

Identifying Elasticities in Autocorrelated Time Series Using Causal Graphs

Silvana Tiedemann^{*1}, Jorge Sanchez Canales¹, Felix Schur², Raffaele Sgarlato¹, Lion Hirth¹, Oliver Ruhnau³, Jonas Peters²

¹Centre for Sustainability, Hertie School

²Department of Mathematics, ETH Zurich

³Department of Economics and Institute of Energy Economics, University of Cologne

September 19th, 2024

Abstract

The price elasticity of demand can be estimated from observational data using instrumental variables (IV). However, naive IV estimators may be inconsistent in settings with autocorrelated time series. We argue that causal time graphs can simplify IV identification and help select consistent estimators. To do so, we propose to first model the equilibrium condition by an unobserved confounder, deriving a directed acyclic graph (DAG) while maintaining the assumption of a simultaneous determination of prices and quantities. We then exploit recent advances in graphical inference to derive valid IV estimators, including estimators that achieve consistency by simultaneously estimating nuisance effects. We further argue that observing significant differences between the estimates of presumably valid estimators can help to reject false model assumptions, thereby improving our understanding of underlying economic dynamics. We apply this approach to the German electricity market, estimating the price elasticity of demand on simulated and real-world data. The findings underscore the importance of accounting for structural autocorrelation in IV-based analysis.

Keywords: instrumental variables; elasticity of demand; directed acyclic graphs; simultaneous equation; structural causal model; autocorrelation; causal inference

^{*}Corresponding author (tiedemann@hertie-school.org)

¹Replication code available at https://github.com/jscanales/elasticities_with_graphs

Contents

1	Introduction	1
2	Causal Representation of Autocorrelated Equilibrium Systems	4
2.1	Equilibrium System with Inertial Demand (Model I)	5
2.2	Alternative Demand Equations	6
3	Graph-based IV Estimation in Equilibrium Models	7
3.1	Conditional IV Estimation in Time Series	8
3.2	Consistent Estimation of Price Elasticities	10
4	Simulating electricity market data	12
5	Estimation of the Price Elasticity of Hourly German Electricity Demand	15
6	Conclusion	19
7	Funding	20
A	Causal Models for Time Series	25
A.1	Structural Causal Models for Time Series	25
A.2	Directed Acyclic Graphs for Time Series	25
A.3	Graph marginalization	26
B	Details of the two-stage-least-squares estimator	26
C	Additional information on Section 4	26
C.1	Additional autocorrelation and error plots for conditional IV	26
C.2	Additional indicator plots for Model II and Model III	27
C.3	Divergence between CIV estimates for simulated data	27
D	Additional information on Section 5	28
D.1	Data	28
D.2	Sensitivity analysis for the price elasticity of German electricity demand	29

1 Introduction

Estimating the price elasticity of demand from observational time series is a fundamental problem in econometrics. Observable prices and quantities are determined simultaneously based on the underlying economic primitives, the demand and supply equations. Identifying the price elasticity of demand from data, therefore, requires non-trivial methodology, such as using credible supply-shifting instrumental variables (IVs) (Wright, 1928; Tinbergen, 1940; J. D. Angrist, Graddy, and Imbens, 2000).

Even simple time dependencies in the economic relationships that cause the data to be autocorrelated, however, may make naive IV estimators inconsistent (e.g., Thams et al., 2022). For simulated data (with known ground truth), Figure 1 shows the importance of appropriately considering time dynamics. Here, demand structurally depends on its past, and an IV estimator that neglects these time dependencies overestimates the elasticity by up to an order of magnitude. The figure also shows other IV estimators that are consistent despite the time dependence. So, how do we determine the consistency of IV estimators when time series are autocorrelated?

To find consistent estimators for causal effects, the statistics and computer science communities have developed a framework based on structural and graphical causal models (Pearl, 2009; Spirtes, Glymour, and Scheines, 2000). In these models, each node represents a variable, and a directed edge indicates a direct causal influence. Assuming the Markov condition (Lauritzen, 1996), graphical criteria such as d-separation can be employed to determine conditional independencies (Brito and Pearl, 2002; Henckel, Buttenschoen, and Maathuis, 2023). Checking d-separation between two nodes works by assessing whether all paths between these nodes are blocked by a conditioning set and can be used to aid in the identification of valid adjustment sets or conditional instruments. For example, an IV estimator may be inconsistent if an unblocked path exists between the instrument and the dependent variable that does not go via the endogenous variable (see Section 3.1 for details). Recently, Thams et al. (2022) have extended the graphical criteria for valid conditional IV estimators to infinite causal time series graphs (Peters, Janzing, and Schölkopf, 2013), enabling the identification of valid conditional IV estimators in settings with time dependence.

To exploit the power of inference via causal time graphs for estimating the price elasticity of demand, we need to find an acyclic representation of a system at equilibrium. In Section 2, we do so by representing the equilibrium condition using an unobserved error term, which leads to a model with unobserved confounding. We then apply the graphical reasoning to three different demand dynamics and, based on their corresponding causal time series graphs, argue under which circumstances the naive IV estimator fails to be consistent.

In Section 3.2, we use the graph-based criteria to derive valid conditional IV estimators for each model. These estimators are based on two different ideas. First, blocking paths via conditioning: Here, the estimators achieve consistency by including lagged terms in the conditioning set; this approach leads to a class of estimators, which includes, for example, the lag-augmented local projection IV estimator (Stock and Watson, 2018; Montiel Olea and Plagborg-Møller, 2021; Montiel Olea, Plagborg-Møller, et al., 2024) which is commonly used in macro-economics. Second, estimating nuisance effects: This approach, developed by Thams et al. (2022), simultaneously estimates effects that are not of primary interest and whose estimates are afterward disregarded.

ELASTICITIES WITH GRAPHS

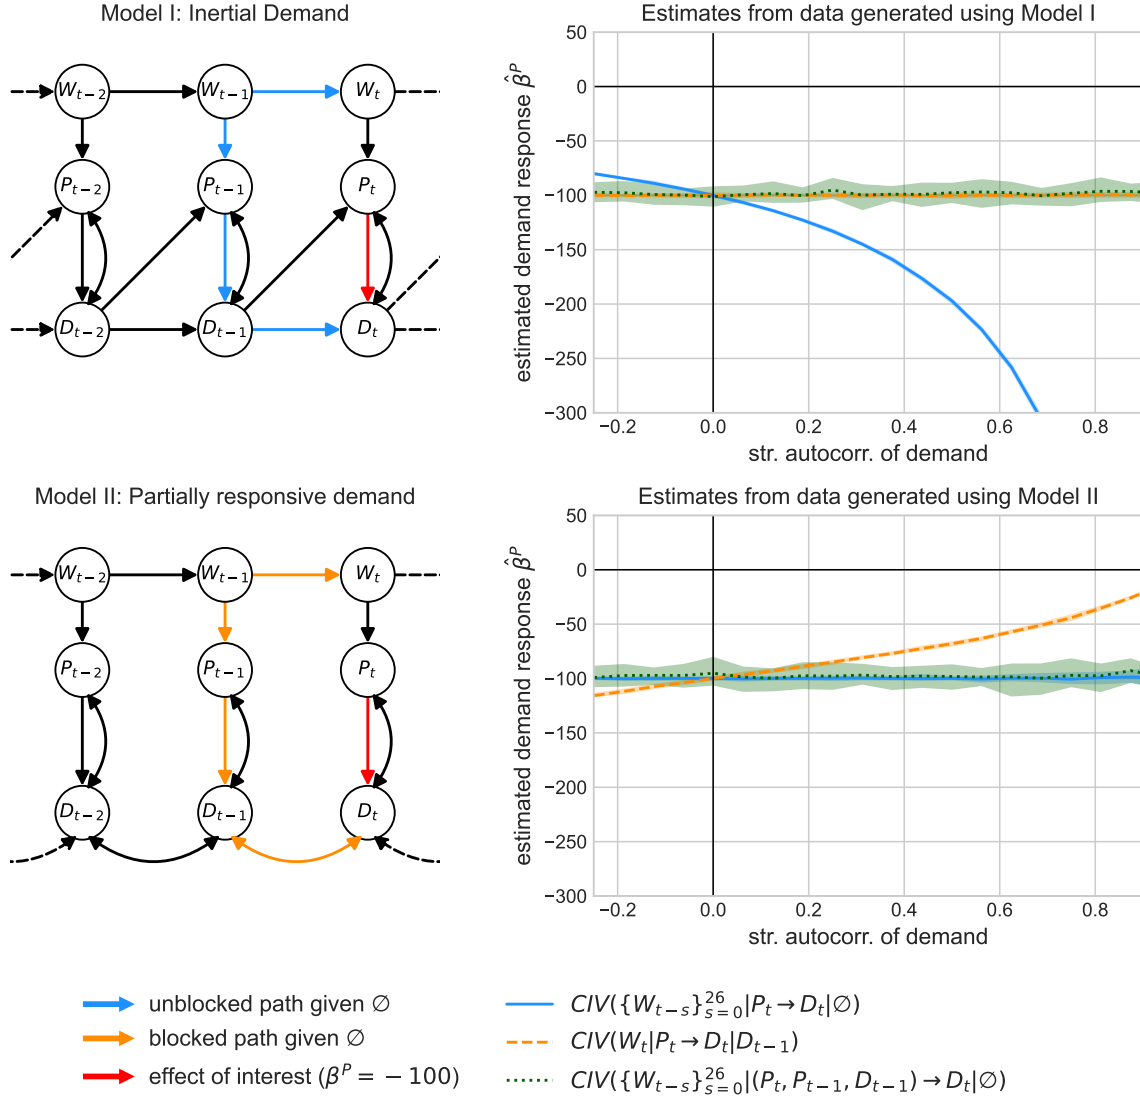


Figure 1: The effect of time dynamics on the validity of IV estimators. The left column shows the causal full time graphs of two underlying data-generating processes; the causal effect of interest is $P_t \rightarrow D_t$ (red edge). Top. In Model I, demand is inertial as D_t structurally depends on D_{t-1} . The blue path is unblocked if W_t is used as an instrument without additional controls. The naive IV estimator is, therefore, generally biased. The right column shows the corresponding estimates $\hat{\beta}^P$ for varying strength of structural autocorrelation in demand. Bottom. In Model II, the price-responsive part of demand is not autocorrelated; only the residual demand is. The orange path is now blocked given \emptyset but unblocked given D_{t-1} : D_{t-1} is a collider in the terminology of causal graphs. Here, the naive IV estimator is consistent. The nuisance IV estimator introduced in this paper (shown in green) is valid under both models.

