i)\top}) = \sum_{k=0}^{\infty} A_i^{k+s} \text{Id} A_i^{k\top}.$$

Observe that

$$A_1^k = \begin{pmatrix} a^k & 0 & 0 & 0 \\ a^{k-1}c & 0 & 0 & 0 \\ a^{k-1}c & 0 & 0 & 0 \\ a^{k-2}bc & 0 & 0 & 0 \end{pmatrix} \quad A_2^k = \begin{pmatrix} a^k & 0 & 0 & 0 \\ 0 & a^k & 0 & 0 \\ 0 & a^{k-1}c & 0 & 0 \\ 0 & a^{k-2}bc & 0 & 0 \end{pmatrix}$$

for $k \geq 2$. Consequently,

$$A_1^{k+s} A_1^{k\top} = \begin{pmatrix} a^{2k+s} & a^{2k+s-1}c & a^{2k+s-1}c & a^{2k+s-2}bc \\ * & a^{2k+s-2}c^2 & a^{2k+s-2}c^2 & a^{2k+s-3}bc^2 \\ * & * & a^{2k+s-2}c^2 & a^{2k+s-3}bc^2 \\ * & * & * & a^{2k+s-4}b^2c^2 \end{pmatrix} \quad \text{and}$$

$$A_2^{k+s} A_2^{k\top} = \begin{pmatrix} a^{2k+s} & 0 & 0 & 0 \\ * & a^{2k+s} & a^{2k+s-1}c & a^{2k+s-2}bc \\ * & * & a^{2k+s-2}c^2 & a^{2k+s-3}bc^2 \\ * & * & * & a^{2k+s-4}b^2c^2 \end{pmatrix}$$

for $k \geq 2$ and $s \geq 0$, where the asterisks are given by symmetry of the matrices. For the case $k = 1, s = 0$, we have:

$$A_1 A_1^\top = \begin{pmatrix} a^2 & a^2c & a^2c & 0 \\ * & c^2 & c^2 & 0 \\ * & * & c^2 & 0 \\ * & * & * & b^2 \end{pmatrix} \quad \text{and} \quad A_2 A_2^\top = \begin{pmatrix} a^2 & 0 & 0 & 0 \\ * & a^2 & ac & 0 \\ * & * & c^2 & 0 \\ * & * & * & b^2 \end{pmatrix},$$

and if $k = 1, s \geq 1$:

$$A_1^{1+s} A_1^\top = \begin{pmatrix} a^{2+s} & a^{1+s}c & a^{1+s}c & 0 \\ a^{1+s}c & a^s c^2 & a^s c^2 & 0 \\ a^{1+s}c & a^s c^2 & a^s c^2 & 0 \\ a^s bc & a^{s-1}c^2b & a^{s-1}c^2b & 0 \end{pmatrix} \quad \text{and} \quad A_2^{1+s} A_2^\top = \begin{pmatrix} a^{2+s} & 0 & 0 & 0 \\ 0 & a^{2+s} & a^{1+s}c & 0 \\ 0 & a^{1+s}c & a^s c^2 & 0 \\ 0 & a^s bc & a^{s-1}c^2b & 0 \end{pmatrix}.$$

For any of the above matrices M , let M_{XY} denote the 2×2 submatrix in the bottom right corner, relating to the X, Y subprocess. In all of the above cases, these coefficients relating to the X, Y subprocess coincide, that is for any $k, s \geq 0$, $(A_1^{k+s} A_1^{k\top})_{XY} = (A_2^{k+s} A_2^{k\top})_{XY}$, and since therefore $(\sum_{k=0}^{\infty} A_1^{k+s} A_1^{k\top})_{XY} = \sum_{k=0}^{\infty} (A_2^{k+s} A_2^{k\top})_{XY}$, it follows that $\mathbb{E}(S_{XY,t}^{(1)} S_{XY,t-s}^{(1)\top}) = \mathbb{E}(S_{XY,t}^{(2)} S_{XY,t-s}^{(2)\top})$ for all $s \geq 0$. \blacksquare

A.3 Structural Causal Models and Interventions

We provide a formal introduction to structural causal models (SCMs) and interventions, which motivates the notion of a causal effect. For a more detailed introduction, see Pearl (2009); Bongers et al. (2021).

An SCM consists of a tuple $\Pi = (\mathcal{S}, P_\varepsilon)$ where \mathcal{S} is a set of structural assignments and P_ε describes the joint distribution of the error terms. For a finite collection of variables S^1, \dots, S^d , with structural assignments $\mathcal{S} := \{f^1, \dots, f^d\}$ and noise distribution $P_\varepsilon = P^1 \otimes \dots \otimes P^d$, for each $j = 1, \dots, d$, the structural equation of S^j is

$$S^j := f^j(\text{PA}_j, \varepsilon^j),$$

where ε^j is distributed according to P^j and the *parents* PA_j is a subset of $\{S^1, \dots, S^d\} \setminus \{S^j\}$. The SCM induces a corresponding graph \mathcal{G} over nodes $\{1, \dots, j\}$, where we draw an edge from j' to j if $S^{j'} \in \text{PA}_j$. We assume that the parent sets are such that \mathcal{G} is acyclic.

Similarly, we interpret the $\text{VAR}(p)$ process described in Equation (1) as a structural causal model over an infinite number of nodes $[S_t]_{t \in \mathbb{Z}}$. In this case, the structural assignments are $S_t := AS_{t-1} + \varepsilon_t$, $t \in \mathbb{Z}$. Here, the error terms ε_t are assumed to be i.i.d. over time, distributed according to P_ε . Furthermore, P_ε is a product distribution, and thus the error terms are jointly independent. The SCM entails an observational distribution on the variables $[S_t]_{t \in \mathbb{Z}}$ which we denote by P_S^Π .

Formally, an intervention on an SCM is a replacement of one or more of the structural assignments at one or more time points. Such a replacement induces a new SCM that we denote by $\tilde{\Pi} = (\tilde{\mathcal{S}}, \tilde{P}_\varepsilon)$. An example of an intervention on the above $\text{VAR}(p)$ process is to fix the value of X for some specific time point t_0 – we write this intervention as $\text{do}(X_{t_0} := x)$. Under this intervention, for $t \neq t_0$, the process still satisfies the original SCM, including assumptions on the noise variables.

The interventional distribution of Π under this intervention is defined as $P_S^{\Pi; \text{do}(X_{t_0} := x)} := P_S^{\tilde{\Pi}}$. In general, the interventional distribution of Π under an intervention is the distribution that is induced by the SCM $\tilde{\Pi}$ obtained by replacing some of the structural assignments. We require that this distribution exists and is unique. Depending on the application at hand, several interventions on the process are useful, including, for example, changing the dynamics of one component for all time points (Peters et al., 2022). In this work, we focus on an intervention at a particular time point. When performing such an intervention $\text{do}(X_{t_0} := x)$ on X_{t_0} , we have that $\frac{\partial}{\partial x} \mathbb{E}_{P_Y^{\Pi; \text{do}(X_{t_0} := x)}}[Y_{t_0+1}] = \alpha_{Y,X}$. This motivates calling $\beta = \alpha_{Y,X}$ the direct causal effect from X to Y in our parametric model (while the effect may be mediated in another parametrization, e.g., at finer temporal resolution). In several applications, the causal effect β is of interest by itself because it yields insight into understanding the causal structure of the problem. The causal effect, however, also comes with another benefit: it is optimal for prediction under intervention. We discuss this point of view in Section 4.3.

A.4 Defining Multivariate Total Causal Effects

In some cases, the effect we want to estimate may be more general than a single entry in one of the coefficient matrices A_k . In Section 2.1 we define the total causal effect of a single variable S_{t-l}^i on S_t^j as

$$\left(\sum_{\substack{1 \leq l_1, \dots, l_m \leq p \\ l_1 + \dots + l_m = l}} A_{l_1} \cdots A_{l_m} \right)_{j,i},$$

where A_l are the parameter matrices of the $\text{VAR}(p)$ process. Using the method of path coefficients (Wright, 1934), we now provide a more general definition, where \mathcal{X} may contain multiple variables.

Definition 15 (Path coefficients). *Let S be a $\text{VAR}(p)$ process, let $Y_t := S_t^{i_0}$ be a subprocess of S and let $\mathcal{X}_t = [S_{t-l_1}^{i_1}, \dots, S_{t-l_m}^{i_m}]^\top$ be a collection of subprocesses of S . For $j = 1, \dots, m$, we define a $S_{t-l_j}^{i_j}$ -causal path to be a directed path from $S_{t-l_j}^{i_j}$ to Y_t in the full time graph of S that does not intersect any other $S_{t-l_{j'}}^{i_{j'}}$, for $j' \neq j$. For a $S_{t-l_j}^{i_j}$ -causal path $\pi : S_{t-l_j}^{i_j} \xrightarrow{e_1} \dots \xrightarrow{e_d} S_t^{i_0}$, we define the path coefficient to be the product of linear coefficients along π , $c_\pi := \prod_{k=1}^d a_k$, where a_k denotes the entry in the coefficient matrix A_v , for the lag v corresponding to the edge e_k .*

We can now define the total causal effect.

Definition 16 (Total Causal Effect). *Let S be a $\text{VAR}(p)$ process, let $Y_t = S_t^{i_0}$ be a subprocess of S and let $\mathcal{X} = [S_{t-l_1}^{i_1}, \dots, S_{t-l_m}^{i_m}]^\top$ be a collection of subprocesses of S . For $j = 1, \dots, m$, the total causal effect, β^j , of $S_{t-l_j}^{i_j}$ on Y_t is the sum of path coefficients c_π over all $S_{t-l_j}^{i_j}$ -causal paths π from $S_{t-l_j}^{i_j}$ to Y_t , $\beta^j := \sum_{S_{t-l_j}^{i_j}\text{-causal paths } p} c_p$. Similarly the total causal effect, β , of \mathcal{X}_t on Y_t is the concatenation of these, $\beta := [\beta^1, \dots, \beta^m]^\top$.*

Appendix B. Additional details for Section 3

B.1 Asymptotic variances for i.i.d. estimators

Drawing on existing results (Hall, 2005), we now provide formulas for the asymptotic variances of the NIV and CIV estimators. If a unique solution (β, α) to Equation (6) exists, the asymptotic distribution of the $\hat{\beta}_{\text{NIV},T}$ estimator in the i.i.d. setting, discussed in Section 3.2, with the weight matrix $W := \mathbb{E}[\mathcal{I}\mathcal{I}^\top]^{-1}$ (which asymptotically is optimal, see Section 3.1) is given by

$$\sqrt{T}(\hat{\beta}_{\text{NIV},T} - \beta) \xrightarrow{d} \mathcal{N}(0, \Sigma_1),$$

for $T \rightarrow \infty$, where the asymptotic variance Σ_1 is given by

$$\Sigma_1 := (\mathbb{E}[\tilde{\mathcal{X}}\tilde{\mathcal{I}}^\top]K^{-1}\mathbb{E}[\tilde{\mathcal{X}}\tilde{\mathcal{I}}^\top])^{-1},$$

where $\tilde{\mathcal{X}} := [\mathcal{X}^\top, \mathcal{Z}^\top]^\top$, $K = \mathbb{E}((Y - \beta\mathcal{X} - \alpha\mathcal{Z})^2)\mathbb{E}(\mathcal{I}\mathcal{I}^\top)$. Σ_1 is a $(d_{\mathcal{X}} + d_{\mathcal{Z}}) \times (d_{\mathcal{X}} + d_{\mathcal{Z}})$ matrix, with the top-left $d_{\mathcal{X}} \times d_{\mathcal{X}}$ sub-matrix describing the asymptotic variance of \mathcal{X} .

Similarly, the asymptotic distribution of $\hat{\beta}_{\text{CIV},T}$ in the i.i.d. setting, discussed in Section 3.1, with weight matrix $W := \mathbb{E}[\text{var}(\mathcal{I}|\mathcal{B})]^{-1}$ is

$$\sqrt{T}(\hat{\beta}_{\text{CIV},T} - \beta) \xrightarrow{d} \mathcal{N}(0, \Sigma_2),$$

where

$$\Sigma_2 := (\mathbb{E}[\text{cov}(\mathcal{X}, \mathcal{I}|\mathcal{B})]K^{-1}\mathbb{E}[\text{cov}(\mathcal{X}, \mathcal{I}|\mathcal{B})^\top])^{-1},$$

and $K = \mathbb{E}[\text{var}((Y - \beta\mathcal{X})^2|\mathcal{B})]\mathbb{E}[\text{var}(\mathcal{I}|\mathcal{B})]$.

B.2 Asymptotic variances for estimators in time series

Closed-form expressions for the asymptotic variances of the NIV and CIV estimators can also be found for the $\text{VAR}(1)$ process presented in Section 4, but these are slightly more involved than for the i.i.d. setting presented in Appendix B.1. Assume the setting as described in Theorem 7 and consider the NIV estimator $\hat{\beta}_{\text{NIV},T}$. We then have

$$\sqrt{T}(\hat{\beta}_{\text{NIV},T} - \beta) \xrightarrow{d} \mathcal{N}(0, \Sigma_1),$$

where the asymptotic variance Σ_1 is given by

$$\Sigma_1 := (\mathbb{E}(\tilde{\mathcal{X}}\mathcal{I}^\top)K^{-1}\mathbb{E}(\tilde{\mathcal{X}}\mathcal{I}^\top)^\top)^{-1},$$

where $\tilde{\mathcal{X}} := [X_{t-1}^\top, Y_{t-1}^\top]^\top$, $\mathcal{I} := \{I_{t-2}, \dots, I_{t-m-1}\}$ and $K := \lim_{T \rightarrow \infty} \text{Var}\left(\frac{1}{\sqrt{T}} \sum_{t=1}^T (Y_t - \beta\mathcal{X} - \alpha\mathcal{Z})\mathcal{I}^\top\right)$ using the optimal choice of weight matrix, $W = K^{-1}$, see Hall (2005, Chapter 3).

Now for the CIV estimator, see Theorem 6, the asymptotic distribution of $\hat{\beta}_{\text{CIV},T}$ is

$$\sqrt{T}(\hat{\beta}_{\text{CIV},T} - \beta) \xrightarrow{d} \mathcal{N}(0, \Sigma_2),$$

where

$$\Sigma_2 := (\mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)]K^{-1}\mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)^\top])^{-1},$$

and $K := \lim_{T \rightarrow \infty} \text{Var}\left(\frac{1}{\sqrt{T}} \sum_{t=1}^T (r_{Y_t} - \beta r_{X_{t-1}})r_I^\top\right)$ with $r_i := i - \mathbb{E}[i | \mathcal{B}_t]$ using the optimal choice of weight matrix, $W = K^{-1}$, see Hall (2005, Chapter 3).

Appendix C. Additional details for Section 4

C.1 Failure of IV $_{X_{t-1} \rightarrow Y_t}(I_{t-2})$

In Section 4, we discuss that using I_{t-2} as an instrument for the effect $X_{t-1} \rightarrow Y_t$ is not valid. This is simply a consequence of requirement (CIV1') not being met, since I_{t-2} correlates with Y_t through other causal paths than $I_{t-2} \rightarrow X_{t-1} \rightarrow Y_t$. Here we formally prove that this procedure fails and derive the bias of the resulting estimate.

Proposition 17 (Failure of naive IV adaption). *Consider a VAR(1) process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2) with $d_I = d_X = d_H = d_Y = 1$. If $\text{cov}(X_{t-1}, I_{t-2}) \neq 0$ and $\alpha_{I,I}\alpha_{Y,Y} \neq 1$, the IV $_{X_{t-1} \rightarrow Y_t}(I_{t-2})$ estimator $\hat{\beta}$ converges in probability to*

$$(1 - \alpha_{I,I}\alpha_{Y,Y})^{-1}\beta.$$

Consequently, $\hat{\beta}$ is in general not consistent for the causal effect β of X_{t-1} on Y_t , unless I or Y do not have any autoregressive structure, that is, $\alpha_{I,I} = 0$ or $\alpha_{Y,Y} = 0$.

C.2 Example of a distribution that does not satisfy the rank requirement in Theorem 8

In Section 4.2.2, we have developed a criterion for identifiability that depends on the parameter matrix of a VAR(1) process satisfying (A2). We have showed in Corollary 9 that if parameter matrices are drawn from a distribution with density with respect to Lebesgue measure, then the identifiability criterion holds almost surely. In this section, we provide an example of a parameter matrix that does not satisfy the criterion.

Example 1. *Consider the case where $d_X > 1$, $d_I = 1$, and $\alpha_{X,X} = \text{diag}(c, \dots, c)$ for a $c \in \mathbb{R}$. By Theorem 8, β is not identifiable by NIV: this follows because A_{XY} is a lower triangular matrix with c, \dots, c (d_X times) and $\alpha_{Y,Y}$ on the diagonal, and the Jordan form J is a diagonal matrix with the same diagonal entries. Hence there are d_X Jordan blocks with the same eigenvalue c so the causal effect β is not identified by NIV. On the contrary, when $\alpha_{X,X} = \text{diag}(c_1, \dots, c_{d_X})$ (see Figure 9) where $c_i \neq c_j$ for all $i \neq j$, β is identified by NIV if also $\alpha_{Y,Y} \neq c_i$ for all i .*




Figure 9: Subgraph of the full-time graph of the structure discussed in Example 1 when $d_X = 2$. For simplicity, we do not draw nodes corresponding to the instrument, I , and the confounder, H .

Example 2. In the case of $d_X = 1, d_I = 1$, and $\alpha_{X,I} \neq 0$, β is identifiable. If, for example, $\alpha_{Y,Y} = \alpha_{X,X} =: \alpha$, we have,

$$A_{XY} = \begin{pmatrix} \alpha & 0 \\ \beta & \alpha \end{pmatrix} = \underbrace{\begin{pmatrix} 0 & 1 \\ \beta & 0 \end{pmatrix}}_{=:P} \begin{pmatrix} \alpha & 1 \\ 0 & \alpha \end{pmatrix} \underbrace{\begin{pmatrix} 0 & \frac{1}{\beta} \\ 1 & 0 \end{pmatrix}}_{=:P^{-1}}.$$

This has only one Jordan block with algebraic multiplicity $m = 2$ and

$$\left(P^{-1} \begin{pmatrix} \alpha_{X,I} \\ 0 \end{pmatrix} \right)_m = [1, 0] \begin{pmatrix} \alpha_{X,I} \\ 0 \end{pmatrix} = \alpha_{X,I} \neq 0,$$

where $(\cdot)_m$ refers to the m -th entry, so by Theorem 8, β is identifiable (and by similar arguments the same holds if $\alpha_{X,X} \neq \alpha_{Y,Y}$).

C.3 Algorithm for Prediction under Interventions

This section contains Algorithm 2, the procedure for using a causal effect to predict under interventions, as discussed in Section 4.3.

Algorithm 2 Linear prediction under the intervention $\text{do}(X_t := x)$

Input: Causal parameter β , sample $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_{t-1}] \in \mathbb{R}^{d_X \times (t-1)}$, $\mathbf{Y} = [\mathbf{Y}_1, \dots, \mathbf{Y}_t] \in \mathbb{R}^{1 \times t}$, interventional value x , lag parameters m and l

- 1: Compute the residual process $r_s := \mathbf{Y}_{s+1} - \beta \mathbf{X}_s$ for $s = 1, \dots, t-1$.
- 2: For $s > \max(k, m)$, linearly regress r_s on $\{\mathbf{X}_{s-k}, k = 1, \dots, m\}$ and $\{\mathbf{Y}_{s-j}, j = 0, \dots, l\}$ to obtain coefficients $\hat{\alpha}_{Y,X}^k$ and $\hat{\alpha}_{Y,Y}^j$.

Output:

- 3: Prediction $\hat{\mathbf{Y}}_{t+1} = \beta x + \sum_{k=1}^m \hat{\alpha}_{Y,X}^k \mathbf{X}_{t-k} + \sum_{j=0}^l \hat{\alpha}_{Y,Y}^j \mathbf{Y}_{t-j}$.
-

C.4 Identifying equations for time-inhomogeneous effects

Throughout the main part, we assume that the target parameter β is a constant. We now show, that if instead $\{\beta_t\}_{t \in \mathbb{Z}}$ is a sequence of random variables (not necessarily i.i.d.), then we can apply our methodology to estimate $\mathbb{E}[\beta_t]$. For a VAR(1) process S , we make the following modification of (A2)

(A2*) Assume that S satisfies (A2) except that for each $t \in \mathbb{Z}$, the structural assignment of Y_t is $Y_t = \beta_t X_{t-1} + \alpha_{Y,H} H_{t-1} + \alpha_{Y,Y} Y_{t-1} + \varepsilon_t^Y$, where $\{\beta_t\}_{t \in \mathbb{Z}}$ is a sequence of random variables.

The following two results generalize Theorems 6 and 7, by stating that the identifying equations used in Theorems 6 and 7 to identify β , identify $\mathbb{E}[\beta_t]$ in the case of a time inhomogeneous effect.

Proposition 18 (Identification with conditioning set and time inhomogeneous effect). *Consider a VAR(1) process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2*), ensuring that all quantities below are well-defined. Let either $\mathcal{B}_t := \{I_{t-3}\}$ or $\mathcal{B}_t := \{I_{t-3}, X_{t-2}, Y_{t-1}\}$ and let $b := \mathbb{E}[\beta_t]$ be the expectation of β_t . If $\beta_t \perp\!\!\!\perp (X_{t-1}, I_{t-2}, \mathcal{B}_t)$, the following two statements hold. (i) b satisfies the CIV moment condition $\mathbb{E}[\text{cov}(Y_t - bX_{t-1}, I_{t-2} | \mathcal{B}_t)] = 0$. (ii) Furthermore, if $\mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)]$ has rank d_X , then b is identified by CIV $X_{t-1} \rightarrow Y_t (I_{t-2} | \mathcal{B}_t)$.*

Proposition 19 (Identification with nuisance regressor and time inhomogeneous effect). *Consider a VAR(1) process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2*), ensuring that all quantities below are well-defined. Let $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$ for an $m \geq 1$ and $\mathcal{Z}_t := \{Y_{t-1}\}$ and let $b := \mathbb{E}[\beta_t]$ be the expectation of β_t . If $\beta_t \perp\!\!\!\perp (X_{t-1}, \mathcal{I}_t)$, the following two statements hold. (i) There exists $\alpha \in \mathbb{R}$ such that b satisfies the NIV moment condition $\mathbb{E}[\text{cov}(Y_t - bX_{t-1} - \alpha Z_t, \mathcal{I}_t)] = 0$. (ii) Further, if $\mathbb{E}[[X_{t-1}^\top, Z_t^\top]^\top \mathcal{I}_t^\top]$ has rank $d_X + d_Y$, b is identified by NIV $X_{t-1} \rightarrow Y_t (\mathcal{I}_t, \mathcal{Z}_t)$.*

The proofs can be found in Appendices E.15 and E.16, respectively. The above considers identification of a single time point. Further assumptions (such as observing each time point several times) may be required to obtain consistent estimators.

Appendix D. Simulation details and additional experiments

D.1 Correlation matrix for errors in Figure 5

In Figure 5, we plot the absolute errors for two CIV estimators and two NIV estimators for several different data generating mechanisms. Here, we print the Spearman correlations between the four estimators.

When $n = 100$, the correlation is

$$\begin{array}{l} \text{CIV}_{I,X,Y} \\ \text{CIV}_I \\ \text{NIV}_1 \text{ lag} \\ \text{NIV}_3 \text{ lags} \end{array} \begin{pmatrix} 1.000 & 0.635 & 0.391 & 0.390 \\ 0.635 & 1.000 & 0.276 & 0.241 \\ 0.391 & 0.276 & 1.000 & 0.393 \\ 0.390 & 0.241 & 0.393 & 1.000 \end{pmatrix}$$

When $n = 1,000$, the correlation is

$$\begin{array}{l} \text{CIV}_{I,X,Y} \\ \text{CIV}_I \\ \text{NIV}_1 \text{ lag} \\ \text{NIV}_3 \text{ lags} \end{array} \begin{pmatrix} 1.000 & 0.728 & 0.400 & 0.476 \\ 0.728 & 1.000 & 0.275 & 0.308 \\ 0.400 & 0.275 & 1.000 & 0.417 \\ 0.476 & 0.308 & 0.417 & 1.000 \end{pmatrix}$$

When $n = 10,000$, the correlation is

$$\begin{array}{l} \text{CIV}_{I,X,Y} \\ \text{CIV}_I \\ \text{NIV}_1 \text{ lag} \\ \text{NIV}_3 \text{ lags} \end{array} \begin{pmatrix} 1.000 & 0.742 & 0.385 & 0.465 \\ 0.742 & 1.000 & 0.280 & 0.308 \\ 0.385 & 0.280 & 1.000 & 0.459 \\ 0.465 & 0.308 & 0.459 & 1.000 \end{pmatrix}$$

Across sample sizes, all estimators are positively correlated, with the strongest correlation present between the two CIV estimators. Yet, for no other pair does the correlation exceed 0.5, indicating that the remaining estimator pairs do not agree on which data generating mechanisms are more ‘difficult’.




Figure 10: Absolute errors of CIV and NIV estimators of the causal effect in the experiment described in Appendix D.2. Each point corresponds to a single instantiation of a time series. The data are simulated using four different data generating mechanisms, with the strength of the instrument ranging from ‘Strong’ to ‘Very Weak’.

D.2 Varying instrument strength for a fixed data generating mechanism

In Section 5.1, we simulate random data generating mechanisms, and consider the error in estimating causal effects. The resulting errors plotted in Figure 5 then reflect both finite sample fluctuations and fluctuations stemming from the fact that estimation is more difficult in some data generating mechanisms than in others.

In this experiment, we consider a fixed data generating mechanism to compare only the finite sample error of the NIV and CIV estimators. We use the simulation setup described in Section 5, except that the entries of the parameter matrix are not random, but fixed at

$$A := \begin{matrix} & \begin{matrix} r^1 \\ r^2 \\ r^3 \\ x^1 \\ x^2 \\ H \\ Y \end{matrix} \\ \begin{matrix} r^1 \\ r^2 \\ r^3 \\ x^1 \\ x^2 \\ H \\ Y \end{matrix} & \begin{pmatrix} 0.25 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.75 & 0 & 0 & 0 & 0 \\ s \cdot 0.75 & s \cdot 0.75 & s \cdot 1 & 0.5 & 0.4 & 0.2 & 0 \\ s \cdot 1 & s \cdot 0.75 & s \cdot 0.75 & 0.5 & 0.5 & 0.2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.7 & 0 \\ 0 & 0 & 0 & 0.6 & 0.5 & 0.2 & 0.4 \end{pmatrix} \end{matrix},$$

where s is a parameter that controls the strength of the instrument to be either ‘Strong’ ($s = 1$), ‘Medium’ ($s = 0.5$), ‘Weak’ ($s = 0.1$) or ‘Very Weak’ ($s = 0.01$).

For each instrument strength, we sample a time series of $T = 1,000$ steps, apply the $\text{CIV}_{I,X,Y}$ and $\text{NIV}_{3 \text{ lags}}$ estimators from Figure 5 (left), and store the absolute deviation from the ground truth causal effect β . We repeat this 300 times and plot the resulting errors in Figure 10.

As expected, the average errors for both CIV and NIV increase as the instruments get weaker. However, for this data generating mechanism, the weakening of the instruments particularly hurts the performance of the CIV estimator: While the performance of NIV and CIV are comparable using the ‘strong’ instruments, the NIV estimator outperforms the CIV estimator in the majority of simulations with the ‘very weak’ instruments.

D.3 Asymptotic behaviour with weak instruments in the asymptotic regime

The instrument process may sometimes only explain a small fraction of the variance in the covariate process, a situation known as ‘weak instruments’ (Stock et al., 2002). When coefficients are kept constant, as is the case for the asymptotics discussed in Appendix B.2, both i.i.d. IV estimators as well as the time series estimators, which this paper focuses on, are asymptotically normal; however,




Figure 11: Average error (in log scale) of NIV, CIV and IV estimators for the causal effect of X_{t-1} on Y_t , where the strength of the instrument vanishes at rate $O(T^{-1/2})$. Similar to the weak instrument setting in i.i.d. data, in this asymptotic regime our estimators do not converge to the true causal effect (e.g., Stock et al., 2002).

if we consider an asymptotic regime where the effect of the instruments on the covariates vanishes at rate $O(T^{-\frac{1}{2}})$, Staiger and Stock (1997) show that in i.i.d. data, estimators are not necessarily asymptotically normal, and may be ill-behaved even for large sample sizes. In this simulation study, we compare the bias and variance of NIV and CIV estimators to a misspecified IV estimator. We sample data from the same data generating process as in Section 5 (see ‘Data generating process’), except that we scale the entries $\alpha_{X,I}$ by $\frac{1}{2\sqrt{T}}$. We use the same estimators as the ones described in Figure 5, and in addition an ordinary (misspecified) IV estimator, $IV_{X_{t-1} \rightarrow Y_t}(I_{t-2})$. In Figure 11, we plot the mean average error. As one would expect for GMM-based IV-methods (Stock et al., 2002), we do not see evidence that the error converges to 0, but rather remains approximately constant. Furthermore, the experiment does not suggest that our estimators are more affected by weak instruments than the ordinary IV estimator. While the $NIV_{1 \text{ lag}}$ estimator exhibits larger average errors, this aligns with its behavior under non-weak instrument conditions, as shown in Figure 5, where we noted that it has heavy-tailed errors due to just-identification.

Appendix E. Proofs

E.1 Proof of Conditional IV

For completeness, we now prove our statement about CIV from Section 3.1. Similar statements have been reported by Brito and Pearl (2002); Pearl (2009); Henckel (2021) but they differ in their precise formulations.