The proposed graphical approach allows us to systematically analyze the validity of several estimators. We further argue that observing significant differences between the estimates of presumably valid estimators can help to reject false model assumptions, thereby improving our understanding of economic dynamics. If the researcher’s assumptions encoded in the structural model are correct, two valid estimators should yield comparable results (up to finite sample errors). We argue that a statistically significant difference (i.e., non-overlapping confidence intervals) between two estimates that were considered valid ex-ante should lead the researcher to reject her model assumptions. This may be seen as an alternative approach to model validation using the reductionist approach (Hendry, 1995).

The above findings allow us to study the problem of estimating the price elasticity of demand in the electricity market using weather-derived instruments, both on simulated data (Section 4) and on real-world data (Section 5). Estimating the price elasticity of electricity demand is important to manage the transition towards a carbon-neutral economy. Weather-dependent supply by renewables replaces flexible fossil fuel-based power plants. The price of electricity becomes highly dependent on the weather and more volatile. If demand reacts to this high-frequency price signal, less (fossil fuel-based) capacity is needed as backup facilities, and, therefore, the overall system cost decreases. Hence, getting a reliable estimate for the elasticity of demand is essential for efficient and reliable system planning. Furthermore, the application models data using high-frequency time series, increasingly available for economic studies (Webel, 2022). These data differ from classical macroeconomic time series by exhibiting high autocorrelation and seasonal correlation, making it crucial to consider dynamic effects, and allowing numerous lagged terms to be included in the conditioning set without sacrificing analytical robustness.

Related literature. There is an ongoing discussion about whether (and if so, how) causal graphs can help to solve problems in economics. Proponents argue that causal graphs representing structural causal models (SCMs) allow for the derivation of valid non-parametric estimators in a transparent and straightforward manner (Spirtes, Glymour, and Scheines, 2000; Pearl, 2009; Pearl and Mackenzie, 2020; Hünermund and Bareinboim, 2023), as exemplified by the derivation of the instrumental variable estimator by Wright (1928) for the price elasticity of demand for vegetable oil via graph-based path analysis.¹ Opponents point out that their utility is limited to problems that can be expressed in terms of acyclic graphs and argue that this precludes a broad class of econometric problems represented by simultaneous equations such as textbook examples of supply and demand. They further argue (see, e.g., Imbens, 2020) that using causal graphs does not solve problems that could not be solved using alternative frameworks such as the potential outcomes framework (Neyman, 1923; Holland, 1986; J. D. Angrist, Graddy, and Imbens, 2000; Rubin, 2005; J. Angrist and Pischke, 2009).

Despite the debate, graphs representing structural equations are not alien to econometrics, especially not to identification in time series. For example, Tinbergen (1940) and Wold (1964) have graphically represented causal dependencies between variables in economic models over time, using what they call ‘arrow scheme’. Arguably, they have not been widely adopted, as they were inherently tied to what Wold (1964) coined the ‘causal chain model’, which posits that economic processes are recursive. In this context, recursive means that each variable of the system depends unidirectionally on the preceding variables in a sequential manner. The

¹See Cunningham (2021) for a historical account of the derivation of the estimator.

arrow schemes thus conflict with the prevailing notion that economic systems are generally in equilibrium and should be modeled using simultaneous equations (see, e.g., Christ, 1994), where variables are determined jointly and can influence each other simultaneously. Our approach is different as it proposes an acyclic representation of a system at equilibrium.

Causal time graphs also speak to the literature on identifying dynamic causal effects in macroeconomic time series. Macroeconomic applications focus on estimating dynamic causal effects, known as impulse response functions at different horizons h , that is, the response of a (macro)economic system at time $t+h$ to a shock at time t . A recent and popular approach for identification is to use a local projection IV estimator (Jordà, Schularick, and Taylor, 2015; Ramey and Zubairy, 2018; Stock and Watson, 2018), which is essentially a sequential multivariate IV regression to identify the effect of price on demand at different time horizons, i.e., of $P_t \rightarrow D_t$, $P_t \rightarrow D_{t+1}$, \dots , $P_t \rightarrow D_{t+h}$. The lag-augmented local projection IV estimator (Plagborg-Møller and Wolf, 2021; Montiel Olea, Plagborg-Møller, et al., 2024) can be represented in the CIV-notation (see estimator #4 in Section 3). The conditional IV criteria provide a means to decide its validity: they offer a non-parametric proof in settings where local projection IV is consistent and motivate an example in which it is not (see Section 3.2). Furthermore, the graph-based analysis provides a straightforward explanation for recent findings in this literature, such as the relevance of the lead-lag exogeneity of the instrument (Stock and Watson, 2018), and the robustness to model misspecification (Plagborg-Møller and Wolf, 2021; Montiel Olea, Plagborg-Møller, et al., 2024).

Hence, formulating the problem of estimating the elasticity of demand using causal time graphs allows us to use simple, graphical criteria to determine which IV-based estimators are generally consistent and to avoid biases introduced by neglecting the time dependencies. Thus, we believe that this paper demonstrates the value of using graphs in econometrics.

2 Causal Representation of Autocorrelated Equilibrium Systems

This section links the economic equilibrium model of autocorrelated price and demand with structural causal models (Spirtes, Glymour, and Scheines, 2000; Pearl, 2009; Peters, Janzing, and Schölkopf, 2017) and (causal) graphical models (Lauritzen, 1996; Spirtes, Glymour, and Scheines, 2000; Pearl, 2009) in the context of time series data (e.g., Peters, Bauer, and Pfister, 2022).

We propose a formalization of the equilibrium relationship that models the system of simultaneous equations as a system with time-instantaneous hidden confounding between price and demand. It serves both as an observational model and as an interventional model (Peters, Bühlmann, and Meinshausen, 2016) considering interventions on the instrument (here wind) and price but not on demand. The formalization allows for a straightforward representation as a marginalized graph without cycles, a directed acyclic graph (DAG).

Section 2.1 introduces the general procedure by applying it to a model that postulates that demand is generated by an autoregressive process in which the current demand structurally depends on current prices and the immediate past of demand, i.e., demand features a direct structural autocorrelation (see Equation (1)). In Section 2.2 we introduce two alternative models with different assumptions on the intertemporal structural relations. Each model

generates a time series of autocorrelated demand (see the simulations in Section 4), but they have different implications for the validity of conditional instrumental variable estimators (see Section 3). While we do not claim that these three models are an exhaustive representation of possible demand structures (see Hendry, 1995, chapter 7), each of them builds on plausible assumptions (see, for example, the discussion on electricity markets in Section 5).

2.1 Equilibrium System with Inertial Demand (Model I)

We first model demand as a structural equation, where current demand D_t depends structurally on the price P_t , past demand D_{t-1} (we refer to this dependence as *direct structural autocorrelation* or simply *structural autocorrelation*), and an i.i.d. error term U_t^D . Assuming linearity, the demand equation then reads

$$\text{Model I: } \quad D_t := D_0 + \beta^P P_t + \beta^{D1} D_{t-1} + U_t^D. \quad (1)$$

Here, β^P is the slope of the demand equation and the causal (or structural) parameter of interest.

Similarly, for all $t \in \mathbb{Z}$ we assume for supply S_t :

$$S_t := S_0 + \gamma^P P_t + \gamma^W W_t + U_t^S, \quad (2)$$

where W_t is exogenous, U_t^S is an i.i.d. error term, and γ^P is the slope of the supply equation. The structural equation of W_t does not depend on P_t and is independent of U_t^S and U_t^D .

At equilibrium, demand equals supply, and the equilibrium price clears the market. This yields a structural constraint, the equilibrium condition for P_t : for all t we have

$$S_t = D_t. \quad (3)$$

The demand and supply equations (1) and (2) are structural in (P_t, D_{t-1}) and (P_t, W_t) , respectively, in that they hold under interventions on these variables. Using the notation of potential outcomes (Neyman, 1923; Rubin, 2005; Imbens and Rubin, 2015; Hernan and Robins, 2020) we can equivalently write them in the form $D_t(p) := D_0 + \beta^P p + \beta^{D1} D_{t-1} + U_t^D$ and $S_t(p, w) := S_0 + \gamma^P p + \gamma^W w + U_t^S$, where for fixed realisations of U_t^S and U_t^D , the quantities $D_t(p)$ and $S_t(p, w)$ denote the potential outcomes of demand and supply at time t for p and w , respectively.

From the structural equations for demand (1) and supply (2), and the equilibrium constraint (3) we get the following expression for prices at equilibrium:

$$\text{Model I: } \quad P_t = \frac{S_0 - D_0}{\beta^P - \gamma^P} + \frac{\gamma^W}{\beta^P - \gamma^P} W_t - \frac{\beta^{D1}}{\beta^P - \gamma^P} D_{t-1} + \frac{U_t^S - U_t^D}{\beta^P - \gamma^P}, \quad (4)$$

where we assume that $\beta^P - \gamma^P \neq 0$, a prerequisite of market clearing.

To construct estimators for β^P , we exploit that the equations for demand and price, together with the equations for W_t and U_t^D form a structural causal model (SCM) (Pearl, 2009; Bongers et al., 2021) or, more precisely, an adaption of an SCM which allows for time series models (Peters, Janzing, and Schölkopf, 2013; Peters, Bauer, and Pfister, 2022). (As an SCM, the system then allows for a graphical representation, with respect to which

the distribution satisfies the Markov condition (Pearl, 2009; Lauritzen et al., 1990).) In short, here, an SCM is a set of structural equations $X_t^j := f^j(\text{PA}(j, t), \epsilon_t^j)$, $j \in \{1, \dots, d\}$, where $\text{PA}(j, t)$ are the parents of X_t^j and ϵ_t^j , $t \in \mathbb{Z}, j \in \{1, \dots, d\}$, are jointly independent random error variables. In the example of Model I, when choosing D_t to be X_t^1 , we have $\text{PA}(1, t) = (P_t, D_{t-1}, U_t^D)$ and, when choosing P_t to be X_t^2 , we get $\text{PA}(2, t) = (W_t, U_t^D, D_{t-1})$. We provide a formal mathematical introduction to structural causal models for time series in Appendix A.1.

The SCM models the observed distribution and intervention distributions, where we intervene on W_t and P_t ; interventions on D_t are not modelled even though D_{t-1} appears in the price Equation (4) – the reason is that we have used (1) in deriving (4).² In this sense, the price equation is not structural in D_t . Thus, here, we use SCMs to construct consistent estimators and not to model all causal relationships. Using an SCM to model interventions on some but not all variables is similar in spirit to the decision-theoretic framework by Dawid (2021).

Importantly, as a result of the equilibrium constraint, the variable U_t^D appears in both the demand and price equations, yielding a confounder between price and demand. Since U_t^D is unobserved in practice, we have thus transformed a system of equilibrium relationships into a system with a hidden confounder U_t^D between P_t and D_t .

We now represent the structure of the SCM as a directed acyclic graph (DAG) \mathcal{G} ; each variable in the SCM corresponds to a vertex in the graph, and vertices in the graph are connected by directed edges. More precisely, we connect two nodes u and v by a directed edge from u to v if u is a causal parent of v . When we construct a graph of a time series, we obtain a full time graph \mathcal{G}_{FT} .³ Connecting the different steps above, the structural causal model defined by the equations (1)–(4) and $W_t = \beta^W W_{t-1} + U_t^W$ induces the marginalized⁴ DAG shown in Figure 2.

2.2 Alternative Demand Equations

We now present two alternative models that can also generate autocorrelated time series but with different structural dependencies in time. First, Model II (heterogeneous demand) assumes that aggregate demand consists of two types: the first, A_t , is price-sensitive but not directly structurally dependent on its own past, and the second, B_t , is price-insensitive and exhibits a direct structural dependency in time. In the case of electricity markets, this distinction by price sensitivity can be motivated by the fact that a part of demand is not exposed to high-frequency price fluctuations, such as retail consumers without real-time pricing. The corresponding structural equations read

$$\begin{aligned} \text{Model II: } \quad D_t &:= A_t + B_t = (A_0 + \beta^P P_t + U_t^A) + (B_0 + \beta^{B1} B_{t-1} + U_t^B) \\ &= D_0 + \beta^P P_t + \beta^{B1} B_{t-1} + U_t^D, \end{aligned} \tag{5}$$

²For example, if we were to consider an intervention on D_t and set it to x , we would obtain the price equation $P_t = (x - S_0 - \gamma^W W_t - U_t^S) / \gamma^P$.

³Appendix A.2 formally introduces directed acyclic graphs for time series.

⁴To simplify visualization, the marginalized DAG omits unobserved variables. Figure 11 in Appendix A.3 shows the non-marginalized graph without any dependencies in time. With respect to the operations we perform in this work, both versions of the graphs are equivalent; see also Richardson (2003).

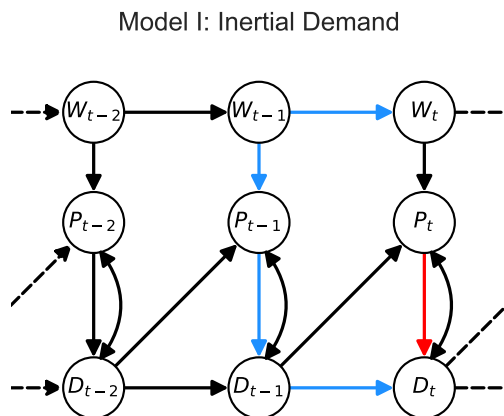


Figure 2: (Marginalized) graphical representation of Model I representing an inertial demand equation given by the structural equations (1) and (4). The causal effect of interest is highlighted in red: the effect of price on demand $P_t \rightarrow D_t$. The equilibrium constraint that supply equals demand results in bidirected edges $P_t \leftrightarrow D_t$. The blue path highlights why a naive IV estimator using W_t as an instrument generally fails to estimate the magnitude of $P_t \rightarrow D_t$ consistently: it does not account for the confounding $W_t \leftarrow W_{t-1} \rightarrow P_{t-1} \rightarrow D_{t-1} \rightarrow D_t$ (see also Section 3).

where $D_0 := A_0 + B_0$ and $U_t^D := U_t^A + U_t^B$. The price equation becomes

$$\text{Model II: } P_t = \frac{S_0 - D_0}{\beta^P - \gamma^P} + \frac{\gamma^W}{\beta^P - \gamma^P} W_t + \frac{\beta^{B1}}{\beta^P - \gamma^P} B_{t-1} + \frac{U_t^S - U_t^D}{\beta^P - \gamma^P}. \quad (6)$$

Figure 3 (left) depicts the corresponding graph. Second, Model III introduces another type of time dependence. In a high-frequency time series, we expect a potential for demand optimization over time because consumption shortly before t may be a substitute for consumption at time t . Consumers may respond to price differentials and shift their demand over time, such as deferring charging an electric vehicle. Figure 3 (right) shows a simplified cross-price elasticity in response to the realization of the price at time $t-1$. The corresponding equations are

$$\text{Model III: } D_t := D_0 + \beta^P P_t + \beta^{P1} P_{t-1} + U_t^D \quad (7)$$

and

$$\text{Model III: } P_t = \frac{S_0 - D_0}{\beta^P - \gamma^P} + \frac{\gamma^W}{\beta^P - \gamma^P} W_t + \frac{\beta^{P1}}{\beta^P - \gamma^P} P_{t-1} + \frac{U_t^S - U_t^D}{\beta^P - \gamma^P}. \quad (8)$$

3 Graph-based IV Estimation in Equilibrium Models

Next, we discuss sufficient conditions for valid instrumental variable (IV) estimation in equilibrium systems with a time series structure (Section 3.1). To do so, we exploit graphical conditions for validity of conditional IV (CIV) estimators (Brito and Pearl, 2002; Henkel,

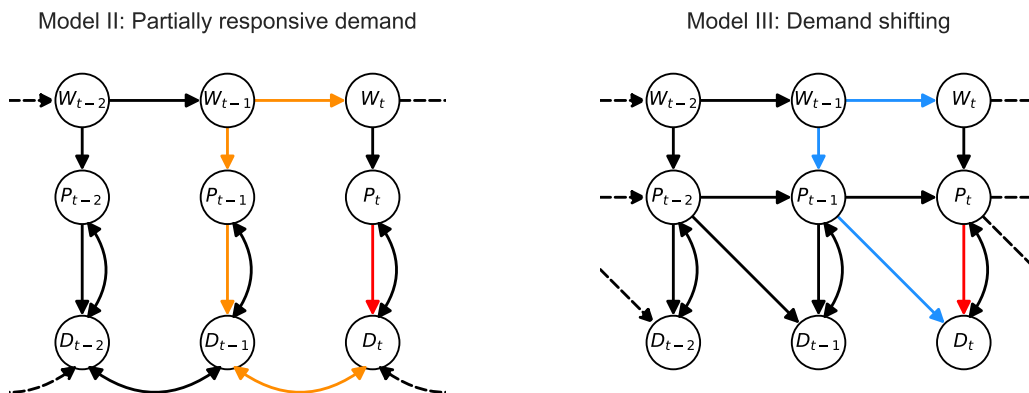


Figure 3: (Marginalized) graphical representation of alternative demand equations. Left: Model II (Partially Responsive Demand): Aggregate demand can be divided into two types as in (5), one being price-exposed and not autocorrelated, one not exposed to prices and with inertia. Right: Model III (Demand shifting): Demand depends on the price in two time steps, current and lagged, as in (7). Thus, the model assumes that processes or systems can be shifted in time, similar to a simplified representation of a cross-price response. The blue path highlights why a naive IV estimator fails to estimate the causal effect from price on demand in the presence of cross-price elasticity (CIV1 is violated, see Section 3); by contrast, the orange path does not yield a similar violation.

Buttenschoen, and Maathuis, 2023), which have recently been extended to time series settings (Thams et al., 2022). In Section 3.2 we apply these CIV principles to the three models introduced in Section 2 and present several valid CIV estimators. We argue that discrepancies between them can inform researchers about the validity of the underlying models and assumptions.

3.1 Conditional IV Estimation in Time Series

For constructing valid estimators we make use of the concept of d-separation (short for directed separation) (Pearl, 2009). D-separation is a purely graphical criterion connected to conditional independences of random variables via the Markov condition (Lauritzen, 1996). Pearl (2009)’s original formulation considers finite graphs; here, we use a version that is adapted to (marginalized) infinite time series graphs.

Definition 3.1 (Pearl’s d-separation (Pearl, 2009)). Let \mathcal{G} be the marginalized version of an acyclic directed full time graph (Peters, Janzing, and Schölkopf, 2013), which contains a node for each variable and time point, with vertices⁵ $V = \mathbb{Z} \times \{1, \dots, d\}$, and let i_1, \dots, i_m be a path between vertices i_1 and i_m in \mathcal{G} . We say that the path i_1, \dots, i_m is *blocked by the set of variables* $S \subseteq V \setminus \{i_1, i_m\}$ if there exists $k \in \{2, \dots, m-1\}$ such that either

- (i) $i_k \in S$ and i_k is not a collider on the path i_1, \dots, i_m , or

⁵Here the \mathbb{Z} component encodes time, e.g., D_t is a vertex denoted as $(t, 1)$.

- (ii) i_k is a collider⁶ on the path and $(\{i_k\} \cup \text{DE}(i_k)) \cap S = \emptyset$, where $\text{DE}(i_k)$ are the descendants of i_k ;

if a path is not blocked by S , we say it is *unblocked, given S* . For three mutually disjoint sets of vertices $A, B, S \subset V$, we say that A and B are *d -separated by S* if all paths between A and B are blocked by S . If this is the case, we write

$$A \perp\!\!\!\perp_{\mathcal{G}} B \mid S.$$

The concept of d -separation helps us to construct valid IV estimators in time series. The following statement is an adapted version of Theorems 5 and 8 of Thams et al., 2022. Let $d, d' \in \mathbb{N}$, $d \leq d'$ and let $V' = \mathbb{Z} \times \{1, \dots, d'\}$. Consider a linear SCM for time series over variables $X_{V'}$ and let \mathcal{G} be the induced marginalized version of an acyclic directed full time graph with vertices $V = \mathbb{Z} \times \{1, \dots, d\}$. Let⁷ $\mathcal{I}, \mathcal{X}, \mathcal{B}, \{Y\} \subseteq V$ have zero means and finite second moments and let $\beta \in \mathbb{R}^{|\mathcal{X}|}$ be the causal coefficient with which $X_{\mathcal{X}}$ enters the structural equation for X_Y , that is,

$$X_Y = \beta^T X_{\mathcal{X}} + \gamma^T X_K + \epsilon^Y, \tag{9}$$

for some variables $X_K \subseteq X_{V'} \setminus X_{\mathcal{X}}$ (some of the components of β can be zero, so not all variables in $X_{\mathcal{X}}$ have to be parents of X_Y).