Proposition 20. *Consider a linear SCM (see Appendix A.3) over variables V , and let $\mathcal{I}, \mathcal{X}, \mathcal{B}, \{Y\} \subseteq V$ be disjoint collections of variables from V , and let \mathcal{G} be the corresponding DAG. Assume that \mathcal{I}, \mathcal{X} and Y have zero mean and finite second moment and let β be the causal coefficient with which \mathcal{X} enters the structural equation for Y (some of the entries of β can be zero, so not all variables in \mathcal{X} have to be parents of Y). We consider the following three requirements on $\mathcal{I}, \mathcal{X}, \mathcal{B}$ and Y*

(CIV1) \mathcal{I} and Y are d -separated given \mathcal{B} in the graph $\mathcal{G}_{\mathcal{X} \nrightarrow Y}$, that is the graph \mathcal{G} where all direct edges from \mathcal{X} to Y are removed,

(CIV2) \mathcal{B} is not a descendant of $\mathcal{X} \cup Y$ in \mathcal{G} , and

(CIV3) the matrix $\mathbb{E}[\text{cov}(\mathcal{X}, \mathcal{I} | \mathcal{B})]$ has rank $d_{\mathcal{X}}$, that is, full row rank.

If requirements (CIV1) and (CIV2) are met, $Y - \beta\mathcal{X} \perp\!\!\!\perp \mathcal{I}|\mathcal{B}$, and in particular β satisfies the CIV moment equation

$$\mathbb{E}[\text{cov}(Y - \beta\mathcal{X}, \mathcal{I}|\mathcal{B})] = 0. \quad (14)$$

If, additionally, requirement (CIV3) is met, β is the unique solution to this equation,

$$\mathbb{E}[\text{cov}(Y - b\mathcal{X}, \mathcal{I}|\mathcal{B})] = 0 \implies b = \beta.$$

Proof Due to the additive, linear structure of the SCM, we can write Y as a

$$Y = \beta\mathcal{X} + \pi\mathcal{B} + \gamma R + \varepsilon^Y,$$

where R are those parents of Y that are not in $\mathcal{X} \cup \mathcal{B}$ and $\pi \in \mathbb{R}^{d_Y \times d_{\mathcal{B}}}$, $\gamma \in \mathbb{R}^{d_Y \times d_R}$ are some coefficients.

We claim (1) that any path from \mathcal{I} to R is blocked by \mathcal{B} and (2) that $\mathbb{E}[\text{cov}(\varepsilon^Y, \mathcal{I}|\mathcal{B})] = 0$. It then follows from the global Markov property (Lauritzen, 1996), that $\mathbb{E}[\text{cov}(R, \mathcal{I}|\mathcal{B})] = 0$, and trivially also $\mathbb{E}[\text{cov}(\mathcal{B}, \mathcal{I}|\mathcal{B})] = 0$. Hence, since $Y - \beta\mathcal{X} = \pi\mathcal{B} + \gamma R + \varepsilon^Y$, it follows that $\mathbb{E}[\text{cov}(Y - \beta\mathcal{X}, \mathcal{I}|\mathcal{B})] = \mathbb{E}[\text{cov}(\pi\mathcal{B} + \gamma R + \varepsilon^Y, \mathcal{I}|\mathcal{B})] = 0$.

For (1), suppose for a contradiction that a path π between \mathcal{I} and R that is unblocked given \mathcal{B} exists. Case 1: π does not contain any edge from \mathcal{X} to Y . Then, the path that concatenates π with the corresponding edge from R to Y is an unblocked path (given \mathcal{B}) in the graph, where the edges from \mathcal{X} to Y are removed. This contradicts requirement (CIV1). Case 2: π contains an edge from $X \in \mathcal{X}$ to Y . Then, π contains either the structure $X \rightarrow Y \leftarrow$ or the structure $X \rightarrow Y \rightarrow$. The first case implies $\mathcal{B} \cap \text{DE}(Y) \neq \emptyset$, violating requirement (CIV2). In the second case, we either have that there is a directed path from Y to \mathcal{I} that is unblocked by \mathcal{B} , violating requirement (CIV1) or, again that $\mathcal{B} \cap \text{DE}(Y) \neq \emptyset$, violating requirement (CIV2).

For (2), we have that neither \mathcal{B} nor \mathcal{I} are descendants of Y : \mathcal{B} cannot be a descendant of Y due to requirement (CIV2), and by requirement (CIV1), \mathcal{I} can only be a descendant if \mathcal{B} is also a descendant, which is not possible. Every variable in the linear SCM can be rewritten as a function only of noise terms corresponding to ancestors. Applying this to \mathcal{B} and \mathcal{I} , we have (since Y is not an ancestor of neither \mathcal{I} nor \mathcal{B} and because ε^Y is independent of all other noise terms in the SCM) that ε^Y is independent of $(\mathcal{B}, \mathcal{I})$ and it follows that $\mathbb{E}[\text{cov}(\varepsilon^Y, \mathcal{I}|\mathcal{B})] = 0$.

When we additionally assume requirement (CIV3), the solution to Equation (14) is unique because the equation can be rewritten to

$$b\mathbb{E}[\text{cov}(\mathcal{X}, \mathcal{I}|\mathcal{B})] = \mathbb{E}[\text{cov}(Y, \mathcal{I}|\mathcal{B})],$$

and by the assumption of full row rank of $\mathbb{E}[\text{cov}(\mathcal{X}, \mathcal{I}|\mathcal{B})]$, this can have at most one solution. \blacksquare

E.2 Proof of Theorem 1

Theorem 1. Consider $p \in \mathbb{N}$, a (weakly) stationary time series S satisfying (A0) for that p , and consider disjoint, finite sets A, B, C of nodes in the full time graph $\mathcal{G}_{\text{full}}$. If A and C are d -separated given B in $\mathcal{G}_{\text{full}}$, then $A \perp\!\!\!\perp C|B$.

Proof Our proof is inspired by Lauritzen et al. (1990, Sec. 6). Given a time series (sub)graph \mathcal{G} over nodes V , we denote by $V_{[s,t]}^i := \{S_v^i | i \in \{1, \dots, d\}, s \leq v \leq t, S_v^i \in V\} \subseteq V$ the nodes in V between times s and t and denote by $\mathcal{G}_{[s,t]}$ the subgraph of \mathcal{G} , where only vertices in $V_{[s,t]}$ are included. Let further s_0, t_0 be the largest and smallest time points respectively such that $A \cup C \cup B \subseteq V_{[s_0, t_0]}$. Let $q \in \mathbb{N}_{>p}$ be such that if two nodes in $V_{[s_0, t_0]}$ are d -connected (with empty conditioning set) in $\mathcal{G}_{\text{full}}$,

then there is a d -connecting path in $\mathcal{G}_{[s_0-q, t_0]}$. Define the set

$\mathcal{A} := \{n \in \mathbb{N} \mid \text{for all } \mathcal{G}^* \text{ that are graphs over nodes } V^* \text{ with time indices between } s_0 - q \text{ and } t_0 \text{ s.t.}$
 all edges in \mathcal{G}^* point ‘forward in time’, i.e., $\forall k \in \mathbb{N}$, there is no edge $S_t^i \rightarrow S_{t-k}^j$ and
 \mathcal{G}^* is a subgr. of \mathcal{G}_f , where \mathcal{G}_f is the full-time gr. induced by $\mathcal{G}^0 := \mathcal{G}_{[s_0-q, s_0-1]}^*$ and
 $AN_{\mathcal{G}_f}(V^+)_{[s_0, t_0]} = V^+$, where $V^+ := V_{[s_0, t_0]}^*$ and
 $|V^+| = n$ and

for all $\text{VAR}(p)$ processes whose structure is specified by \mathcal{G}^0 and

for all $A^*, B, C^* \subseteq \mathcal{G}^*$ such that

$(A^+ \cup B \cup C^+) = V^+$ (where $A^+ := A_{[s_0, t_0]}^*$, $C^+ := C_{[s_0, t_0]}^*$, $\mathcal{G}^+ := \mathcal{G}_{[s_0, t_0]}^*$), and

$A^* = A^+ \cup (\text{PA}_{\mathcal{G}^*}(A^+) \cap V^0)$ and $C^* = C^+ \cup (\text{PA}_{\mathcal{G}^*}(C^+) \cap V^0)$

where V^0 are the nodes of \mathcal{G}^0

we have

$$A^* \perp_{\mathcal{G}^*} C^* | B \Rightarrow A^* \perp\!\!\!\perp C^* | B,$$

where $\perp_{\mathcal{G}}$ indicates d -separation in \mathcal{G} .

We show below, by induction, that $\mathcal{A} = \mathbb{N}$. This suffices to prove the statement of the theorem because of the following line of arguments. Let $V^0 := V_{[s_0-q, s_0-1]}$ and $V^+ := AN_{\mathcal{G}_{\text{full}}}(A \cup C \cup B)_{[s_0, t_0]}$. Let \mathcal{G}^* be the graph $\mathcal{G}_{\text{full}}$ restricted to the nodes in $V^* := V^0 \cup V^+$. Then, $A \perp_{\mathcal{G}^*} C | B$ (as \mathcal{G}^* is a subgraph of $\mathcal{G}_{\text{full}}$). If $V^+ \neq (A \cup C \cup B)$, we enlarge A and C to the disjoint sets A^+ and C^+ such that $A^+ \perp_{\mathcal{G}^*} C^+ | B$ and $V^+ = A^+ \cup B \cup C^+$. Let us define $A^* := A^+ \cup (\text{PA}_{\mathcal{G}^*}(A^+) \cap V^0)$ and $C^* := C^+ \cup (\text{PA}_{\mathcal{G}^*}(C^+) \cap V^0)$. Importantly, these two sets are disjoint (otherwise, A^+ and C^+ would have a joint parent not in B , violating $A^+ \perp_{\mathcal{G}^*} C^+ | B$). We then have that $A^* \perp_{\mathcal{G}^*} C^* | B$ (Indeed, if there is an open path from a node in $a \in A^*$ to a node in $c \in C^*$, given B , then there is an open path between a node in A^+ (either a itself or its child in A^+) to a node in C^+ (either c itself or its child in C^+), violating $A^+ \perp_{\mathcal{G}^*} C^+ | B$). But then $\mathcal{A} = \mathbb{N}$ implies $A^* \perp\!\!\!\perp C^* | B$. And this implies that $A \perp\!\!\!\perp C | B$, as A and C are subsets of A^* and C^* , respectively.

Let us now prove that $\mathcal{A} = \mathbb{N}$ by, (1), proving $1 \in \mathcal{A}$ and $2 \in \mathcal{A}$ and, (2), proving $n \in \mathcal{A}$ implies $n+1 \in \mathcal{A}$.

(1) We now prove that $1 \in \mathcal{A}$ and $2 \in \mathcal{A}$.

The only non-trivial statement occurs when $A^* \neq \emptyset$ and $C^* \neq \emptyset$ and $B = \emptyset$. Because $A^* \perp_{\mathcal{G}^*} C^*$, we have $AN_{\mathcal{G}^*}(A^*) \cap AN_{\mathcal{G}^*}(C^*) = \emptyset$. This implies $AN_{\mathcal{G}_{\text{full}}}(A^*) \cap AN_{\mathcal{G}_{\text{full}}}(C^*) = \emptyset$ because of the repetitive structure in a full time graph and the choice of q .

(2) We now prove that $n \in \mathcal{A}$ implies $n+1 \in \mathcal{A}$.

Assume $n \in \mathcal{A}$ and consider \mathcal{G}^* , \mathcal{G}^0 , \mathcal{G}^+ , V^+ , A^* , A^+ , B , C^* , C^+ as described in set \mathcal{A} with $|V^+| = n+1$. Consider a node $\lambda \in V^+$ that is a sink node in \mathcal{G}^* .

First, assume that $\lambda \in A^+$. Then $\text{PA}_{\mathcal{G}^*}(\lambda) \subseteq (A^* \setminus \{\lambda\}) \cup B$ (because d -separation would be violated if $C^* \cap \text{PA}(\lambda) \neq \emptyset$). Thus, it follows that

$$\lambda \perp\!\!\!\perp C^* | B \cup (A^* \setminus \{\lambda\}). \quad (15)$$

(Indeed, $\text{PA}_{\mathcal{G}_{\text{full}}}(\lambda) = \text{PA}_{\mathcal{G}^*}(\lambda)$ and thus there exist a coefficient vector $\gamma \in \mathbb{R}^{|\text{PA}_{\mathcal{G}^*}(\lambda)|}$ such that $\lambda = \gamma^\top \text{PA}_{\mathcal{G}^*}(\lambda) + \varepsilon^\lambda$; it then follows from the $MA(\infty)$ representation of S , see Brockwell and Davis (1991, Sec. 11.3) or Hamilton (1994), that $\lambda \perp\!\!\!\perp C^* \cup (B \cup (A^* \setminus \{\lambda\}) \setminus \text{PA}_{\mathcal{G}^*}(\lambda)) | \text{PA}_{\mathcal{G}^*}(\lambda)$; the claimed independence then follows with the weak union property.) If $A^* = \{\lambda\}$, then this already implies $A^* \perp\!\!\!\perp C^* | B$. Otherwise, we observe that $A^* \setminus \{\lambda\} \perp_{\mathcal{G}^{*m}} C^* | B$, where \mathcal{G}^{*m} denotes moralization of graph \mathcal{G} (Lauritzen, 1996), as d -separation is equivalent to separation in the moralized graph. But then, $A^* \setminus \{\lambda\} \perp_{(\mathcal{G}^{*m})_{V^* \setminus \{\lambda\}}} C^* | B$, as this graph contains no more edges. And therefore, $(A^* \setminus$

$\{\lambda\} \perp_{(\mathcal{G}_{V^* \setminus \{\lambda\}}^*)^m} C^* | B$ as, again, the graph contains no more edges. By the induction hypothesis $n \in \mathcal{A}$ and thus

$$A^* \setminus \{\lambda\} \perp C^* | B. \quad (16)$$

Combining (15) and (16) by the contraction property, it follows that

$$A^* \perp C^* | B.$$

Second, assume that $\lambda \in C^+$. The argument follows in the same way as in the case $\lambda \in A^+$.

Third, assume that $\lambda \in B$ (these are all cases since $A^+ \cup B \cup C^+ = V^+$). Since B separates A^* and C^* in $(\mathcal{G}^*)^m$, then $B \setminus \{\lambda\}$ also separates A^* and C^* in $(\mathcal{G}^{*m})_{V^* \setminus \{\lambda\}}$ (as it has no more edges), and therefore $B \setminus \{\lambda\}$ also separates A^* and C^* in $(\mathcal{G}_{V^* \setminus \{\lambda\}}^*)^m$, since this graph, again, has no more edges. By the induction hypothesis $n \in \mathcal{A}$, so

$$A^* \perp C^* | (B \setminus \{\lambda\}). \quad (17)$$

We now prove a second independence statement. We now make a case distinction (a) Assume that

$$\text{PA}_{\mathcal{G}^+}(\lambda) \cap A^* \neq \emptyset \quad \text{or} \quad \text{AN}_{\mathcal{G}^*}(\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}) \cap \text{AN}_{\mathcal{G}^*}(\text{PA}_{\mathcal{G}^*}(A^*)_{[s_0-q, s_0-1]}) \neq \emptyset.$$

Then, it follows that

$$\text{PA}_{\mathcal{G}^+}(\lambda) \cap C^* = \emptyset \quad \text{and} \quad \text{AN}_{\mathcal{G}^*}(\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}) \cap \text{AN}_{\mathcal{G}^*}(\text{PA}_{\mathcal{G}^*}(C^*)_{[s_0-q, s_0-1]}) = \emptyset. \quad (18)$$

(Indeed, if the statement on the left-hand side would be false, then there is a d -connecting path between A^* and C^* , given B : this goes from the element in $\text{PA}_{\mathcal{G}^+}(\lambda) \cap C^*$ to λ (which is in B) and then either to the element in $\text{PA}_{\mathcal{G}^+}(\lambda) \cap A^*$ or to the common ancestor of $\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}$ and $\text{PA}_{\mathcal{G}^*}(A^*)_{[s_0-q, s_0-1]}$ and then to the corresponding element in A^* . If the statement on the right-hand side would be false, then we can use the same path but this time going via the common ancestor of $\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}$ and $\text{PA}_{\mathcal{G}^*}(C^*)_{[s_0-q, s_0-1]}$.)