One then considers the following three requirements on $\mathcal{I}, \mathcal{X}, \mathcal{B}$ and Y , the CIV criteria:

- (CIV1) \mathcal{I} and Y are d -separated given \mathcal{B} in the graph $\mathcal{G}_{\mathcal{X} \nrightarrow Y}$, which denotes the graph obtained when removing all direct edges from \mathcal{X} to Y from \mathcal{G} .
- (CIV2) \mathcal{B} is not a descendant of $\mathcal{X} \cup \{Y\}$ in \mathcal{G} .
- (CIV3) the matrix $\mathbb{E}[\text{cov}(X_{\mathcal{X}}, X_{\mathcal{I}} | X_{\mathcal{B}})]$ has rank $|\mathcal{X}|$, that is, full row rank.

(CIV1) and (CIV2) are properties of the underlying causal graph, (CIV3) is a property of the induced distribution.

Assume we have observations at time points $1, \dots, T$. If an estimator is based on $\mathcal{I}, \mathcal{X}, \mathcal{B}$ such that requirements (CIV1), (CIV2) and (CIV3) are met, we call it a *valid CIV estimator*. In this case,

$$\hat{\beta} := \arg \min_b \|\text{cov}(X_Y - b^T X_{\mathcal{X}}, X_{\mathcal{I}} | X_{\mathcal{B}})\|, \tag{10}$$

is a consistent estimator for β if $T \rightarrow \infty$ (for a detailed definition of $\hat{\beta}$ and its closed form solution see Appendix B). We call $\hat{\beta}$ the *conditional instrumental variable (CIV) estimator* and denote it by $\text{CIV}(\mathcal{I} \mid \mathcal{X} \rightarrow Y \mid \mathcal{B})$.

If at least one of the conditions is not met, then we call the estimator invalid and, in general, it may not be consistent (Henckel, Buttenschoen, and Maathuis, 2023) (but exceptions exist). In the special case, however, that (CIV1) and (CIV2) are satisfied but (CIV3) is not (this is related to weak instrument settings), any estimation technique that comes with valid inference yields unbounded confidence sets containing the true causal effect with large probability. In this sense, it is therefore not necessary to test for the validity of (CIV3).

⁶Let \mathcal{G} be a marginalized version of an acyclic directed full time graph with vertices V . For such graphs we define v to be a *collider* on a path whenever two consecutive edges have arrowheads at v (for example, $u_1 \leftrightarrow v \leftarrow u_2$).

⁷Depending on the context, and slightly overloading notation, X_S for $S \subseteq V$ is either the set random variables $\{X_s \mid s \in S\}$ or the random vector $(X_{s_1}, \dots, X_{s_n})$, where s_i are in lexicographic order.

3.2 Consistent Estimation of Price Elasticities

We now apply the CIV criteria to derive valid estimators for the Models I–III presented in Section 2. For example, consider Model I from Section 2.1. The model is described by Equations (1) and (4), and the corresponding marginalized time graph in Figure 2. The naive IV estimator #1 that neglects time dependencies can be written as $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid \emptyset)$, that is, we set $X_{\mathcal{X}} = (P_t)$, $X_Y = \{D_t\}$, $X_{\mathcal{B}} = \emptyset$ and $X_{\mathcal{I}} = (W_t)$. We see that (CIV1) is not satisfied: After removing the edge $P_t \rightarrow D_t$ in Figure 2, we obtain the graph $\mathcal{G}_{P_t \not\rightarrow D_t}$. In this graph, there are several unblocked paths given \emptyset through W_{t-1} that connect the instrument W_t to the dependent variable D_t . An example is the blue path in Figure 2, $W_t \leftarrow W_{t-1} \rightarrow P_t \rightarrow D_{t-1} \rightarrow D_t$, given $X_{\mathcal{B}} = \emptyset$. W_t is therefore not d-separated from D_t .

We provide an overview of some available estimators and their validity in Table 1. We group CIV estimators along two main concepts. The first approach uses a selection of path-blocking variables as the conditioning set. Such a conditioning set can contain past realizations of the instrument W_t , the endogenous variable P_t , or the outcome variable D_t . For example, the estimator #3 $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid D_{t-1})$ contains only one lag of the outcome variable; it is valid for Model I. The estimator #2, $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid \{W_{t-s}\}_{s=1}^L)$, contains L lags of the instrument and is valid for Models I–III. Intuitively, by including the past realizations of the instrument in the conditioning set, we block the path starting with $W_t \leftarrow W_{t-k}$ for some $k \geq 1$. The estimator #2 is valid under a strictly more general model class: for example, it is even valid for any model yielding a (marginalized) graph, in which $P_t \rightarrow D_t$ is the only directed path between P_t and D_t , and the set of parents of W_t is a subset of $\{W_{t-1}, \dots, W_{t-L}\}$. This model class contains Models I, II, and III. We therefore consider #2 our benchmark estimator. The second approach relies on the concept of simultaneously estimating nuisance effects. Estimator #8, $\text{CIV}(\{W_{t-s}\}_{s=0}^L \mid (P_t, P_{t-1}, D_{t-1}) \rightarrow D_t \mid \emptyset)$, for example, adds two additional nuisance covariates to \mathcal{X} . It estimates a three-dimensional effect and later ignores two of its components.

Estimator	Model I Figure 2	Model II Figure 3 (left)	Model III Figure 3 (right)
Naive IV			
#1 $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid \emptyset)$	x	✓	x
Conditional IV			
#2 $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid \{W_{t-s}\}_{s=1}^L)$	✓	✓	✓
#3 $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid D_{t-1})$	✓	x	x
#4 $\text{CIV}(W_t \mid P_t \rightarrow D_t \mid \{W_{t-s}, P_{t-s}, D_{t-s}\}_{s=1}^L)$	✓	✓	✓
Nuisance IV			
#5 $\text{CIV}(\{W_{t-s}\}_{s=0}^L \mid (P_t, D_{t-1}) \rightarrow D_t \mid \emptyset)$	✓	✓	x
#6 $\text{CIV}(\{W_{t-s}\}_{s=0}^L \mid (P_t, P_{t-1}) \rightarrow D_t \mid \emptyset)$	x	✓	✓
#7 $\text{CIV}(\{W_{t-s}\}_{s=0}^L \mid (P_t, P_{t-1}) \rightarrow D_t \mid D_{t-1})$	✓	x	✓
#8 $\text{CIV}(\{W_{t-s}\}_{s=0}^L \mid (P_t, P_{t-1}, D_{t-1}) \rightarrow D_t \mid \emptyset)$	✓	✓	✓

Table 1: Estimators and their validity for the causal effect β^P from $P_t \rightarrow D_t$ for Models I–III

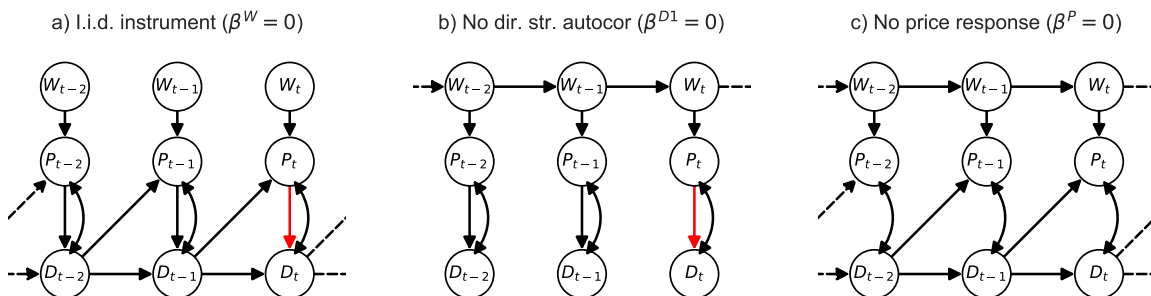


Figure 4: Graphical representation of special cases of Model I, for which naive IV is a valid estimator. Left: The instrument is i.i.d. i.e., not autocorrelated. Middle: The demand function is not autocorrelated. Right: The demand is not responsive to prices.

Special cases for a valid naive IV estimator. Even though the naive IV estimator does not necessarily satisfy (CIV1), there are special cases in which the naive estimator is valid, as visualized in Figure 4. These occur i) when the instrumental time series is i.i.d., or ii) when demand is not directly affected by its past realizations, or iii) when there is no response to prices, i.e., $\beta = 0$. In all three cases, no unblocked path exists (given \emptyset) from W_t to D_t through the past.

Local Projection IV Estimator. Estimator #4 in Table 1 is the lag-augmented local projection IV estimator⁸ from the macroeconomic literature on identifying impulse response functions (Stock and Watson, 2018; Montiel Olea and Plagborg-Møller, 2021). In CIV notation, the local projection IV estimator for the time horizon $h = 0$, i.e., the effect of $P_t \rightarrow D_t$, can be written as $\text{CIV}(W_t | P_t \rightarrow D_t | \{W_{t-s}, P_{t-s}, D_{t-s}\}_{s=1}^L)$; thus, the local projection IV estimator is a special case of a CIV estimator with lags of the instrumental, endogenous, and dependent variable in the conditioning set, since it also relies on the concept of blocking path to achieve validity. To achieve validity, however, including the lags of the instrument would be sufficient (see estimator #2). If, however, in Model I, there is additional hidden confounding between W_t and P_{t-1} , and D_t and P_{t-1} , #4 is invalid, while #2 would still be valid. Estimator #2 is valid across a larger model class.

Model validation. The existence of several valid IV estimators allows us to learn about the structure of the demand response. We can assume a particular structure of the demand equation, but we do not know ex-ante whether this assumption holds. The ability to derive multiple supposedly valid estimators for a structural model can be used to test this assumption. For example, if Model II describes the true time dependencies, the confidence intervals of all estimators but number #3 in Table 1 should overlap with high probability. This can be tested from observational data. If the confidence intervals do not overlap, one should reject Model II.

⁸The general formulation of the local projection IV estimator includes lags up to infinity (see e.g., Montiel Olea and Plagborg-Møller, 2021, chap. 3.3), which in practice translates into a 'large number of lags', which we represent by L .

4 Simulating electricity market data

We now evaluate the CIV estimators of Section 3.2 in a controlled environment motivated by the German electricity market (see Section 5). We empirically quantify the bias that arises when using invalid estimators and compare the empirical performances of the CIV estimators (including nuisance IV estimators). We also show that the autocorrelation of the response time series does not suffice to correct for the bias of invalid estimators.

Data generation. We generate multiple datasets containing three hourly time series: a supply-shifting instrumental time series, wind generation⁹, the equilibrium electricity demand,¹⁰ and the equilibrium electricity price. We model wind generation as an AR(L) process with autocorrelation coefficients derived from observed wind generation data in Germany in 2019 (see Appendix D.1 for data sources). The supply equation corresponds to Equation (2) with $\gamma^P = +500$ MW/(EUR/MWh). It is parameterized so that every unit of generated wind electricity is immediately offered to the market ($\gamma^W = 1$). The constant $S_0 = 25,000$ MWh/h implies that there is supply even if prices drop below zero, a typical characteristic of real-world electricity markets. The errors U_t^S are i.i.d. $\mathcal{N}(0, 1)$. Finally, we equate the supply and the model-specific demand equation (Equation (1), (5) or (7)) to obtain hourly equilibrium prices and quantities. Unless otherwise specified, the demand equation is parameterized as follows: the own price elasticity is $\beta^P = -100$ MW/(EUR/MWh), the error terms of the demand equation are i.i.d. $\mathcal{N}(0, 2000)$, and the constant of the demand series D_0 is chosen so that the resulting demand is stationary around the mean value of German electricity load (approximately 60 GWh/h).

CIV estimation. All CIV estimators are calculated with the Python package `linearmodels` (Sheppard and al., 2024), using the kernel correction for heteroscedasticity and autocorrelation robust errors (HAC).

Validity of estimators. We validate the theoretical findings of Section 3. In Figure 5, we apply the CIV estimators specified in Table 1 to the data simulated from Model I.¹¹ We find that the confidence sets of all valid estimators overlap (compare the overlap of the confidence intervals with the check marks for Model III on the right of Figure 5). Conversely, if one had assumed Model II, one would be forced to reject the associated structural assumptions since estimators #1 and #6 deviate from other valid estimators (e.g., #2 and #8). The simulation thus underscores our argument that comparing the overlap of several presumably valid estimators provides a means of rejecting incorrect structural assumptions.

Magnitude of bias. We quantify the bias caused by ignoring any direct structural autocorrelation. In linear SCMs, the bias (explicitly computed by Thams et al. (2022, Prop. 17, for a special case)) depends on the path coefficients along unblocked paths (which may be unblocked due to wrongfully including descendants of colliders in the conditioning set). Figure 6 shows the empirical percentage error of the naive IV estimator #1 applied to estimate

⁹For a discussion on wind generation being a valid instrument see Section 5.

¹⁰In the electricity market, equilibrium demand is also referred to as load or electricity consumption.

¹¹Appendix C provides simulation results for Models II in Figure 16 and for Model III in Figure 17.

ELASTICITIES WITH GRAPHS




Figure 5: CIV estimates $\hat{\beta}^P$ of the true demand response $\beta^P = -100$ MW/(EUR/MWh) when data is simulated based on Model I (Figure 2). The structural autocorrelation of demand is $\beta^{D1} = 0.7$. The columns on the right indicate the validity according to the CIV criteria (see Table 1).

the demand response for data generated from Model I. We vary the direct structural autocorrelations of the instrumental time series (β^W) and the demand series (β^{D1}). For low levels of instrument autocorrelation ($\beta^W < 0.2$), ignoring time dependencies leads to a moderate bias ($< 20\%$). However, the resulting bias is substantial for high-frequency time series. For example, in real-world data (see Section 5), the autocorrelation of wind generation (at lag 1) can have values above 0.9. Hence, even for small levels of direct structural autocorrelation of demand, this could lead to a bias of more than 25%. We observe the same pattern for other combinations of models and invalid estimators (see Figure 12).

Structural vs. observed autocorrelation. Knowing the observed autocorrelation of the demand is not sufficient for correcting for the bias of the naive IV estimator #1. In Figure 7, we plot the percentage error of estimator #1 against the observed autocorrelation of demand (which we denote by α^{D1} to distinguish it from the structural autocorrelation denoted by β^{D1}). The intensity of the color indicates the value of the respective structural coefficients. The figure shows that the same level of observed autocorrelation (at lag one) can result from different structural dependencies, yielding different estimation errors. Low levels of observed autocorrelation can result in a substantial bias of 100% (e.g., Model III at $\alpha^{D1} \approx 0.2$), and observed autocorrelation up to $\alpha^{D1} \approx 0.9$ can be associated with no bias at all if the data is generated by a different model (e.g., Model II). Consequently, knowing the observed autocorrelation of demand (at lag one) is insufficient for deriving a structural model or predicting a bias.

Empirical performance of different CIV estimators. We analyze the empirical performance of the three CIV estimators that are valid across a wide class of models: the nuisance IV estimator #8, the CIV estimator #2, and the CIV estimator #4, which corresponds to a lag-augmented local projection IV estimator. We do so based on three indicators: *Coverage*, which indicates how often the 95%-confidence interval of $\hat{\beta}^P$ contains the true value β^P ; *aver-*

ELASTICITIES WITH GRAPHS

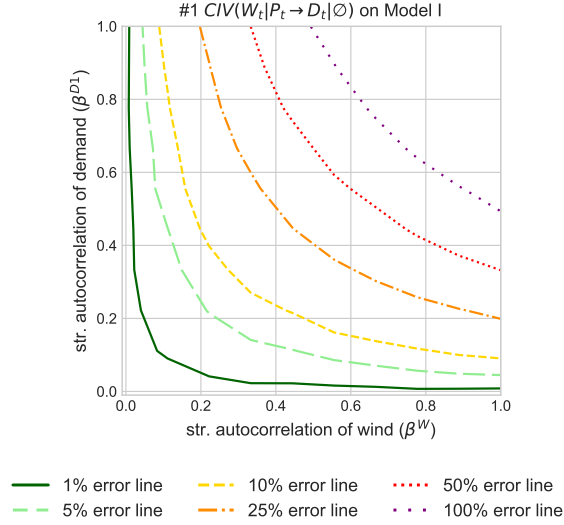


Figure 6: Isolines showing the average absolute percentage error of the point estimate as a function of the structural autocorrelation of wind (instrument) and the structural autocorrelation of demand (dependent variable). For the simulations, we divide both axes into ten equidistantly spaced autocorrelation coefficients between 0 and 1 each, based on which we run 20 simulations with five years of data ($T = 43,800$).




Figure 7: The percentage error of the naive IV estimator #1 is shown as a function of the observed autocorrelation of demand (α^{D1}). In each simulation we vary the structural dependencies: In Models I and II, the strengths of the structural autocorrelations, β^{D1} and β^{B1} , are both varied between -0.25 and 0.99 . In each simulation we vary the structural dependencies: In Models I and II, the strengths of the structural autocorrelations, β^{D1} and β^{B1} , are both varied between -0.25 and 0.99 . In Model III, β^{P1} is varied between -250 MW/(EUR/MWh) and $+250$ MW/(EUR/MWh). The color intensity increases in tandem with the aforementioned ranges. For each model, we conduct 400 simulations with a sample period of five years ($T = 43,800$).

ELASTICITIES WITH GRAPHS

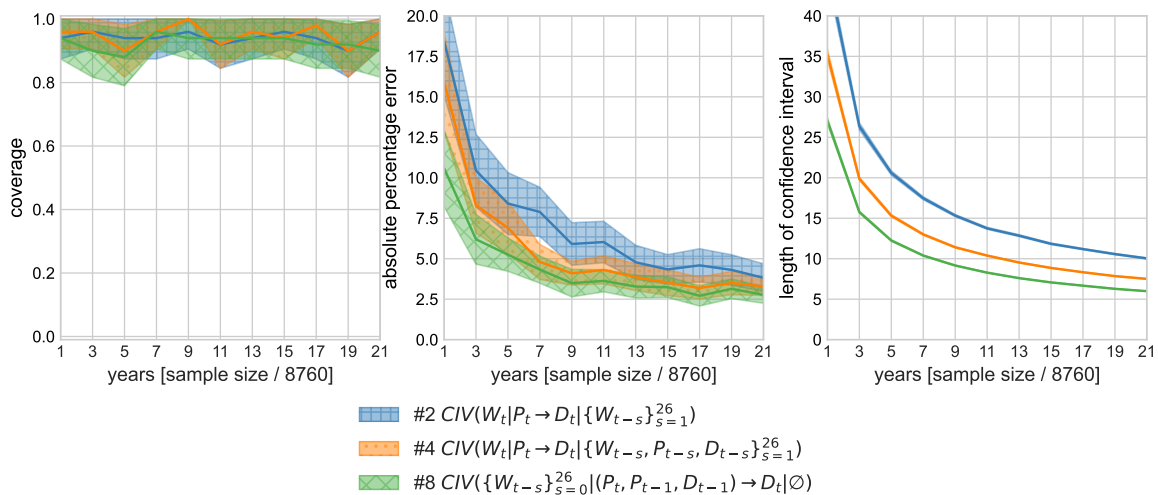


Figure 8: Evaluation of the statistical performance of conditional vs. nuisance IV estimators for Model I based on the performance indicators: (left) coverage, (middle) the average absolute percentage error, and (right) the length of the confidence interval. ($\beta^{D^1} = 0.9$.) The shaded areas represent the 95% confidence intervals. For each year, we run 50 simulations.

age absolute percentage error, which averages $100 \cdot |\hat{\beta}^P - \beta^P| / |\beta^P|$ over different simulations; and the *length of the confidence interval*.

Figure 8 plots the indicators as a function of the sample size for data generated from Model I. All estimators have coverage (left), confirming these approaches' validity. However, the nuisance IV estimator #8 has a smaller average absolute percentage error (middle) and yields shorter confidence intervals at smaller sample sizes (right) than both the conditional IV estimators of the most general class, #2, and the lag-augmented local projection IV estimator, #4. The benefit of having a large sample size becomes marginal beyond five years but holds regardless of the structural model (see Figure 14 and Figure 15). Thus, while the estimators that rely on blocking the blue path in Figure 2 by conditioning, #2 and #4, are conceptually simple and retain validity across a large class of models, the estimator that relies on estimating nuisance effects, #8, in the above setting outperforms them, thereby providing a useful alternative for smaller sample sizes.