But then it follows that

$$\lambda \perp C^* | A^* \cup (B \setminus \{\lambda\}). \quad (19)$$

(Indeed, noting that $\text{PA}_{\mathcal{G}^+}(\lambda) \subseteq A^* \cup B \setminus \{\lambda\}$, we can replace the left-hand side by the $\text{MA}(\infty)$ representation of $\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}$. For C^* , we repeatedly use the structural equations except for variables in $B \setminus \{\lambda\}$ or variables in $\text{PA}_{\mathcal{G}^*}(C^*)_{[s_0-q, s_0-1]}$ (other variables will not occur: If there was a variable in A^* , for example, there would be a directed path from A^* to C^*). We then use the $\text{MA}(\infty)$ representation of $\text{PA}_{\mathcal{G}^*}(C^*)_{[s_0-q, s_0-1]}$. The statement then follows from the fact that $\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}$ and $\text{PA}_{\mathcal{G}^*}(C^*)_{[s_0-q, s_0-1]}$ do not have common ancestors, see Equation (18).

Combining (17) and (19) using the contraction property, it follows that $C^* \perp (\{\lambda\} \cup A^*) | (B \setminus \{\lambda\})$, and by the weak union property that

$$C^* \perp A^* | B.$$

(b) Now assume that

$$\text{PA}_{\mathcal{G}^+}(\lambda) \cap A^* = \emptyset \quad \text{and} \quad \text{AN}_{\mathcal{G}^*}(\text{PA}_{\mathcal{G}^*}(\lambda)_{[s_0-q, s_0-1]}) \cap \text{AN}_{\mathcal{G}^*}(\text{PA}_{\mathcal{G}^*}(A^*)_{[s_0-q, s_0-1]}) = \emptyset.$$

Similarly as in case (a) it follows that

$$\lambda \perp A^* | C^* \cup (B \setminus \{\lambda\}). \quad (20)$$

Combining (17) and (20) using the contraction property, it follows that $(\{\lambda\} \cup C^*) \perp A^* | (B \setminus \{\lambda\})$, and by the weak union property that

$$C^* \perp A^* | B.$$

This concludes the proof. ■

E.3 Proof of Theorem 2

Theorem 2 (Nuisance IV). *Consider a linear SCM (see Appendix A.3) over variables V , and let $\mathcal{I}, \mathcal{X}, \mathcal{Z}, \mathcal{B}, \{Y\} \subseteq V$ be disjoint collections of variables from V , and let \mathcal{G} be the corresponding DAG. Assume that $\mathcal{I}, \mathcal{X}, \mathcal{Z}$ and Y have zero mean and finite second moment and let β and α be the causal coefficients with which \mathcal{X} and \mathcal{Z} enter the structural equation for Y , respectively (some of the entries of β and α can be zero, so not all variables in \mathcal{X} and \mathcal{Z} have to be parents of Y). Let $\tilde{\mathcal{X}} := \mathcal{X} \cup \mathcal{Z}$. If requirements (CIV1) to (CIV3) are satisfied in \mathcal{G} for $\mathcal{I}, \tilde{\mathcal{X}}, \mathcal{B}$ and Y , the causal effect β of \mathcal{X} on Y is identified by $\text{NIV}_{\mathcal{X} \rightarrow Y}(\mathcal{I}, \mathcal{Z}|\mathcal{B})$.*

Proof By satisfaction of requirements (CIV1) to (CIV3), the causal effect $\tilde{\beta} = [\beta, \alpha]$ of $\tilde{\mathcal{X}} = \mathcal{X} \cup \mathcal{Z}$ on Y is identified by the instrument \mathcal{I} and the conditioning set \mathcal{B} by Proposition 20 in Appendix E.1. In particular, also the sub-vector of the IV estimate corresponding to \mathcal{X} is identified. \blacksquare

E.4 Proof of Proposition 3

Proposition 3. *If an effect can be identified by CIV and by NIV, then the estimators cannot be strictly sorted in terms of asymptotic variance. More specifically, there exist data generating processes, for which CIV has strictly smaller asymptotic variance and others, for which NIV has strictly smaller asymptotic variance.*

Proof We show this by considering two SCMs over 6 variables $S = [H, I, X, Y, Z, B]$ given by $S := AS + \varepsilon$ where A is such that the resulting graph is acyclic and admits the graphical model in Figure 3 (right) and $\varepsilon \sim \mathcal{N}(0, \Gamma)$; we provide two concrete choices for A and Γ below. We consider both the $\text{CIV}_{X \rightarrow Y}(I|B)$ and the $\text{NIV}_{X \rightarrow Y}([I, B], Z)$ estimates of β , the causal effect of X on Y , and provide two sets of parameters (A^I, Γ^I) and (A^{II}, Γ^{II}) such that if X is generated according to (A^I, Γ^I) , the CIV estimator has a lower asymptotic variance than the NIV estimator, and if S is generated according to (A^{II}, Γ^{II}) , the CIV estimator has a higher asymptotic variance than the NIV estimator.

$$A^I := \begin{matrix} & \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} \\ \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1.185 \\ 21.095 & 6.885 & 0 & 0 & 0 & -5.969 \\ -7.244 & 0 & 16.499 & 0 & -1.892 & 0 \\ 1.921 & 0 & 0 & 0 & 0 & 2.62 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{matrix} \Gamma^I := \begin{matrix} & \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} \\ \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} & \begin{pmatrix} 0.2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2.2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1.2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2.2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.2 \end{pmatrix} \end{matrix},$$

and

$$A^{II} := \begin{matrix} & \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} \\ \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -2.918 \\ -22.439 & 3.519 & 0 & 0 & 0 & 4.282 \\ 19.964 & 0 & 4.737 & 0 & 4.011 & 0 \\ 0.884 & 0 & 0 & 0 & 0 & -7.97 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{matrix} \Gamma^{II} := \begin{matrix} & \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} \\ \begin{matrix} \text{H} \\ \text{I} \\ \text{X} \\ \text{Y} \\ \text{Z} \\ \text{B} \end{matrix} & \begin{pmatrix} 3.2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 3.2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2.2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1.2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2.2 \end{pmatrix} \end{matrix}.$$

We can now use the formulas from Appendix B.1 to get the asymptotic variances, using that $\mathbb{E}[S] = 0$ and $\mathbb{E}[SS^\top] = (1 - A)^{-1}\Gamma(1 - A)^{-\top}$. When the data generating mechanism is (A^I, Γ^I) , the asymptotic distributions are specified by $\sqrt{T}(\hat{\beta}_{\text{CIV}} - \beta) \sim \mathcal{N}(0, 524.4)$ and $\sqrt{T}(\hat{\beta}_{\text{NIV}} - \beta) \sim \mathcal{N}(0, 522.7)$. Furthermore, when the data generating mechanism is (A^{II}, Γ^{II}) , the asymptotic distributions are $\sqrt{T}(\hat{\beta}_{\text{CIV}} - \beta) \sim \mathcal{N}(0, 320.0)$ and $\sqrt{T}(\hat{\beta}_{\text{NIV}} - \beta) \sim \mathcal{N}(0, 575.4)$, respectively. \blacksquare

E.5 Proof of Theorem 5

Theorem 5 (Time series IV by marginalization). *Consider a process $S = [S_t]_{t \in \mathcal{Z}}$ satisfying (A1) with full time graph $\mathcal{G}_{\text{full}}$. Let Y be some node in $\mathcal{G}_{\text{full}}$ and let $\mathcal{X}, \mathcal{I}, \mathcal{Z}$, and \mathcal{B} be disjoint collections of nodes from $\mathcal{G}_{\text{full}}$. Let $\tilde{\mathcal{X}} := \mathcal{X} \cup \mathcal{Z}$ and define $M := \{Y\} \cup \mathcal{X} \cup \mathcal{I} \cup \mathcal{Z} \cup \mathcal{B}$. Assume that requirements (CIV1') and (CIV2) are satisfied for $\mathcal{I}, \tilde{\mathcal{X}}, \mathcal{B}$ and Y in \mathcal{G}_M (see Definition 4). Then, the following three statements hold. (i) The total causal effect $[\beta, \alpha]$ of $[\mathcal{X}^\top, \mathcal{Z}^\top]^\top$ on Y satisfies the NIV moment equation*

$$\mathbb{E}[\text{cov}(Y - b\mathcal{X} - a\mathcal{Z}, \mathcal{I}|\mathcal{B})] = 0. \quad (8)$$

(ii) Further, if requirement (CIV3) is satisfied for $\mathcal{I}, \tilde{\mathcal{X}}, \mathcal{B}$, then $[\beta, \alpha]$ is the unique solution to Equation (8). (iii) If, additionally, $\mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{Z}$ and \mathbf{B} are observations of $\mathcal{X}, Y, \mathcal{I}, \mathcal{Z}$ and \mathcal{B} at T time points, W is a positive definite matrix, and

$$[\hat{b}, \hat{a}] := \arg \min_{b, a} \|\text{cov}(\mathbf{Y} - b\mathbf{X} - a\mathbf{Z}, \mathbf{I}|\mathbf{B})\|_W^2, \quad (9)$$

then \hat{b} is a consistent estimator for β .

Proof We first show part (i). Due to the additive, linear structure in (A1), we can rewrite Y as a linear combination of the parents of Y in $\mathcal{G}_{\text{full}}$ plus some additive noise, $Y = a_1 p_1 + \dots + a_m p_m + \varepsilon^Y$, where a_1, \dots, a_m are coefficients and p_1, \dots, p_m are nodes from $\mathcal{G}_{\text{full}}$. Similarly, we can recursively decompose the parents into their parents (in $\mathcal{G}_{\text{full}}$) and noise without replacing variables in $\mathcal{X} \cup \mathcal{Z} \cup \text{ND}(\mathcal{X} \cup \mathcal{Z})_{\mathcal{G}_{\text{full}}} \cup \mathcal{B}$, until the first time we have a decomposition

$$Y = \beta\mathcal{X} + \alpha\mathcal{Z} + \pi\mathcal{B} + \gamma R + \varepsilon,$$

where R are d_R variables from $\mathcal{G}_{\text{full}}$ that are not in \mathcal{B} or on any directed path from $\mathcal{X} \cup \mathcal{Z}$ to Y in $\mathcal{G}_{\text{full}}$,

$$R \cap (\mathcal{B} \cup \mathcal{X} \cup \mathcal{Z} \cup \text{DE}(\mathcal{X} \cup \mathcal{Z})_{\mathcal{G}_{\text{full}}}) = \emptyset, \quad (21)$$

(R may include descendants of \mathcal{B}), ε are all the weighted noise variables accumulated when doing the decomposition in parents, $[\beta, \alpha]$ is the joint total causal effect of \mathcal{X} and \mathcal{Z} on Y , and $\pi \in \mathbb{R}^{1 \times d_{\mathcal{B}}}, \gamma \in \mathbb{R}^{1 \times d_R}$ are some coefficients. The coefficients in front of \mathcal{X} and \mathcal{Z} are indeed β and α , because the joint total causal effect of $[\mathcal{X}^\top, \mathcal{Z}^\top]^\top$ is the product along all paths (in $\mathcal{G}_{\text{full}}$) from $\mathcal{X} \cup \mathcal{Z}$ to Y (see Appendix A.4), and by the assumption that $\mathcal{B} \subseteq \text{ND}(\mathcal{X} \cup \mathcal{Z})_{\mathcal{G}_M}$ (which implies $\mathcal{B} \subseteq \text{ND}(\mathcal{X} \cup \mathcal{Z})_{\mathcal{G}_{\text{full}}}$), no directed path from $\mathcal{X} \cup \mathcal{Z}$ to Y is blocked by \mathcal{B} .

We claim (1) that any path in $\mathcal{G}_{\text{full}}$ from \mathcal{I} to R is blocked by \mathcal{B} and (2) that $\mathbb{E}[\text{cov}(\varepsilon, \mathcal{I}|\mathcal{B})] = 0$. It then follows from Theorem 1, that $\mathbb{E}[\text{cov}(R, \mathcal{I}|\mathcal{B})] = 0$, and trivially also $\mathbb{E}[\text{cov}(\mathcal{B}, \mathcal{I}|\mathcal{B})] = 0$. Hence, since $Y - \beta\mathcal{X} - \alpha\mathcal{Z} = \pi\mathcal{B} + \gamma R + \varepsilon$, it follows that $\mathbb{E}[\text{cov}(Y - \beta\mathcal{X} - \alpha\mathcal{Z}, \mathcal{I}|\mathcal{B})] = \mathbb{E}[\text{cov}(\pi\mathcal{B} + \gamma R + \varepsilon, \mathcal{I}|\mathcal{B})] = 0$.

For (1), suppose for a contradiction that there exist $i \in \mathcal{I}$ and $m \in R$ and a path $p : i - v_1 - \dots - v_n - m$ in $\mathcal{G}_{\text{full}}$ that is unblocked given \mathcal{B} , where an undirected edge indicates a directed edge that can have any orientation. Since p is unblocked given \mathcal{B} , the non-colliders of p are disjoint from \mathcal{B} and for every collider v_k on p , there is a node $b^k \in \mathcal{B}$ such that $b^k \in \text{DE}(v_k)$ in $\mathcal{G}_{\text{full}}$. (Observe that the end node m is not in M : It is not in $\{Y\}$ (since it was found among the ancestors of Y), and by construction of R , it is not in $\mathcal{B} \cup \mathcal{X} \cup \mathcal{Z}$. Also $m \notin \mathcal{I}$: By the construction of R through recursive rewriting as parents, there exists a directed path from m to Y that does not intersect $\mathcal{X} \cup \mathcal{Z} \cup \mathcal{B}$; if $m \in \mathcal{I}$, this path would violate requirement (CIV1').)

Let w_1, \dots, w_L be those vertices among v_1, \dots, v_n that appear in the marginalized graph \mathcal{G}_M and let $w_0 := i$ and $w_{L+1} := Y$. We now, a) construct a path including the nodes w_0, \dots, w_{L+1} (and possibly some additional colliders, see below) from \mathcal{I} to Y in \mathcal{G}_M ; we then show that, b), this path

is still unblocked, given \mathcal{B} , if we remove all edges as specified by requirement (CIV1'), creating a contradiction to requirement (CIV1').

For a), consider those $0 \leq k \leq L$ for which w_k and w_{k+1} are not directly connected in p . For such k , consider the segment of p (as a path in $\mathcal{G}_{\text{full}}$) from $w_k - \dots - w_{k+1}$, where \dots represent edges from p that are not in M . If there are no colliders on this segment, by Definition 4 at least one of the edges $w_k \rightarrow w_{k+1}$, $w_k \leftarrow w_{k+1}$ or $w_k \leftrightarrow w_{k+1}$ are present in \mathcal{G}_M . If there is exactly one collider on $w_k - \dots - w_{k+1}$ (as part of p in $\mathcal{G}_{\text{full}}$), the segment must be one of the following four options:

$$\begin{aligned} & w_k \rightarrow \dots \rightarrow v_i \leftarrow \dots \leftarrow w_{k+1}, \\ & w_k \leftarrow \dots \leftarrow v_{j_1} \rightarrow \dots \rightarrow v_i \leftarrow \dots \leftarrow w_{k+1}, \\ & w_k \rightarrow \dots \rightarrow v_i \leftarrow \dots \leftarrow v_{j_2} \rightarrow \dots \rightarrow w_{k+1}, \\ & w_k \leftarrow \dots \leftarrow v_{j_1} \rightarrow \dots \rightarrow v_i \leftarrow \dots \leftarrow v_{j_2} \rightarrow \dots \rightarrow w_{k+1}, \end{aligned}$$

where v_{j_1}, v_{j_2} also are nodes on p and \dots is a short-hand for edges with the same orientation. But since $\text{DE}(v_i)_{\mathcal{G}_{\text{full}}} \cap \mathcal{B} \neq \emptyset$, this implies that at least one of the following paths are present in \mathcal{G}_M :

$$\begin{aligned} & w_k \rightarrow b_{k,1} \leftarrow w_{k+1} \\ & w_k \leftrightarrow b_{k,1} \leftarrow w_{k+1} \\ & w_k \rightarrow b_{k,1} \leftrightarrow w_{k+1} \\ & w_k \leftrightarrow b_{k,1} \leftrightarrow w_{k+1}, \end{aligned}$$

where $b_{k,1} \in \mathcal{B}$ (there is no node a from M on the path from v_i to $b_{k,1}$, because if $a \in \mathcal{X} \cup \mathcal{Z} \cup \{Y\}$, requirement (CIV2) would be violated, and if $a \in \mathcal{I}$, this would constitute a path in $\mathcal{G}_{\text{full}}$ from Y to \mathcal{I} that is unblocked given \mathcal{B} , using the same argument as for the original path). Similarly, if there are several colliders on $w_k - \dots - w_{k+1}$, a path $w_k \rightarrow b_{k,1} \leftrightarrow \dots \leftrightarrow b_{k,L} \leftarrow w_{k+1}$ (or one of the configurations $\rightarrow \dots \leftrightarrow$, $\leftrightarrow \dots \leftarrow$ or $\leftrightarrow \dots \leftrightarrow$ as first and last edge) is present in \mathcal{G}_M , where $b_{k,1}, \dots, b_{k,L} \in \mathcal{B}$.

We now construct a path p_M in \mathcal{G}_M that is d -connecting \mathcal{I} and Y , given \mathcal{B} : For $k = 0, \dots, L-1$, paste together the segments (in \mathcal{G}_M) from w_k to w_{k+1} including those possible colliders $b_{k,j}$ discussed above. Further, add the edge $w_L \rightarrow y$ or $w_L \leftrightarrow y$, depending on the orientation of the edge $w_L - m$ in p . If w_k was a collider on p , it is also a collider on p_M . Since p was unblocked in $\mathcal{G}_{\text{full}}$, given \mathcal{B} , p_M is unblocked in \mathcal{G}_M , given \mathcal{B} : consider a collider on p_M ; either it is one of the $b_{k,j}$ (in this case, it does not block p_M) or it is one of the w_k (in this case, w_k is also a collider on p and thus has a descendant in \mathcal{B} , which is still a descendant of w_k in \mathcal{G}_M ; again, it does not block p_M).

We now turn to b) and argue that the path is still unblocked in \mathcal{G}_M , given \mathcal{B} , if we remove the edges specified in requirement (CIV1'). First, assume for some $w_k \in \mathcal{X} \cup \mathcal{Z}$ that p_M contained the segment $w_k \rightarrow w_{k+1}$. Then there would be some $k+2 \leq k' \leq L+1$ such that $w_{k'} \notin \text{DE}(\mathcal{X} \cup \mathcal{Z})_{\mathcal{G}_{\text{full}}}$ (because otherwise m would be a descendant of $\mathcal{X} \cup \mathcal{Z}$, violating (21)). But this would imply that p_M has a collider, which is a descendant of $\mathcal{X} \cup \mathcal{Z}$ and an ancestor of \mathcal{B} (in \mathcal{G}_M), which is not possible by requirement (CIV2). Thus, p_M cannot contain an edge $w_k \rightarrow w_{k+1}$ where $w_k \in \mathcal{X} \cup \mathcal{Z}$.

Second, assume that p_M does not contain an edge $w_k \rightarrow w_{k+1}$ with $w_k \in \mathcal{X} \cup \mathcal{Z}$ and that for some $w_k \in \mathcal{X} \cup \mathcal{Z}$, p_M contains an edge $w_{k-1} \xleftarrow{e} w_k$ that satisfies the criterion for removal under requirement (CIV1'). We can choose w_k such that e is the first edge on p_M satisfying the criterion and, by definition of the criterion, there exists a directed path $w_{k-1} \rightarrow u_1 \rightarrow \dots \rightarrow u_t \rightarrow Y$ in \mathcal{G}_M where $u_i \notin \mathcal{X} \cup \mathcal{Z}$ and $u_i \notin \mathcal{B}$ (because otherwise $\mathcal{B} \cap \text{DE}(w_k)_{\mathcal{G}_M} \neq \emptyset$). Because p_M is unblocked given \mathcal{B} in \mathcal{G}_M and because no edge is removed on $\tilde{p}_M = w_0 - \dots - w_{k-1} \rightarrow u_1 \rightarrow \dots \rightarrow Y$, \tilde{p}_M is a violation of requirement (CIV1'), and therefore p_M does not contain any edge $w_{k-1} \leftarrow w_k$ that satisfies the criterion for removal under requirement (CIV1').

Combining 'First' and 'Second', the path p_M does not contain any edges that would be removed under the criterion in requirement (CIV1'). Also, since p_M is unblocked given \mathcal{B} , any collider is an ancestor of some $b \in \mathcal{B}$, and that collider is still an ancestor of $b \in \mathcal{B}$ after removing edges satisfying

the criterion in requirement (CIV1') (otherwise $b \in \text{DE}(\mathcal{X} \cup \mathcal{Z})_{\mathcal{G}_M}$, violating requirement (CIV2)). In conclusion, p_M is also unblocked given \mathcal{B} in the graph where we remove edges satisfying the criterion in requirement (CIV1') from \mathcal{G}_M . This concludes the proof of (1).

A similar argument proves (2), that is, $\mathbb{E}[\text{cov}(\varepsilon, \mathcal{I}|\mathcal{B})] = 0$: Each ε^i was accumulated as a noise variable of an ancestor (in $\mathcal{G}_{\text{full}}$) of Y , $A^i \notin \mathcal{X} \cup \mathcal{Z} \cup \mathcal{B}$; by construction, a directed path from $\mathcal{Z} \cup \mathcal{X}$ through A^i to Y exists in $\mathcal{G}_{\text{full}}$ that does not intersect $\mathcal{X} \cup \mathcal{Z} \cup \mathcal{B}$, except at the first node of this path. Hence, \mathcal{B} does not contain a descendant of A^i in $\mathcal{G}_{\text{full}}$ (because that would imply \mathcal{B} containing a descendant of $\mathcal{X} \cup \mathcal{Z}$ in \mathcal{G}_M , violating requirement (CIV2)). Also, A^i is not an ancestor of any node in \mathcal{I} in $\mathcal{G}_{\text{full}}$, because that would imply an unblocked path from \mathcal{I} via A_i to Y in $\mathcal{G}_{\text{full}}$ that does not contain any node in \mathcal{B} and therefore this corresponds to an unblocked path in \mathcal{G}_M , too. Using the MA(∞)-representation of \mathcal{B} and \mathcal{I} , see Brockwell and Davis (1991, Sec. 11.3) or Hamilton (1994), ε^i is independent of $(\mathcal{B}, \mathcal{I})$, and so $\mathbb{E}[\text{cov}(\varepsilon^i, \mathcal{I}|\mathcal{B})] = 0$. This concludes the proof of the first part.

Part (ii) follows because if the $(d_{\mathcal{X}} + d_{\mathcal{Z}}) \times d_{\mathcal{I}}$ matrix $\mathbb{E}[\text{cov}(\tilde{\mathcal{X}}, \mathcal{I}|\mathcal{B})]$ has rank $d_{\mathcal{X}} + d_{\mathcal{Z}}$, then if a solution to the moment equation $\mathbb{E}[\text{cov}(Y, \mathcal{I}|\mathcal{B})] = \beta \mathbb{E}[\text{cov}(\tilde{\mathcal{X}}, \mathcal{I}|\mathcal{B})]$ exists, it is unique.

For part (iii), let $\tilde{\mathbf{X}} := [\mathbf{X}^\top \quad \mathbf{Z}^\top]^\top$. By Equation (5), we have to show that

$$\hat{\mathbb{E}}[r_{\mathbf{Y}} r_{\mathbf{I}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}} r_{\tilde{\mathbf{X}}}^\top] \left(\hat{\mathbb{E}}[r_{\tilde{\mathbf{X}}} r_{\mathbf{I}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}} r_{\tilde{\mathbf{X}}}^\top] \right)^{-1} \xrightarrow{P} \gamma$$

with $\gamma := [\beta^\top, \alpha^\top]^\top$. From Equation (2) we have that empirical moments converge in probability to the population moment, and thus, using Slutsky's Theorem, we get that

$$\hat{\mathbb{E}}[r_{\mathbf{Y}} r_{\mathbf{I}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}} r_{\tilde{\mathbf{X}}}^\top] \left(\hat{\mathbb{E}}[r_{\tilde{\mathbf{X}}} r_{\mathbf{I}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}} r_{\tilde{\mathbf{X}}}^\top] \right)^{-1} \xrightarrow{P} \mathbb{E}[r_{Y_t} r_{\mathcal{I}_t}^\top] W \mathbb{E}[r_{\mathcal{I}_t} r_{\tilde{\mathcal{X}}_t}^\top] \left(\mathbb{E}[r_{\tilde{\mathcal{X}}_t} r_{\mathcal{I}_t}^\top] W \mathbb{E}[r_{\mathcal{I}_t} r_{\tilde{\mathcal{X}}_t}^\top] \right)^{-1},$$

where $r_{\tilde{\mathcal{X}}} = \tilde{\mathcal{X}} - \mathbb{E}[\tilde{\mathcal{X}}|\mathcal{B}]$ and similarly for $r_{\mathcal{Y}}$ and $r_{\mathcal{I}}$. We can rewrite $\mathbb{E}[r_{Y_t} r_{\mathcal{I}_t}^\top]$ by adding and subtracting $\gamma r_{\tilde{\mathcal{X}}_t}$:

$$\begin{aligned} \mathbb{E}[r_{Y_t} r_{\mathcal{I}_t}^\top] &= \mathbb{E}[(r_{Y_t} - \gamma r_{\tilde{\mathcal{X}}_t}) r_{\mathcal{I}_t}^\top] + \gamma \mathbb{E}[r_{\tilde{\mathcal{X}}_t} r_{\mathcal{I}_t}^\top] \\ &= \mathbf{0} + \gamma \mathbb{E}[r_{\tilde{\mathcal{X}}_t} r_{\mathcal{I}_t}^\top]. \end{aligned}$$

The first term is zero due to the conditional uncorrelation established in (i) and we can thus conclude that

$$\begin{aligned} \mathbb{E}[r_{Y_t} r_{\mathcal{I}_t}^\top] W \mathbb{E}[r_{\mathcal{I}_t} r_{\tilde{\mathcal{X}}_t}^\top] \left(\mathbb{E}[r_{\tilde{\mathcal{X}}_t} r_{\mathcal{I}_t}^\top] W \mathbb{E}[r_{\mathcal{I}_t} r_{\tilde{\mathcal{X}}_t}^\top] \right)^{-1} \\ = \gamma \mathbb{E}[r_{\tilde{\mathcal{X}}_t} r_{\mathcal{I}_t}^\top] W \mathbb{E}[r_{\mathcal{I}_t} r_{\tilde{\mathcal{X}}_t}^\top] \left(\mathbb{E}[r_{\tilde{\mathcal{X}}_t} r_{\mathcal{I}_t}^\top] W \mathbb{E}[r_{\mathcal{I}_t} r_{\tilde{\mathcal{X}}_t}^\top] \right)^{-1} \\ = \gamma. \end{aligned}$$

■

E.6 Proof of Proposition 17

Proposition 17 (Failure of naive IV adaption). *Consider a VAR(1) process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2) with $d_I = d_X = d_H = d_Y = 1$. If $\text{cov}(X_{t-1}, I_{t-2}) \neq 0$ and $\alpha_{I,I} \alpha_{Y,Y} \neq 1$, the IV $X_{t-1} \rightarrow Y_t(I_{t-2})$ estimator $\hat{\beta}$ converges in probability to*

$$(1 - \alpha_{I,I} \alpha_{Y,Y})^{-1} \beta.$$

Consequently, $\hat{\beta}$ is in general not consistent for the causal effect β of X_{t-1} on Y_t , unless I or Y do not have any autoregressive structure, that is, $\alpha_{I,I} = 0$ or $\alpha_{Y,Y} = 0$.




Figure 12: (left) Marginalization of the full time graph to nodes I_{t-2} , I_{t-3} , X_{t-1} and Y_t . (middle) Marginalization of the full time graph to nodes I_{t-2} , X_{t-1} and Y_t and their lagged values. (right) Marginalization to m instrument nodes $I_{t-2}, \dots, I_{t-m-1}$, and X_{t-1} , Y_t , and Y_{t-1} .