5 Estimation of the Price Elasticity of Hourly German Electricity Demand

We estimate the own-price elasticity of the aggregate German electricity demand using the CIV estimators derived in Section 3. Furthermore, we discuss what the divergence between the estimators can reveal about the structure of the dynamic response.

Price elasticity of electricity demand. Historically, electricity demand was considered inelastic, and power plants adjusted their supply according to whatever demand occurred (Borenstein, 2002; Cramton, Ockenfels, and Stoft, 2013). Liberalization, digitalization, and

decarbonization over the last three decades have started to change that. After liberalization, it soon became clear that a low elasticity of demand inflates the ability of producers to abuse market power, and policymakers took measures to promote more flexible customer behavior, e.g., by increasing the exposure of consumers to real-time pricing (Borenstein, 2002; Fabra et al., 2021). In addition, as the role of wind and solar expands globally, price volatility increases, and the benefits of flexible demand become larger (Ciarreta, Pizarro-Irizar, and Zarraga, 2020; Maniatis and Milonas, 2022; Hosius et al., 2023). Policy scenario and system planning studies assume future demand to be increasingly price-elastic (e.g., ACER, 2020). Therefore, an accurate estimate of the price elasticity of electricity demand is important for reliable system planning and appropriate policy decisions.

The electricity market. The electricity system must always be in balance for physical reasons. The primary mechanism to equalize demand and supply is the day-ahead auction of the wholesale market. This is where generators, retail suppliers, and large industrial consumers submit their price-volume bids for each hour of the following day.¹² Because electricity storage is expensive and limited in capacity, wholesale electricity prices are much more volatile than other commodities (Ciarreta, Pizarro-Irizar, and Zarraga, 2020; Maniatis and Milonas, 2022; Hosius et al., 2023).

Data. We estimate the own-price elasticity on electricity market data between October 23, 2017, and December 31, 2020 (see Table 2 for summary statistics). Electricity demand in megawatt-hours per hour (MWh/h) is defined as consumption in the joint German-Luxembourg market. The electricity price in Euro per megawatt-hour (EUR/MWh) corresponds to the clearing price of the day-ahead auction on EPEX Spot, the largest marketplace for the market zone. Wind generation is in gigawatt-hours per hour (GWh/h). A detailed description of the data and sources can be found in Appendix D.1. Whenever we exclude hours and/or split the data set, we determine the relevant lags and then perform the split.

	Mean	Std. Dev.	Min	Median	Max	Skewness	Kurtosis
Consumption [GWh]	59.9	11.3	34.0	59.5	91.7	0.16	-0.7
Price [€/MWh]	38.0	17.6	-90.0	38.0	200.0	-0.35	3.8
Wind Generation [GWh]	14.1	10.1	0.1	11.6	46.1	0.87	-0.1

Table 2: Summary Statistics; $T = 27,072$ after removing holidays.

Identification. To overcome the endogeneity problem induced by the equilibrium condition, we use wind power generation as an exogenous time series. Wind speed and wind power generation are commonly used as supply shifting instruments because they are relevant and credibly exogenous (Bönte et al., 2015; Knaut and Paulus, 2017; Fabra et al., 2021; Hirth,

¹²The bid volume does not correspond to aggregate demand, and the bid curves do not correspond to the demand equation. This is because market participants can submit net bids (pool-based bidding) and have the option of trading on other marketplaces. Hence, the demand elasticity is not observable but must be estimated. However, the clearing price is informative about the equilibrium price (Knaut and Paulus, 2017).

Khanna, and Ruhnau, 2024). The assumption that the structural equation of wind power generation does not depend on the price of electricity is based on the argument that renewable generators have marginal costs close to zero. Additionally, most German wind power generators also receive a subsidy per unit of electricity produced, which induces an opportunity cost of not producing, even if prices are negative.¹³

We follow Hirth, Khanna, and Ruhnau (2024) in proposing that wind power is independent of the noise terms of demand and supply after conditioning on suitable covariates. While it is reasonable to assume that wind power does not affect demand in any way other than through the price of electricity, we expect the existence of common causes such as seasonality and weather events. We therefore condition on the following covariates: seasonal dummies (hour of the day, day of the week, and month of the year) and weather controls (sunlight, and heating and cooling degrees). Furthermore, we include other covariates to increase the precision of the estimate, namely commodity prices (natural gas, coal, and emission allowances), solar PV generation, and other calendar controls (school vacations by state and a variable for the last week of the year).

In Models I–III, we assume W_t to have a structural autocorrelation of one lag, and in that case, for many of the estimators, including a single time lag suffices. However, observed wind power exhibits a higher degree of autocorrelation (see Figure 18), which needs to be reflected in the construction of the different estimators (the argument remains the same). Given the observed autocorrelation pattern, we include up to 50 lags of wind power generation in the conditioning set of the CIV estimator #2 and #4, and as instruments for the nuisance IV estimators #4 to #8. We also observe a significant autocorrelation pattern in the demand time series (this is expected if there is indeed a causal effect from wind power generation on demand, but it may have other sources, too).

Models of electricity demand. Electricity demand features two main mechanisms that can imply structural correlations in time (for a comprehensive classification of demand response in the electricity market, see Albadi and El-Saadany, 2008). First, some electricity-consuming processes run for extended periods and cannot be switched on and off by the hour. This includes many industrial activities spanning an entire work shift and residential activities such as washing machines with multi-hour programs. Therefore, it is likely that the demand at time t depends on the load of the previous hours, as in our Model I. Second, some electricity-consuming processes can be shifted in time to exploit power price differentials. Such ‘load shifting’ may be done by scheduling industrial processes during low-price hours and by postponing charging electric vehicles, for example. Model III represents a simplified version of the shifting dynamic, where demand D_t also depends on the lagged price P_{t-1} .

Our alternative demand Model II assumes that processes or systems can be divided into two types: one that is price-exposed and reacts instantaneously and one that is inertial and not exposed to price variations. This model is motivated by the observation that most retail consumers are not exposed to real-time prices and that only the remaining consumers can respond to prices. For simplicity, we assume that these remaining consumers can regulate their processes individually for each hour.

¹³Negative prices represent approximately 2.5 percent of the observations. If wind generators stopped producing at negative prices, this would indicate that the instrument is not valid for that price range. However, we also observe wind generation at negative prices.

Additional assumptions about the demand equation. The models we consider make additional assumptions about the demand response. First, they posit that the functional form of the resulting demand equation depends linearly on price. To show that dynamics matter irrespective of the functional form, we include results for the assumption of an exponential relationship in Figure 9¹⁴. Second, the models assume that the elasticity is constant across hours and seasons. This implies that consumers who are active during the day have, on average, the same elasticity as consumers who are active at night. We investigate the relevance of this assumption in more detail in a robustness analysis in Appendix D.2.

Quantitative results. Figure 9 gives the quantitative results. The CIV estimators #2 and #4, and the nuisance IV estimator #8, which are valid across a large model class, have overlapping confidence intervals. All three estimate a linear demand response $\hat{\beta}^D$ of approximately $-200 \text{ MW}/(\text{EUR}/\text{MWh})$.¹⁵ The log-log specification yields a unitless estimate of -0.1 .¹⁶ An intervention on supply by a quantity similar to the magnitude of one standard deviation of the instrument, 693 MW,¹⁷ leads to a predicted price change of EUR -0.64 ,¹⁸ and thus an increase in the equilibrium quantities of 140 MW.

Overlap pattern. We can also analyze the pattern of the different CIV estimators to learn about the real-world dynamics of electricity demand response. First, we can reject a linear model without relevant structural autocorrelation. Figure 9 shows that the pattern of estimators does not correspond to the pattern that we would expect if real-world dynamics would not bias a naive estimate: The confidence intervals of the benchmark estimator and the naive IV estimator do not overlap, for example. Also, we can reject that the observed autocorrelation in the time series is only caused by an autocorrelation in the residual demand, as in Model II: the estimators #1 and #5 yield smaller absolute values than estimators #2 and #8 for the own-price elasticity, with non-overlapping confidence sets. Second, our results indicate that the dynamics of electricity markets are more complex than suggested by the simple Models I, II, and III because, taken in isolation, none of the models is sufficient to explain the disagreement between the estimators. Our approach only allows us to falsify structural assumptions but not to confirm them.

Robustness. As mentioned at the beginning of this section, we make the strong assumption that the demand elasticity is constant across hours. However, different consumers are active in the electricity market at other times of the day, resulting in a changing composition of consumer elasticities. To account for this temporal heterogeneity, Knaut and Paulus (2017) investigate the level of demand response by the hour of the day, and Hirth, Khanna, and Ruhnau (2024) additionally analyze heterogeneity by weekday and season. In Appendix D.2, we provide a robustness analysis that divides the hours into on-peak and off-peak periods (with lags of the variables extending into the respective other phase as needed). The results

¹⁴For the log-log transformation, we exclude prices of zero and below.

¹⁵#2 estimates a value of -220 , #4 of -183 , and #8 of -181 . All estimates in $\text{MW}/(\text{EUR}/\text{MWh})$.

¹⁶#2 estimates a value of -0.1 , #4 of -0.11 , and #8 of -0.08 . All estimates are unitless.

¹⁷The standard deviation is obtained from the residual time series of wind production, i.e., after being regressed on the whole conditioning set of estimator #2.

¹⁸We obtain the price prediction by multiplying the standard deviation with the first stage coefficient $\hat{\pi} = -0.00096 \text{ (EUR/MWh)/MW}$.

show that the overlap pattern holds primarily for on-peak demand. This period is also when we would expect the highest demand response, and the quantitative results of the benchmark estimator #2 are correspondingly higher. In contrast, the structural dynamics during off-peak periods (i.e., at night) suggest a different dynamic, which could be investigated further with our proposed method.