Proof Since ε_t^Y and H_{t-1} are both independent of I_{t-2} (for H_{t-1} , this follows from (A2) and Theorem 1), it follows that

$$\begin{aligned} \mathbb{E}[Y_t I_{t-2}] &= \alpha_{Y,Y} \alpha_{I,I} \mathbb{E}[Y_{t-1} I_{t-3}] + \beta \mathbb{E}[X_{t-1} I_{t-2}] \\ \implies \mathbb{E}[Y_t I_{t-2}] &= (1 - \alpha_{Y,Y} \alpha_{I,I})^{-1} \beta \mathbb{E}[X_{t-1} I_{t-2}], \end{aligned}$$

where in the last step we use that by covariance stationarity we have $\mathbb{E}[Y_{t-1} I_{t-3}] = \mathbb{E}[Y_t I_{t-2}]$. The $\text{IV}_{X_{t-1} \rightarrow Y_t}(I_{t-2})$ moment equation is $\mathbb{E}[(Y_t - b X_{t-1}) I_{t-2}] = 0$, which has the solution (because $d_I = d_X = d_Y = 1$)

$$b = \frac{\mathbb{E}[I_{t-2} Y_t]}{\mathbb{E}[I_{t-2} X_{t-1}]}.$$

By plugging in the expression for $\mathbb{E}[Y_t I_{t-2}]$ above, we get $b = (1 - \alpha_{Y,Y} \alpha_{I,I})^{-1} \beta$. \blacksquare

E.7 Proof of Theorem 6

Theorem 6 (Identification with conditioning set). *Consider a VAR(1) process $S = [S_t]_{t \in \mathbb{Z}}$ with $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2). Let either $\mathcal{B}_t := \{I_{t-3}\}$ or $\mathcal{B}_t := \{I_{t-3}, X_{t-2}, Y_{t-1}\}$. Then, the following three statements hold. (i) The causal effect β of X_{t-1} on Y_t satisfies the CIV moment condition $\mathbb{E}[\text{cov}(Y_t - \beta X_{t-1}, I_{t-2} | \mathcal{B}_t)] = 0$. (ii) Furthermore, if $\mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)]$ has rank d_X , then β is identified by $\text{CIV}_{X_{t-1} \rightarrow Y_t}(I_{t-2} | \mathcal{B}_t)$. (iii) If, additionally, \mathbf{X}_t , \mathbf{Y}_t , \mathbf{I}_t , and \mathbf{B}_t are observations of X, Y, I and \mathcal{B} at T time points, then β can be consistently estimated as $T \rightarrow \infty$ by $\text{CIV}_{\mathbf{X}_{t-1} \rightarrow \mathbf{Y}_t}(\mathbf{I}_{t-2} | \mathbf{B}_t)$, that is, the output of Algorithm 1.*

Proof By Theorem 5, it suffices for part (i) to show that requirements (CIV1') and (CIV2) are satisfied for $\mathcal{X}_t := \{X_{t-1}\}$, $\mathcal{I}_t := \{I_{t-2}\}$, \mathcal{B}_t , and Y_t in the marginalized graph \mathcal{G}_{M_t} with $M_t := \mathcal{X}_t \cup \mathcal{I}_t \cup \mathcal{B}_t \cup \{Y_t\}$ (with \mathcal{B}_t being either of the two sets from the theorem), see Figure 12 left and middle. For either choice of \mathcal{B}_t , \mathcal{B}_t is not a descendant of X_{t-1} and Y_t , so requirement (CIV2) is satisfied (see Figure 12 left and middle). To show that requirement (CIV1') holds, we argue that every path from I_{t-2} to Y_t is blocked by I_{t-3} in $\mathcal{G}_{M_t(X_{t-1} \not\rightarrow Y_t)}$, the graph obtained from \mathcal{G}_{M_t} by removing the directed edge from X_{t-1} to Y_t . For either graph, we have that any path from I_{t-2} to Y_t either contains the non-collider I_{t-3} or the collider X_{t-1} . Since I_{t-3} is in the conditioning

set \mathcal{B}_t (for either definition of \mathcal{B}_t) and $(\{X_{t-1}\} \cup \text{DE}(X_{t-1})) \cap \mathcal{B}_t = \emptyset$, any path from I_{t-2} to Y_t in $\mathcal{G}_{M_t(X_{t-1} \not\rightarrow Y_t)}$ is blocked by \mathcal{B}_t .

Parts (ii) and (iii) follow directly from Theorem 5. ■

E.8 Proof of Theorem 7

Theorem 7 (Identification with nuisance regressor). *Consider a VAR(1) process $S = [S_t]_{t \in \mathbb{Z}}$ with $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2). Let $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$ for an $m \geq 1$ and $\mathcal{Z}_t := \{Y_{t-1}\}$. Then, the following three statements hold. (i) There exists $\alpha \in \mathbb{R}$ such that the causal effect β of X_{t-1} on Y_t satisfies the NIV moment condition $\text{cov}(Y_t - \beta X_{t-1} - \alpha \mathcal{Z}_t, \mathcal{I}_t) = 0$. (ii) Further, if $\mathbb{E}[[X_{t-1}^\top, \mathcal{Z}_t^\top]^\top \mathcal{I}_t^\top]$ has rank $d_X + d_Y$, β is identified by $\text{NIV}_{X_{t-1} \rightarrow Y_t}(\mathcal{I}_t, \mathcal{Z}_t)$. (iii) If, additionally, $\mathbf{X}_t, \mathbf{Y}_t, \mathbf{I}_t$, and \mathbf{Z}_t are observations of X, Y, \mathcal{I} and \mathcal{Z} at T time points, then β can be consistently estimated as $T \rightarrow \infty$ by $\text{NIV}_{\mathbf{X}_{t-1} \rightarrow \mathbf{Y}_t}(\mathbf{I}_t, \mathbf{Z}_t)$, that is, the output of Algorithm 1.*

Proof Let $[\beta, \alpha]$ be the total causal effect of $[X_{t-1}, \mathcal{Z}_t]$ on Y_t . By Theorem 5, it suffices for part (i) to show that requirements (CIV1') and (CIV2) are satisfied for $\mathcal{X}_t := \{X_{t-1}\}, \mathcal{Z}_t, \mathcal{I}_t$, and Y_t in the marginalized graph \mathcal{G}_{M_t} with $M_t := \mathcal{X}_t \cup \mathcal{I}_t \cup \mathcal{Z}_t \cup \{Y_t\}$ (see Figure 12 right). Since $\mathcal{B} = \emptyset$, requirement (CIV2) is trivially satisfied. It remains to argue that requirement (CIV1) requirement (CIV1') is satisfied, that is, that \mathcal{I}_t is d -separated from Y_t in the marginalized graph $\mathcal{G}_{M_t(X_{t-1}, Y_{t-1} \not\rightarrow Y_t)}$ obtained from \mathcal{G}_{M_t} by removing the edges $Y_{t-1} \rightarrow Y_t$ and $X_{t-1} \rightarrow Y_t$. Let $s \in \{t-m-1, \dots, t-2\}$. Every path from I_s to Y_t must go through either the collider $\rightarrow X_{t-1} \leftarrow$ or the collider $\rightarrow Y_{t-1} \leftarrow$ and since the conditioning set is empty, those paths are blocked.

Parts (ii) and (iii) follow directly from Theorem 5. ■

E.9 Proof of Theorem 8

A BRIEF REVIEW OF JORDAN CANONICAL FORMS

If M is an arbitrary square matrix of size $d \times d$, there exists a unique (up to row or column permutations) square invertible matrix Q of the same dimension such that $M = QJQ^{-1}$ where J is a $d \times d$ block diagonal matrix

$$J = J_{m_1}(\lambda_1) \oplus \dots \oplus J_{m_k}(\lambda_k) := \text{diag}(J_{m_1}(\lambda_1), \dots, J_{m_k}(\lambda_k)) \quad (22)$$

with each *Jordan block* $J_{m_i}(\lambda_i)$ being an $m_i \times m_i$ matrix having one value λ_i on the diagonal and ones on the superdiagonal (and zeros elsewhere): that is, for all $m \in \mathbb{N}_{>0}$,

$$J_m(\lambda) := \begin{pmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{pmatrix}.$$

We sometimes write J_m instead of $J_m(\lambda)$ and call Equation (22) the Jordan canonical form. Jordan forms and the involved matrices satisfy the following properties (Horn and Johnson, 1985).

- Let N_m (we simply write N if the dimension is obvious) be the canonical nilpotent matrix of degree m , that is the $m \times m$ -matrix with ones in the superdiagonal and zeroes elsewhere. Then $J_m(\lambda) = \lambda I + N$ and by the binomial formula $J_m^n = \sum_{i=0}^n \binom{n}{i} \lambda^{n-i} N^i$.

- Every diagonal value of a Jordan block is an eigenvalue of M and for every eigenvalue λ of M , there is at least one Jordan block with diagonal λ . There may however be more than one Jordan block for the same eigenvalue.
- The geometric multiplicity of an eigenvalue λ is the number of corresponding Jordan blocks
- The algebraic multiplicity of an eigenvalue λ is the sum of the sizes m_i of the corresponding Jordan blocks.
- If M is diagonalizable, all Jordan blocks are of size one, which is equivalent to the algebraic and geometric multiplicities being equal.

SOME LEMMATA FOR THE PROOF OF THEOREM 8

We say that a vector v in a d -dimensional vector space is *cyclic of the $d \times d$ matrix J* if $v, Jv, \dots, J^{d-1}v$ constitute a basis for the vector space.

Lemma 21. *Let $J = J_{m_1}(\lambda_1) \oplus \dots \oplus J_{m_k}(\lambda_k)$ be a block Jordan form over \mathbb{C} for a square matrix J . If two or more blocks have the same eigenvalue, no vector $v \in \mathbb{R}^{\sum_{i=1}^k m_i}$ is cyclic of J .*

Proof It suffices to consider the case $J = J_{m_1}(\lambda) \oplus J_{m_2}(\lambda)$ where without loss of generality $m_1 \geq m_2$. For $J_{m_1}(\lambda) = \lambda 1 + N_{m_1}$ and $J_{m_2}(\lambda) = \lambda 1 + N_{m_2}$ the degree m_1 minimal polynomial $p(x) = (x - \lambda)^{m_1}$ annihilates J such that $p(J) = 0$. Consequently J^{m_1} can be written as a linear combination of J^0, \dots, J^{m_1-1} . In particular $J^0 v, \dots, J^{m_1+m_2-1} v$ cannot be linearly independent. ■

Lemma 22. *Let $J = J_{m_1}(\lambda_1) \oplus \dots \oplus J_{m_k}(\lambda_k)$ be a block Jordan form over \mathbb{C} for a square matrix J , with each block corresponding to a distinct eigenvalue λ_i . Then $v \in \mathbb{C}^{\sum_{i=1}^k m_i}$ is a cyclic vector for J if and only if for each $d = 1, \dots, k$ the entry $v_{\sum_{i=1}^d m_i}$ is non-zero.*

Proof We first show by contraposition that if v is cyclic for J , the corresponding entries will be non-zero. If it does not hold that for each $d = 1, \dots, k$ the entry $v_{\sum_{i=1}^d m_i}$ is non-zero, we may, without loss of generality, assume that the last entry of v is zero, such that $v = [u, 0]^\top$ for suitable $u \in \mathbb{C}^{-1+\sum_{i=1}^k m_i}$. Denote $\lambda = \lambda_k$ the eigenvalue corresponding to the last Jordan block $J_{m_k}(\lambda_k)$; observe that the bottom row of J^n is $[0, \dots, 0, \lambda^n]$ for any power n , and so the last entry of $J^n v$ is 0 for every n and consequently the matrix $[J^0 v, J^1 v, \dots, J^{(\sum_{i=1}^k m_i)-1} v]$ has a 0-row. Consequently, v is not cyclic of J . This shows that if v is cyclic, the entries $v_{\sum_{i=1}^d m_i}$ are non-zero.

Now we show the other implication by induction over k . Assume first that $k = 1$, i.e., $J = J_m(\lambda)$ consists of a single Jordan block. A vector $v = [v_1, \dots, v_m]^\top$ is cyclic of J if $v_m \neq 0$. Indeed, consider coefficients a_0, \dots, a_{m-1} such that $0 = \sum_{n=0}^{m-1} a_n J^n v$. Recall that $J^n = \sum_{i=0}^n \binom{n}{i} \lambda^{n-i} N^i$, which implies

$$0 = \sum_{n=0}^{m-1} a_n J^n v = \sum_{n=0}^{m-1} a_n \left(\sum_{i=0}^n \binom{n}{i} \lambda^{n-i} N^i \right) v = \sum_{i=0}^{m-1} \left(\sum_{n=i}^{m-1} a_n \binom{n}{i} \lambda^{n-i} \right) N^i v,$$

where in the final equality, we swap the order of summation, using that the pairs (n, i) where $n \in \{0, \dots, m-1\}$ and $i \in \{0, \dots, n\}$, are the same as the pairs (n, i) where $i \in \{0, \dots, m-1\}$ and $n \in \{i, \dots, m-1\}$. Since $v_m \neq 0$ the collection $N^0 v, \dots, N^{m-1} v$ are linearly independent: they form an upper-triangular matrix with v_m on the diagonal. This implies that, in particular, the coefficient on $N^{m-1} v$ must be 0. But this coefficient equals $\sum_{n=m-1}^{m-1} a_n \binom{n}{m-1} \lambda^{n-(m-1)} = a_{m-1}$, and so $a_{m-1} = 0$. Substituting this into the coefficient on $N^{m-2} v$, one obtains $a_{m-2} = 0$ and so forth.

Therefore, $a_n = 0$ for all n and thus $J^0v, \dots, J^{m-1}v$ are linearly independent, so v is cyclic of J if $v_m \neq 0$.

Next assume that the induction hypothesis holds for any matrix with k Jordan blocks $J = J_{m_1}(\lambda_1) \oplus \dots \oplus J_{m_k}(\lambda_k)$ with distinct eigenvalues for each block and for all $v = [v_1, \dots, v_{\sum_{i=1}^k m_i}]^\top$ where for every $d = 1, \dots, k$: $v_{\sum_{i=1}^d m_i} \neq 0$. Now consider the additional Jordan block $D = J_{m_{k+1}}(\lambda_{k+1})$ where $\lambda_{k+1} \neq \lambda_1, \dots, \lambda_k$ and the vector u whose last entry $u_{m_{k+1}}$ is non-zero, and let $\tilde{J} = J \oplus D$, $\tilde{v} = [v^\top, u^\top]^\top$.