Intervention. Causal inference is concerned with predicting the effect of interventions on a system. $\hat{\beta}^D$ is an estimate of the structural coefficient β^D of the demand equation. It can be interpreted as the demand response to a hypothetical intervention on the equilibrium price. If β^D were indeed -200 MW/(EUR/MWh) and it were possible to reduce the price by EUR 1, demand would increase by 200 MWh/h (see Figure 10 (middle)). We call this intervention hypothetical because a direct and isolated intervention on the price only would yield a market that is out of balance: If one were to set the price to p_{fix} , not only would the quantity demanded change, but so would the quantity supplied, resulting in an imbalanced market ($S(p_{\text{fix}}) \neq D(p_{\text{fix}})$). One may also consider demand or supply shocks that shift or alter the demand or supply curve. For example, a supply shock that occurs regularly is changing weather conditions that affect the availability of renewable energy. Figure 10 (right) illustrates a supply shock that leads to the equilibrium quantities Q . Unlike the example of a hypothetical intervention on the price, in this scenario, equilibrium quantity and price are endogenous, ensuring that the market clears.

6 Conclusion

This paper demonstrates the advantages of using directed acyclic graphs to estimate the price elasticity of demand with autocorrelated instruments. While it is well known that autocorrelation can lead to biased estimates, relying solely on observed autocorrelation and statistical tests can mislead researchers seeking valid estimators. Causal graphs allow us to express structural assumptions transparently, understand biasing dynamics, derive multiple valid estimators, and ultimately test the validity of assumptions about structural dependencies over time. Suppose economists wish to benefit from causal time graphs. In that case, they can do so without abandoning the idea of simultaneous determination of supply and demand: equilibrium relationships can be represented in DAGs similarly to a time-instantaneous hidden confounding.

The graphical CIV criteria provide two different approaches to valid IV estimators. One approach blocks paths by including lagged terms in the conditioning set, and the other simultaneously estimates nuisance effects. While the former is more robust to model misspecification, the latter can be more powerful at small sample sizes.

We have applied the above to estimate the own-price elasticity of electricity demand under three competing structural assumptions: a model in which demand exhibits inertia, a model in which demand additionally depends on previous prices and a model of heterogeneous response. We have shown through simulations that each model can manifest identical levels of observed demand autocorrelation, illustrating that the observed correlation alone cannot predict potential bias. The estimates diverge when conditional IV estimators are applied to German electricity demand. This suggests that the widely used IV estimator, which ignores time dynamics, cannot identify the own-price elasticity without significant bias.

Looking ahead, the similarity between the CIV estimators derived from causal graphs and the lag-augmented local projection IV estimator calls for an extension of causal time graphs to estimate impulse response functions. Causal full time graphs can help strengthen the intuition behind recent results in this literature. To give just three examples: First, local projection IV estimators are found to identify the same impulse response function as a vector autoregressive model with the instrument ordered first (Plagborg-Møller and Wolf, 2021) . Second, local projection IV is found to be robust to misspecification (Montiel Olea, Plagborg-Møller, et al., 2024), which could be related to the lag-augmented local projection IV estimator blocking the same unblocked path multiple times. Third, Montiel Olea and Plagborg-Møller (2021) prove that for a local projection IV estimator, the use of Eicker-Huber-White heteroskedasticity-robust standard errors is sufficient without further needing to correct for autocorrelation. This result could be because the lagged terms in the conditioning set already sufficiently account for structural and, hence, bias-inducing autocorrelation. Going beyond recent findings, the graphical approach would further provide nonparametric proof, and the CIV criteria allow for a transparent analysis of the validity of the local projection IV estimator under hidden confounding.

Thus, as this paper has shown, the graphical approach is a powerful addition to the toolkit of every economist working with data structured as a time series, with many applications yet to be explored.

7 Funding

Researchers from the Hertie School gratefully acknowledge financial support from the German Federal Ministry of Education and Research through the ARIADNE Project (FKZ 03SFK5K0). Lion Hirth is owner and director of Neon Neue Energieökonomik GmbH, an energy economics consulting firm. Oliver Ruhnau is owner and director of tournesol energy GmbH, an energy consulting and software firm.

References

- ACER (2020). *Decision No 29/2020 on the methodology and assumptions that are to be used in the bidding zone review process and for the alternative bidding zone configurations to be considered*. Tech. rep. Ljubljana, Slovenia: European Union Agency for the Cooperation of Energy Regulators (cit. on p. 16).
- Albadi, M.H. and E.F. El-Saadany (2008). “A summary of demand response in electricity markets”. In: *Electric Power Systems Research* 78.11, pp. 1989–1996. DOI: [10.1016/j.epsr.2008.04.002](https://doi.org/10.1016/j.epsr.2008.04.002) (cit. on p. 17).
- Angrist, J. D., K. Graddy, and G. W. Imbens (2000). “The Interpretation of Instrumental Variables Estimators in Simultaneous Equations Models with an Application to the Demand for Fish”. In: *The Review of Economic Studies* 67.3, pp. 499–527 (cit. on pp. 1, 3).
- Angrist, J.D. and J.-S. Pischke (2009). *Mostly harmless econometrics: an empiricist’s companion*. Princeton: Princeton University Press (cit. on p. 3).
- Bongers, S. et al. (2021). “Foundations of Structural Causal Models with Cycles and Latent Variables”. In: *Annals of Statistics* 49.5, pp. 2885–2915 (cit. on p. 5).
- Bönte, W. et al. (2015). “Price elasticity of demand in the EPEX spot market for electricity—New empirical evidence”. In: *Economics Letters* 135, pp. 5–8. DOI: [10.1016/j.econlet.2015.07.007](https://doi.org/10.1016/j.econlet.2015.07.007) (cit. on p. 16).
- Borenstein, S. (2002). “The Trouble With Electricity Markets: Understanding California’s Restructuring Disaster”. en. In: *Journal of Economic Perspectives* 16.1, pp. 191–211. DOI: [10.1257/0895330027175](https://doi.org/10.1257/0895330027175) (cit. on pp. 15, 16).
- Brito, C. and J. Pearl (2002). “Generalized instrumental variables”. In: *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 85–93 (cit. on pp. 1, 7).
- Christ, C. F. (1994). “The Cowles Commission’s Contributions to Econometrics at Chicago, 1939-1955”. In: *Journal of Economic Literature* 32.1, pp. 30–59 (cit. on p. 4).
- Ciarreta, A., C. Pizarro-Irizar, and A. Zarraga (2020). “Renewable energy regulation and structural breaks: An empirical analysis of Spanish electricity price volatility”. In: *Energy Economics* 88, p. 104749. DOI: [10.1016/j.eneco.2020.104749](https://doi.org/10.1016/j.eneco.2020.104749) (cit. on p. 16).
- Cramton, P., A. Ockenfels, and S. Stoft (2013). “Capacity Market Fundamentals”. In: *Economics of Energy & Environmental Policy* 2.2. DOI: [10.5547/2160-5890.2.2.2](https://doi.org/10.5547/2160-5890.2.2.2) (cit. on p. 15).
- Cunningham, S. (2021). *Causal inference: the mixtape*. New Haven, London: Yale University Press (cit. on p. 3).
- Dawid, P. (2021). “Decision-theoretic foundations for statistical causality”. In: *Journal of Causal Inference* 9.1, pp. 39–77. DOI: [doi:10.1515/jci-2020-0008](https://doi.org/10.1515/jci-2020-0008) (cit. on p. 6).
- Destatis (2023). *Population by nationality and federal states*. en (cit. on p. 28).
- ENTSO-E (2023). *ENTSO-E Transparency Platform* (cit. on p. 28).
- Eurostat (Feb. 2023). *Heating and cooling degree days*. en (cit. on p. 28).
- Fabra, N. et al. (2021). “Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market”. In: *AEA Papers and Proceedings* 111, pp. 425–429. DOI: [10.1257/pandp.20211007](https://doi.org/10.1257/pandp.20211007) (cit. on p. 16).
- Feiertagskalender (2023). *School holidays Germany* (cit. on p. 28).

- Friedman, N. and D. Koller (2003). “Being Bayesian about Bayesian Network Structure: A Bayesian Approach to Structure Discovery in Bayesian Networks.” In: *Machine Learning* 50, pp. 95–125 (cit. on p. 25).
- Global Modeling And Assimilation Office and S. Pawson (2015). “MERRA-2 tavg3_3d_tdt_Np: 3d, 3-Hourly, Time-Averaged, Pressure-Level, Assimilation, Temperature Tendencies V5.12.4”. In: *NASA Goddard Earth Sciences Data and Information Services Center*. DOI: [10.5067/9NCR9DDDOPFI](https://doi.org/10.5067/9NCR9DDDOPFI) (cit. on p. 28).
- Henckel, L., M. Buttenschoen, and M. H. Maathuis (Nov. 2023). “Graphical tools for selecting conditional instrumental sets”. en. In: *Biometrika*, asad066. DOI: [10.1093/biomet/asad066](https://doi.org/10.1093/biomet/asad066) (cit. on pp. 1, 7, 9).
- Hendry, D. F. (Feb. 1995). *Dynamic Econometrics*. en. 1st ed. Oxford University Press Oxford. DOI: [10.1093/0198283164.001.0001](https://doi.org/10.1093/0198283164.001.0001) (cit. on pp. 3, 5).
- Hernan, M. and J. Robins (2020). *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC (cit. on p. 5).
- Hirth, L., T. M. Khanna, and O. Ruhnau (July 2024). “How aggregate electricity demand responds to hourly wholesale price fluctuations”. In: *Energy Economics* 135, p. 107652. DOI: [10.1016/j.eneco.2024.107652](https://doi.org/10.1016/j.eneco.2024.107652) (cit. on pp. 16–18).
- Holland, P. W. (Dec. 1986). “Statistics and Causal Inference”. In: *Journal of the American Statistical Association* 81.396, pp. 945–960 (cit. on p. 3).
- Hosius, E. et al. (2023). “The impact of offshore wind energy on Northern European wholesale electricity prices”. In: *Applied Energy* 341, p. 120910. DOI: [10.1016/j.apenergy.2023.120910](https://doi.org/10.1016/j.apenergy.2023.120910) (cit. on p. 16).
- Hünermund, P. and E. Bareinboim (Mar. 2023). “Causal inference and data fusion in econometrics”. In: *The Econometrics Journal*. DOI: [10.1093/ectj/utad008](https://doi.org/10.1093/ectj/utad008) (cit. on p. 3).
- Imbens, G. W. (2020). “Potential Outcome and Directed Acyclic Graph Approaches to Causality: Relevance for Empirical Practice in Economics”. In: *Journal of Economic Literature* 58.4, pp. 1129–79 (cit. on p. 3).
- Imbens, G. W. and D. B. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. New York, NY: Cambridge University Press (cit. on p. 5).
- Investing.com (2023). *Commodities Prices*. en (cit. on p. 28).
- Jordà, Ò., M. Schularick, and A. M. Taylor (July 2015). “Betting the house”. In: *Journal of International Economics*. 37th Annual NBER International Seminar on Macroeconomics 96, S2–S18. DOI: [10.1016/j.jinteco.2014.12.011](https://doi.org/10.1016/j.jinteco.2014.12.011) (cit. on p. 4).
- Knaut, A. and S. Paulus (2017). “When are consumers responding to electricity prices? An hourly pattern of demand elasticity”. In: *EWI Working Papers* (cit. on pp. 16, 18).
- Lauritzen, S. L. (1996). *Graphical Models*. New York, NY, USA: Oxford University Press (cit. on pp. 1, 4, 8, 25).
- Lauritzen, S. L. et al. (1990). “Independence properties of directed markov fields”. In: *Networks* 20, pp. 491–505 (cit. on p. 6).
- Maniatis, G. I. and N. T. Milonas (2022). “The impact of wind and solar power generation on the level and volatility of wholesale electricity prices in Greece”. In: *Energy Policy* 170, p. 113243. DOI: [10.1016/j.enpol.2022.113243](https://doi.org/10.1016/j.enpol.2022.113243) (cit. on p. 16).
- Montel, M. and A. Yakovets (Sept. 2023). *holidays: Generate and work with holidays in Python* (cit. on p. 28).