Define the polynomial p of degree $\sum_{i=1}^k m_i$ by $p(x) = (x - \lambda_1)^{m_1} (x - \lambda_2)^{m_2} \dots (x - \lambda_k)^{m_k}$. Observe that $p(\lambda_{k+1}) \neq 0$ and so $p(D)$ is an upper triangular matrix with $p(\lambda_{k+1})$ on the diagonal. Hence the last entry of the vector $p(D)u$ is $p(\lambda_{k+1})u_{m_{k+1}}$ which is non-zero, and so $p(D)u$ is cyclic of D (by the initial step of the induction proof). Further observe that p annihilates each of the previous blocks because $J_{m_i}(\lambda_i) = \lambda_i 1 + N$ so $(J_{m_i}(\lambda_i) - \lambda_i 1)^{m_i} = N^{m_i} = 0$. Consequently,

$$\begin{aligned} p(\tilde{J}) &= p(J_{m_1}(\lambda_1)) \oplus \dots \oplus p(J_{m_k}(\lambda_k)) \oplus p(D) \\ &= 0 \oplus \dots \oplus 0 \oplus p(D), \end{aligned}$$

so that

$$p(\tilde{J})\tilde{v} = [0^\top, \dots, 0^\top, (p(D)u)^\top]^\top. \quad (23)$$

Now to show \tilde{v} is cyclic of \tilde{J} , we take any vector $x \in \mathbb{C}^{\sum_{i=1}^k m_i}$ and $y \in \mathbb{C}^{m_{k+1}}$. Our aim is to show that $[x^\top, y^\top]^\top$ is in the span of $\tilde{J}^0\tilde{v}, \dots, \tilde{J}^{(\sum_{i=1}^{k+1} m_i)-1}\tilde{v}$. Since v is cyclic of J , x can be expressed as a linear combination of $J^0v, \dots, J^{\sum_{i=1}^k m_i-1}v$. Taking the same linear combination of $\tilde{J}^0\tilde{v}, \dots, \tilde{J}^{\sum_{i=1}^k m_i-1}\tilde{v}$ yields the vector $[x^\top, z^\top]^\top$ for some $z \in \mathbb{C}^{m_{k+1}}$. Since $p(D)u$ is cyclic of D , we can write $y - z$ as a linear combination of $D^0p(D)u, D^1p(D)u, \dots, D^{m_{k+1}-1}p(D)u$. It follows from Equation (23) that by taking the same linear combination of $\tilde{J}^0p(\tilde{J})\tilde{v}, \dots, \tilde{J}^{m_{k+1}-1}p(\tilde{J})\tilde{v}$ one obtains $[0^\top, (y - z)^\top]^\top$. Since p is a polynomial of degree $\sum_{i=1}^k m_i$, it follows that both $[x^\top, z^\top]^\top$ and $[0^\top, (y - z)^\top]^\top$ lie in the span of $\tilde{J}^0\tilde{v}, \dots, \tilde{J}^{\sum_{i=1}^{k+1} m_i-1}\tilde{v}$, and so does $[x^\top, y^\top]^\top$. Since x and y were arbitrary, the entire space is spanned, completing the induction step. \blacksquare

PROOF OF THEOREM 8

Theorem 8. *Consider a process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2). Assume that $d_I = d_Y = 1$ and let $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$, where $m = d_X + d_Y$. Let A_{XY} and A_I be defined as in Equation (10). The following three statements are equivalent:*

- i) $\text{rank } \mathbb{E}[[X_{t-1}^\top, Y_{t-1}^\top]^\top \mathcal{I}_t^\top] = d_X + d_Y$.
- ii) The matrix $[A_{XY}^0 A_I, A_{XY}^1 A_I, \dots, A_{XY}^{d_X} A_I]$ is invertible, where A_{XY}^0 is the identity matrix of size $(d_X + d_Y) \times (d_X + d_Y)$.
- iii) Different Jordan blocks of J have different eigenvalues and for all $q \in \{1, \dots, k\}$, the coefficient $w_{\sum_{i=1}^q m_i}$ is non-zero; here, $J = Q^{-1} A_{XY} Q$ is the Jordan normal form¹¹ of A_{XY} , with k Jordan blocks $J = \text{diag}(J_{m_1}(\lambda_1), \dots, J_{m_k}(\lambda_k))$, each with size m_i and eigenvalue λ_i , and w are the coefficients of A_I in the basis of the generalized eigenvectors Q , that is, $w = Q^{-1} A_I$.

Proof First observe that

$$\begin{pmatrix} X_t \\ Y_t \end{pmatrix} = A_I I_{t-1} + A_{XY} \begin{pmatrix} X_{t-1} \\ Y_{t-1} \end{pmatrix} + \underbrace{\begin{pmatrix} \nu_X \\ \nu_Y \end{pmatrix} H_{t-1} + \varepsilon_t^{X,Y}}_{\text{uncorrelated to I}}$$

11. See Appendix E.9 for the definition of Jordan normal forms and the notation that we use.

and consequently:

$$\mathbb{E} \left[\begin{pmatrix} X_t \\ Y_t \end{pmatrix} I_t \right] = \mathbb{E} \left[\left(A_I I_{t-1} + A_{XY} \begin{pmatrix} X_{t-1} \\ Y_{t-1} \end{pmatrix} \right) \alpha_{I,I} I_{t-1} \right].$$

From this, we obtain

$$\mathbb{E} \left[\begin{pmatrix} X_t \\ Y_t \end{pmatrix} I_t \right] = \underbrace{\mathbb{E}[I_t^2]}_{=:v_I} \alpha_{I,I} \underbrace{(1 - \alpha_{I,I} A_{XY})^{-1}}_{=:B^{-1}} A_I.$$

This expression is justified as B is invertible. (Indeed, if $\alpha_{I,I} = 0$, this is trivial. If $\alpha_{I,I} \neq 0$, since $e_1^\top A_1 = \alpha_{I,I} e_1^\top$, where A_1 is coefficient matrix assumed in (A2) and $e_1 = [1, 0, \dots, 0]^\top$ is the first unit vector, $\alpha_{I,I}$ is an eigenvalue of A_1^\top and thus of A_1 and in particular, by (A1), it has absolute value strictly smaller than 1. B is degenerate if and only if $A_{XY} - \frac{1}{\alpha_{I,I}}$ is, but this would imply that $\frac{1}{\alpha_{I,I}}$ would be an eigenvalue of A_{XY} , but since the eigenvalues of A_{XY} are also eigenvalues of A_1 (if $A_{XY}v = \lambda v$, then $A_1[0, v^\top]^\top = \lambda[0, v^\top]^\top$) and belong to the interior of the unit circle, $\frac{1}{\alpha_{I,I}}$ cannot be an eigenvalue of A_{XY} .)

By performing the same expansion for $\mathbb{E}[[X_t^\top, Y_t^\top]^\top I_{t-j}]$ for $j \geq 1$ and plugging in the above, we obtain:

$$\begin{aligned} \mathbb{E} \left(\begin{pmatrix} X_t \\ Y_t \end{pmatrix} I_{t-1} \right) &= A_I v_I + A_{XY} \mathbb{E} \left(\begin{pmatrix} X_{t-1} \\ Y_{t-1} \end{pmatrix} I_{t-1} \right) = v_I B^{-1} A_I \\ \mathbb{E} \left(\begin{pmatrix} X_t \\ Y_t \end{pmatrix} I_{t-2} \right) &= v_I [A_{XY} B^{-1} + \alpha_{I,I} 1] A_I \\ \mathbb{E} \left(\begin{pmatrix} X_t \\ Y_t \end{pmatrix} I_{t-3} \right) &= v_I [A_{XY}^2 B^{-1} + \alpha_{I,I} A_{XY} + \alpha_{I,I}^2 1] A_I \end{aligned}$$

and in general:

$$\mathbb{E} \left(\begin{pmatrix} X_t \\ Y_t \end{pmatrix} I_{t-1-j} \right) = v_I \left[A_{XY}^j B^{-1} + \sum_{k=0}^{j-1} \alpha_{I,I}^{j-k} A_{XY}^k \right] A_I. \quad (24)$$

The columns (denote the j 'th column by col_j) of $\Sigma := \mathbb{E}[[X_{t-1}, Y_{t-1}] \mathcal{I}_t^\top]$ are exactly those given by Equation (24). If we deduct $\alpha_{I,I} \text{col}_{j-1}$ from col_j we obtain:

$$\begin{aligned} &v_I \left[\left(\sum_{k=0}^{j-1} \alpha_{I,I}^{j-k} A_{XY}^k + A_{XY}^j B^{-1} \right) - \alpha_{I,I} \left(\sum_{k=0}^{j-2} \alpha_{I,I}^{j-1-k} A_{XY}^k + A_{XY}^{j-1} B^{-1} \right) \right] A_I \\ &= v_I \left[\alpha_{I,I} A_{XY}^{j-1} + A_{XY}^j B^{-1} - \alpha_{I,I} A_{XY}^{j-1} B^{-1} \right] A_I \\ &= v_I (1 - \alpha_{I,I}^2) A_{XY}^j B^{-1} A_I. \end{aligned}$$

Since deducting columns from each other does not change the determinant, we can create a simpler matrix, Σ_{equiv} , with the same determinant: for $j \in \{2, \dots, k\}$ we deduct $\alpha_{I,I} \text{col}_{j-1}$ from col_j (starting with the largest j , that is first deducting $\alpha_{I,I} \text{col}_{k-1}$ from col_k , etc.), and obtain

$$\Sigma_{\text{equiv}} = v_I \left[A_{XY}^0 B^{-1} A_I, \quad (1 - \alpha_{I,I}^2) A_{XY}^1 B^{-1} A_I, \quad \dots, \quad (1 - \alpha_{I,I}^2) A_{XY}^{d_X} B^{-1} A_I \right].$$

By the Laplace expansion, removing $(1 - \alpha_{I,I}^2)$ terms appearing in all but the first column scales the determinant by a factor $\frac{1}{(1 - \alpha_{I,I}^2)^{d_X}}$, but it will not change its invertibility (from the requirement on the eigenvalues in (A1), it follows that $1 - \alpha_{I,I}^2 > 0$). The same applies to $v_I = \mathbb{E}[I_t^2]$. Hence, Σ is invertible if and only if

$$\Sigma_{\text{equiv},2} := \left[A_{XY}^0 B^{-1} A_I, \quad A_{XY}^1 B^{-1} A_I, \quad \dots, \quad A_{XY}^{d_X} B^{-1} A_I \right]$$

is invertible. Now observe that B^{-1} commutes with A_{XY}^j . This follows because $BA_{XY} = (1 - \alpha_{I,I}A_{XY})A_{XY} = A_{XY}(1 - \alpha_{I,I}A_{XY}) = A_{XY}B$. This implies $B^{-1}A_{XY} = A_{XY}B^{-1}$, because for any matrix M where $MB = BM$, it follows that

$$M = MBB^{-1} = BMB^{-1} \implies B^{-1}M = B^{-1}BMB^{-1} = MB^{-1}.$$

This implies that

$$\Sigma_{\text{equiv},2} = B^{-1} \underbrace{\left[A_{XY}^0 A_I, \dots, A_{XY}^{d_x} A_I \right]}_{=:\Sigma_{\text{equiv},3}}.$$

Since B^{-1} is invertible, it has non-zero determinant, and again invertibility of Σ is equivalent to invertibility of $\Sigma_{\text{equiv},3}$.

Let $A_{XY} = QJQ^{-1}$ be the Jordan block factorization. Observe that

$$\left\{ A_{XY}^0 A_I, \dots, A_{XY}^{d_x} A_I \right\} = Q \left\{ J^0 Q^{-1} A_I, \dots, J^{d_x} Q^{-1} A_I \right\}.$$

And finally, since Q is invertible, invertibility of Σ is equivalent to $Q^{-1}A_I$ being cyclic of J . According to Lemma 21 if two or more Jordan blocks have the same eigenvalue, no vector can be cyclic, so in particular not $Q^{-1}A_I$. If on the contrary no eigenvalue is shared across Jordan blocks (equivalently, the geometric multiplicity of every eigenvalue is 1), it follows from Lemma 22 that $Q^{-1}A_I$ is a cyclic vector of J if and only if the vector $Q^{-1}A_I$ is non-zero in the entries indexed by $\sum_{i=1}^d m_i$ for all $d = 1, \dots, k$. Writing $A_I = Qa$ in the basis of the columns of Q for some coefficient vector $a \in \mathbb{C}^{d_x+1}$, this means $Q^{-1}A_I$ is a cyclic vector for J exactly when the coefficients $a_{\sum_{i=1}^d m_i}$ are non-zero for all $d = 1, \dots, k$. This concludes the proof. \blacksquare

E.10 Proof of Corollary 9

Corollary 9. *Consider a VAR(1) process S with $d_I = 1$ and parameter matrix A , and assume that sparsity pattern of A is given by (A2) and that the non-zero entries of A are drawn from any distribution which has density with respect to Lebesgue measure. Then β is identifiable with probability 1.*

Proof We check the conditions of Theorem 8. Since the entries are drawn from a density with respect to Lebesgue measure, the eigenvalues are almost surely distinct. Thus, taking into account the sparsity pattern, A_{XY} can almost surely be diagonalized and the corresponding Jordan form has blocks of size one, all with distinct eigenvalues.

Also with probability one, $w = Q^{-1} \begin{pmatrix} \alpha_{X,I} \\ 0 \end{pmatrix}$ does not have any zeroes: Q is determined from A_{XY} (so Q depends only on $\alpha_{X,X}$, $\alpha_{Y,Y}$ and β), and so the probability that $[\alpha_{X,I}^\top, 0]^\top$ is orthogonal to any of the rows of Q^{-1} is 0. \blacksquare

E.11 Proof of Corollary 10

Corollary 10. *Consider a VAR(1) process S satisfying (A2) with $d_I > 1$ instrument processes $I^{(1)}, \dots, I^{(d_I)}$. Assume that there is at least one instrument process $I^{(j)}$ such that both of the following conditions hold.*

- i) $I_t^{(j)}$ is independent of $I_s^{(i)}$ for all t, s and $i \neq j$, and
- ii) the requirements of Theorem 8 are satisfied for the reduced process $(I^{(j)}, X, Y)$.

Then β is identifiable.

Proof Although the instruments $I^{(i)}, j \neq i$ are observed, we may treat them as latent, being part of the latent process $\tilde{H}_t := (H_t, I_t^{(i, i \neq j)})$. By i), $I^{(j)}$ is independent of \tilde{H} . By ii), Theorem 7, and Theorem 8, β is identifiable in the reduced process $(I^{(j)}, X, Y)$, and the solution is therefore also unique in the full system (I, X, Y) . \blacksquare

E.12 Proof of Proposition 11

Proposition 11 (Identification with conditioning set relaxing the VAR(p) assumption). *Consider a process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying Equation (11) and (A1')–(A3'). Let \mathcal{B}_t be a set of variables satisfying $\text{PA}(I_{t-2}) \subseteq \mathcal{B}_t \subseteq \text{ND}(Y_t) \cap \text{ND}(I_{t-2})$ in $\mathcal{G}_{\text{full}}$. Then, (i), (ii) and (iii) from Theorem 6 hold.*

Proof To obtain $\mathbb{E}[\text{cov}((Y_t - \beta X_{t-1})I_{t-2}^\top | \mathcal{B}_t)] = 0$, we show that

$$Y_t - \beta X_{t-1} \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t.$$

Using that $Y_t - \beta X_{t-1} = \alpha_{Y,Y} Y_{t-1} + g(\varepsilon_t^Y, H_{t-1})$ it suffices to show that $Y_{t-1} \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t$ and $(\varepsilon_t^Y, H_{t-1}) \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t$ since g is measurable. For the first conditional independence we use that $Y_{t-1}, \mathcal{B}_t \subseteq \text{ND}(I_{t-2})$ and $\text{PA}(I_{t-2}) \subseteq \mathcal{B}_t$ to conclude $Y_{t-1} \perp_d I_{t-2} | \mathcal{B}_t$ in $\mathcal{G}_{\text{full}}$ (Indeed, any path from I_{t-2} to Y_{t-1} leaves I_{t-2} either through a parent of I_{t-2} or must contain a collider that is a descendant of I_{t-2} .) By the global Markov property, contained in (A3'), this implies $Y_{t-1} \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t$. For the second conditional independence, we show $\varepsilon_t^Y \perp\!\!\!\perp I_{t-2} | (\mathcal{B}_t, H_{t-1})$ and $H_{t-1} \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t$ and use the contraction property of conditional independence to obtain $(\varepsilon_t^Y, H_{t-1}) \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t$. Now, $H_{t-1} \perp\!\!\!\perp I_{t-2} | \mathcal{B}_t$ holds by the global Markov property since $H_{t-1} \in \text{ND}(I_{t-2})$ and $\text{PA}(I_{t-2}) \subseteq \mathcal{B}_t$. To show $\varepsilon_t^Y \perp\!\!\!\perp I_{t-2} | (\mathcal{B}_t, H_{t-1})$, we use that by (A3') ε_t^Y is independent of any finite subset of $\text{ND}(Y_t)$ in $\mathcal{G}_{\text{full}}$. We have that $\mathcal{B}_t \cup \{H_{t-1}\} \cup \{I_{t-2}\} \subseteq \text{ND}(Y_t)$ and thus by weak union we get $\varepsilon_t^Y \perp\!\!\!\perp I_{t-2} | (\mathcal{B}_t, H_{t-1})$ as desired. This proves part (i).

Part (ii) follows because the moment equation $\mathbb{E}[\text{cov}((Y_t - \beta X_{t-1})I_{t-2}^\top | \mathcal{B}_t)] = 0$ is the same as in Theorem 6, and so the rank requirement for identifiability is also the same.

To show part (iii), we have to show that

$$\hat{\mathbb{E}}[r_{\mathbf{Y}_t} r_{\mathbf{I}_{t-2}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}_{t-2}} r_{\mathbf{X}_{t-1}}^\top] \left(\hat{\mathbb{E}}[r_{\mathbf{X}_{t-1}} r_{\mathbf{I}_{t-2}}^\top] W \hat{\mathbb{E}}[r_{\mathbf{I}_{t-2}} r_{\mathbf{X}_{t-1}}^\top] \right)^{-1} \xrightarrow{P} \beta.$$

This is analogous to the argument in the proof of Theorem 5, except for that the convergence of empirical moments is now guaranteed by (A2') (instead of (A1)). Hence, using Slutsky's Theorem and rewriting Y_t as in Theorem 5, gives the desired convergence. \blacksquare

E.13 Proof of Proposition 12

Proposition 12 (Identification with nuisance regressor relaxing the VAR(p) assumption). *Consider a process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying Equation (11) and (A1'), (A2'), and (A4'). Let $\mathcal{Z}_t := \{Y_{t-1}\}$ and $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$ for an $m \geq 1$. Then, (i), (ii), and (iii) from Theorem 7 hold.*

Proof By Equation (11), we have that $Y_t - \beta X_{t-1} - \alpha_{Y,Y} Y_{t-1} = g(\varepsilon_t^Y, H_{t-1})$. Furthermore, by (A4') $(\varepsilon_t^Y, H_{t-1}) \perp\!\!\!\perp \mathcal{I}_t$. Combining this (and using measurability of g) we obtain

$$Y_t - \beta X_{t-1} - \alpha_{Y,Y} Y_{t-1} \perp\!\!\!\perp \mathcal{I}_t.$$

Thus $\mathbb{E}[(Y_t - \beta X_{t-1} - \alpha_{Y,Y} Y_{t-1}) \mathcal{I}_t^\top] = 0$ for $\alpha = \alpha_{Y,Y}$, and part (i) follows.

Part (ii) follows because the moment equation $\mathbb{E}[(Y_t - \beta X_{t-1} - \alpha_{Y,Y} Y_{t-1}) \mathcal{I}_t^\top] = 0$ is the same as in Theorem 7, and so the rank requirement for identifiability is also the same.

To show part (iii), let $\bar{\mathbf{X}}_{t-1} := [\mathbf{X}_{t-1}^\top, \mathbf{Y}_{t-1}^\top]^\top$. We have to show that

$$\hat{\mathbb{E}}[\mathbf{Y}_t \mathcal{I}_t^\top] W \hat{\mathbb{E}}[\mathcal{I}_t \bar{\mathbf{X}}_{t-1}^\top] \left(\hat{\mathbb{E}}[\bar{\mathbf{X}}_{t-1} \mathcal{I}_t^\top] W \hat{\mathbb{E}}[\mathcal{I}_t \bar{\mathbf{X}}_{t-1}^\top] \right)^{-1} \xrightarrow{P} \gamma$$

with $\gamma := [\beta, \alpha_{Y,Y}]$. (A2') ensures convergence of the empirical moments to population moments, and using Slutsky's Theorem in combination with the expression for Y_t , the statement follows as in the proof of Theorem 5. ■

E.14 Proof of Proposition 13

Proposition 13. *Consider a VAR(1) process $S = [S_t]_{t \in \mathbb{Z}}$ satisfying (A2). Let β be the causal effect from X_t to Y_{t+1} , and let, for an arbitrary $m, \ell \in \mathbb{N}$, $(\alpha_{Y,X}, \alpha_{Y,Y})$ be the population vector of coefficients when regressing $Y_{s+1} - \beta X_s$ on $\{X_{s-k}, k = 1, \dots, m\} \cup \{Y_{s-j}, j = 0, \dots, \ell\}$. Then*

$$(\alpha_{Y,Y}, \beta, \alpha_{Y,X}) = \arg \min_{a,b,c} \mathbb{E}_{\text{do}(X_t:=x)} \left\{ Y_{t+1} - \sum_{j=0}^{\ell} a_j Y_{t-j} - b X_t - \sum_{k=1}^m c_k X_{t-k} \right\}^2.$$

Proof Recall that β and $\alpha_{Y,Y}, \alpha_{Y,H}$ denote the causal effects from X_t, Y_t and H_t , respectively, to Y_{t+1} . We have

$$\begin{aligned}
 & \min_{a,b,c} \mathbb{E}_{do(X_t:=x)} \left(Y_{t+1} - \sum_{j=0}^{\ell} a_j Y_{t-j} - bX_t - \sum_{k=1}^m c_k X_{t-k} \right)^2 \\
 &= \min_{a,b,c} \mathbb{E}_{do(X_t:=x)} \left(\{ \beta X_t + \alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y \} - \sum_{j=0}^{\ell} a_j Y_{t-j} - bX_t - \sum_{k=1}^m c_k X_{t-k} \right)^2 \\
 &= \min_{a,b,c} \mathbb{E}_{do(X_t:=x)} (\beta X_t - bX_t)^2 \\
 &\quad + \mathbb{E}_{do(X_t:=x)} (\beta X_t - bX_t) \left(\alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right) \\
 &\quad + \mathbb{E}_{do(X_t:=x)} \left(\alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right)^2 \\
 &= \min_{a,b,c} \mathbb{E}_{do(X_t:=x)} (\beta X_t - bX_t)^2 \\
 &\quad + \mathbb{E}_{do(X_t:=x)} (\beta X_t - bX_t) \left(\alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right) \\
 &\quad + \mathbb{E} \left(\alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right)^2 \\
 &= \min_{a,b,c} \mathbb{E}_{do(X_t:=x)} (\beta X_t - bX_t)^2 + (\beta x - bx) \mathbb{E} \left(\alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right) \\
 &\quad + \mathbb{E} \left(\alpha_{Y,Y} Y_t + \alpha_{Y,H} H_t + \varepsilon_{t+1}^Y - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right)^2 \\
 &= \min_{a,c} \mathbb{E} \left(Y_{t+1} - \beta X_t - \sum_{j=0}^{\ell} a_j Y_{t-j} - \sum_{k=1}^m c_k X_{t-k} \right)^2.
 \end{aligned}$$

Here, the third and fourth equality signs hold because the joint distribution of the variables $H_t, \varepsilon_{t+1}^Y, Y_t, \dots, Y_{t-\ell}$, and X_{t-1}, \dots, X_{t-m} is the same under the observational and the interventional distribution – as the variables are all non-descendants of X_t . Further, the minimum is obtained for $b = \beta$ and a and c being the coefficients after (linearly) projecting $Y_{t+1} - \beta X_t$ on the space spanned by $Y_t, \dots, Y_{t-\ell}$ and X_{t-1}, \dots, X_{t-m} . \blacksquare

E.15 Proof of Proposition 18

Proposition 18 (Identification with conditioning set and time inhomogeneous effect). *Consider a VAR(1) process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2*), ensuring that all quantities below are well-defined. Let either $\mathcal{B}_t := \{I_{t-3}\}$ or $\mathcal{B}_t := \{I_{t-3}, X_{t-2}, Y_{t-1}\}$ and let $b := \mathbb{E}[\beta_t]$ be the expectation of β_t . If $\beta_t \perp\!\!\!\perp (X_{t-1}, I_{t-2}, \mathcal{B}_t)$, the following two statements hold. (i) b satisfies the CIV moment*

condition $\mathbb{E}[\text{cov}(Y_t - bX_{t-1}, I_{t-2} | \mathcal{B}_t)] = 0$. (ii) Furthermore, if $\mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)]$ has rank d_X , then b is identified by $\text{CIV}_{X_{t-1} \rightarrow Y_t}(I_{t-2} | \mathcal{B}_t)$.

Proof The full-time graph of S is the same as that in Theorem 6, and therefore it follows by the same proof as of Theorem 6 that $\mathbb{E}[\text{cov}(Y_t - \beta_t X_{t-1}, I_{t-2} | \mathcal{B}_t)] = 0$. From this, one can derive:

$$\begin{aligned} 0 &= \mathbb{E}[\text{cov}(Y_t - \beta_t X_{t-1}, I_{t-2} | \mathcal{B}_t)] \\ &= \mathbb{E}[\text{cov}(Y_t, I_{t-2} | \mathcal{B}_t)] - \mathbb{E}[\text{cov}(\beta_t X_{t-1}, I_{t-2} | \mathcal{B}_t)] \\ &= \mathbb{E}[\text{cov}(Y_t, I_{t-2} | \mathcal{B}_t)] - \mathbb{E}[\beta_t] \mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)] \\ &= \mathbb{E}[\text{cov}(Y_t - \mathbb{E}[\beta_t] X_{t-1}, I_{t-2} | \mathcal{B}_t)]. \end{aligned}$$

Here the third equality follows because

$$\begin{aligned} \mathbb{E}[\text{cov}(\beta_t X_{t-1}, I_{t-2} | \mathcal{B}_t)] &= \mathbb{E} \left[\mathbb{E}[\beta_t X_{t-1} I_{t-2}^\top | \mathcal{B}_t] - \mathbb{E}[\beta_t X_{t-1} | \mathcal{B}_t] \mathbb{E}[I_{t-2}^\top | \mathcal{B}_t] \right] \\ &= \mathbb{E} \left[\mathbb{E}[\beta_t | \mathcal{B}_t] \mathbb{E}[X_{t-1} I_{t-2}^\top | \mathcal{B}_t] - \mathbb{E}[\beta_t | \mathcal{B}_t] \mathbb{E}[X_{t-1} | \mathcal{B}_t] \mathbb{E}[I_{t-2}^\top | \mathcal{B}_t] \right] \\ &= \mathbb{E} \left[\mathbb{E}[\beta_t] \mathbb{E}[X_{t-1} I_{t-2}^\top | \mathcal{B}_t] - \mathbb{E}[\beta_t] \mathbb{E}[X_{t-1} | \mathcal{B}_t] \mathbb{E}[I_{t-2}^\top | \mathcal{B}_t] \right] \\ &= \mathbb{E}[\beta_t] \mathbb{E}[\text{cov}(X_{t-1}, I_{t-2} | \mathcal{B}_t)], \end{aligned}$$

where we use the independence $\beta_t \perp\!\!\!\perp (X_{t-1}, I_{t-2}, \mathcal{B}_t)$ to conclude $\beta_t \perp\!\!\!\perp (X_{t-1}, I_{t-2}) | \mathcal{B}_t$ and $\beta_t \perp\!\!\!\perp X_{t-1} | \mathcal{B}_t$ (using the ‘Weak Union’ property, Lauritzen (1996)) and $\beta_t \perp\!\!\!\perp \mathcal{B}_t$.

This shows statement (i), and statement (ii) follows by the same arguments as the proof of Theorem 6. \blacksquare

E.16 Proof of Proposition 19

Proposition 19 (Identification with nuisance regressor and time inhomogeneous effect). *Consider a VAR(1) process $S = [I_t^\top, X_t^\top, H_t^\top, Y_t^\top]_{t \in \mathbb{Z}}^\top$ satisfying (A2*), ensuring that all quantities below are well-defined. Let $\mathcal{I}_t := \{I_{t-2}, \dots, I_{t-m-1}\}$ for an $m \geq 1$ and $\mathcal{Z}_t := \{Y_{t-1}\}$ and let $b := \mathbb{E}[\beta_t]$ be the expectation of β_t . If $\beta_t \perp\!\!\!\perp (X_{t-1}, \mathcal{I}_t)$, the following two statements hold. (i) There exists $\alpha \in \mathbb{R}$ such that b satisfies the NIV moment condition $\mathbb{E}[\text{cov}(Y_t - bX_{t-1} - \alpha \mathcal{Z}_t, \mathcal{I}_t)] = 0$. (ii) Further, if $\mathbb{E}[[X_{t-1}^\top, \mathcal{Z}_t^\top]^\top \mathcal{I}_t^\top]$ has rank $d_X + d_Y$, b is identified by $\text{NIV}_{X_{t-1} \rightarrow Y_t}(\mathcal{I}_t, \mathcal{Z}_t)$.*

Proof The full-time graph of S is the same as that in Theorem 7, and therefore it follows by the same proof as of Theorem 7 that there exists α such that $\text{cov}(Y_t - \beta_t X_{t-1} - \alpha \mathcal{Z}_t, \mathcal{I}_t) = 0$. From this, one can derive:

$$\begin{aligned} 0 &= \text{cov}(Y_t - \beta_t X_{t-1} - \alpha \mathcal{Z}_t, I_{t-2}) \\ &= \text{cov}(Y_t - \alpha \mathcal{Z}_t, I_{t-2}) - \text{cov}(\beta_t X_{t-1}, I_{t-2}) \\ &= \text{cov}(Y_t - \alpha \mathcal{Z}_t, I_{t-2}) - \mathbb{E}[\beta_t] \text{cov}(X_{t-1}, I_{t-2}) \\ &= \text{cov}(Y_t - \mathbb{E}[\beta_t] X_{t-1} - \alpha \mathcal{Z}_t, I_{t-2}). \end{aligned}$$

Here, the third equality follows because

$$\begin{aligned} \text{cov}(\beta_t X_{t-1}, I_{t-2}) &= \mathbb{E}[\beta_t X_{t-1} I_{t-2}^\top] - \mathbb{E}[\beta_t X_{t-1}] \mathbb{E}[I_{t-2}^\top] \\ &= \mathbb{E}[\beta_t] \mathbb{E}[X_{t-1} I_{t-2}^\top] - \mathbb{E}[\beta_t] \mathbb{E}[X_{t-1}] \mathbb{E}[I_{t-2}^\top] \\ &= \mathbb{E}[\beta_t] \text{cov}(X_{t-1}, I_{t-2}). \end{aligned}$$

This shows statement (i) , and statement (ii) follows by the same arguments as the proof of Theorem 7. ■