- Montiel Olea, J.L. and M. Plagborg-Møller (July 2021). “Local Projection Inference Is Simpler and More Robust Than You Think”. In: *Econometrica* 89.4, pp. 1789–1823. DOI: [10.3982/ECTA18756](https://doi.org/10.3982/ECTA18756) (cit. on pp. 1, 11, 20).
- Montiel Olea, J.L., M. Plagborg-Møller, et al. (May 2024). “Double Robustness of Local Projections and Some Unpleasant VARithmetic”. In: *National Bureau of Economic Research Working Paper Series*. DOI: [10.3386/w32495](https://doi.org/10.3386/w32495) (cit. on pp. 1, 4, 20).
- Neyman, J. (1923). “On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9 (translated)”. In: *Statistical Science* 5, pp. 465–480 (cit. on pp. 3, 5).
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. 2nd. New York, USA: Cambridge University Press (cit. on pp. 1, 3–6, 8, 25).
- Pearl, J. and D. Mackenzie (2020). *The book of why: the new science of cause and effect*. First trade paperback edition. New York: Basic Books. 418 pp. (cit. on p. 3).
- Peters, J., S. Bauer, and N. Pfister (Mar. 2022). “Causal Models for Dynamical Systems”. In: *Probabilistic and Causal Inference: The Works of Judea Pearl*. 1st ed. Vol. 36. New York, NY, USA: Association for Computing Machinery, pp. 671–690 (cit. on pp. 4, 5).
- Peters, J., P. Bühlmann, and N. Meinshausen (2016). “Causal inference using invariant prediction: identification and confidence intervals”. In: *Journal of the Royal Statistical Society: Series B (with discussion)* 78.5, pp. 947–1012 (cit. on p. 4).
- Peters, J., D. Janzing, and B. Schölkopf (2013). “Causal Inference on Time Series using Structural Equation Models”. In: *Advances in Neural Information Processing Systems 26 (NeurIPS)*. Curran Associates, Inc., pp. 585–592 (cit. on pp. 1, 5, 8, 25).
- (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. Cambridge, MA, USA: MIT Press (cit. on p. 4).
- Plagborg-Møller, M. and C. K. Wolf (2021). “Local Projections and VARs Estimate the Same Impulse Responses”. en. In: *Econometrica* 89.2, pp. 955–980. DOI: [10.3982/ECTA17813](https://doi.org/10.3982/ECTA17813) (cit. on pp. 4, 20).
- Ramey, V. A. and S. Zubairy (Apr. 2018). “Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data”. In: *Journal of Political Economy* 126.2, pp. 850–901. DOI: [10.1086/696277](https://doi.org/10.1086/696277) (cit. on p. 4).
- Richardson, T. (2003). “Markov properties for acyclic directed mixed graphs”. In: *Scandinavian Journal of Statistics* 30.1, pp. 145–157 (cit. on p. 6).
- Rubin, D. B. (2005). “Causal inference using potential outcomes”. In: *Journal of the American Statistical Association* 100, pp. 322–331 (cit. on pp. 3, 5).
- Sheppard, K. and et al. (Apr. 2024). *bashstage/linearmodels: Release 6.0*. DOI: [10.5281/ZENODO.10981685](https://doi.org/10.5281/ZENODO.10981685) (cit. on p. 12).
- Spirtes, P., C. Glymour, and R. Scheines (2000). *Causation, Prediction, and Search*. 2nd. MIT Press (cit. on pp. 1, 3, 4).
- Stock, J. H. and M. W. Watson (May 2018). “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments”. en. In: *The Economic Journal* 128.610, pp. 917–948. DOI: [10.1111/eoj.12593](https://doi.org/10.1111/eoj.12593) (cit. on pp. 1, 4, 11).
- Stopa, K. (Aug. 2019). *suntime: Simple sunset and sunrise time calculation python library*. (Cit. on p. 28).

ELASTICITIES WITH GRAPHS

- Thams, N. et al. (2022). “Identifying Causal Effects using Instrumental Time Series: Nuisance IV and Correcting for the Past”. In: *Journal of Machine Learning Research (accepted)*” arXiv:2203.06056. arXiv: [2203.06056\[stat\]](#) (cit. on pp. [1](#), [8](#), [9](#), [12](#), [25](#), [26](#)).
- Tinbergen, J. (1940). “Econometric Business Cycle Research”. In: *The Review of Economic Studies* 7.2, pp. 73–90. DOI: [10.2307/2967472](#) (cit. on pp. [1](#), [3](#)).
- Verma, T. S. and J. Pearl (1991). “Equivalence and synthesis of causal models”. In: *Proceedings of the 6th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 255–270 (cit. on p. [26](#)).
- Webel, K. (2022). “A Review of Some Recent Developments in the Modelling and Seasonal Adjustment of Infra-Monthly Time Series”. en. In: *SSRN Electronic Journal*. DOI: [10.2139/ssrn.4201921](#) (cit. on p. [3](#)).
- Wold, H. O. A. (1964). *Econometric model building; essays on the causal chain approach*. Amsterdam: North-Holland Pub. Co. (cit. on p. [3](#)).
- Wright, P. G. (1928). *The Tariff on Animal and Vegetable Oils*. Investigations in International Commercial Policies. New York, NY: Macmillan (cit. on pp. [1](#), [3](#)).

A Causal Models for Time Series

This appendix provides a formal introduction to causal inference for time series. First, we give a formal definition of structural causal models (SCM), followed by an introduction to directed acyclic graphs (DAGs) for time series. Finally, we introduce the concept of marginalization, which is a way to simplify the representation of graphs.

A.1 Structural Causal Models for Time Series

Let $\{U_t^j\}_{t \in \mathbb{Z}, j \in \{1, \dots, d\}}$ be a set of d jointly independent random error variables. The index t encodes time, and the index j is the time series component.

For all $j \in \{1, \dots, d\}$ let $\text{PA}(j) \subseteq \{(k, s) \mid k \in \{0, \dots, q\}, s \in \{1, \dots, d\}\}$ and $p_j := |\text{PA}(j)|$ be the number of elements in $\text{PA}(j)$, and let $f^j : \mathbb{R}^{p_j} \times \mathbb{R} \rightarrow \mathbb{R}$ be a measurable function. The sets $\text{PA}(j)$ encode the (causal) parents of X_t^j , which we assume do not depend on t . For all $j \in \{1, \dots, d\}$ and all $t \in \mathbb{Z}$ we consider the structural equations

$$X_t^j := f^j \left(\text{PA}(j, t), U_t^j \right), \quad (11)$$

where the *parents* of X_t^j are defined by $\text{PA}(j, t) := (X_{t-k_1}^{s_1}, \dots, X_{t-k_{p_j}}^{s_{p_j}})$ with $\{(k_1, s_1), \dots, (k_{p_j}, s_{p_j})\} = \text{PA}(j)$, where the (k_i, s_i) are in lexicographic order.¹⁹

We further assume that there is a unique solution²⁰ to (11) that is weakly stationary and that the process realization is covariance ergodic. We further assume that the induced distribution is Markov with respect to the induced graph (see Section A.2). For linear functions f^1, \dots, f^d , which we consider in this paper, these properties are well-studied (see, for example, Thams et al., 2022, Theorem 1).

A.2 Directed Acyclic Graphs for Time Series

We now introduce some standard graph terminology (e.g., Lauritzen, 1996; Friedman and Koller, 2003; Pearl, 2009), adapted to the case of infinitely many nodes. Let $V \subseteq \mathbb{Z} \times \{1, \dots, d\}$. A *directed graph* $\mathcal{G} = (V, E)$ is composed of nodes V (also called vertices) and edges $E \subseteq V \times V$ with the restriction that for all $u, v \in V$ either $(u, v) \notin E$ or $(v, u) \notin E$ or both. We sometimes use the notation $v \rightarrow u$, meaning that $(v, u) \in E$.

A *path* in \mathcal{G} is a sequence of distinct vertices v_1, \dots, v_m , such that an edge connects each consecutive pair of vertices. If $v_k \rightarrow v_{k+1}$ for all $k \in \{1, \dots, m-1\}$, we say that the path is *directed from v_1 to v_m* . A directed path $v_1 \rightarrow \dots \rightarrow v_m$ such that $(v_m, v_1) \in E$ is called a *cycle*. A directed graph \mathcal{G} is classified as a *directed acyclic graph (DAG)* if it has no cycles.

Because we deal with time series, the structural equations induce a graph with infinitely many nodes, i.e., a *full time graph* (Peters, Janzing, and Schölkopf, 2013) \mathcal{G}_{FT} over nodes $V = \{(t, s) \mid t \in \mathbb{Z}, s \in \{1, \dots, d\}\}$, where we draw an edge from (t_1, s_1) to (t_2, s_2) if $(t_1, s_1) \in \text{PA}(t_2, s_2)$. We assume that the parent sets are such that \mathcal{G}_{FT} is directed and acyclic. We say $v \in V$ is a *descendant* of $u \in V$ if $v \neq u$, and there exists a directed path from u to v . We denote the set of descendants of u by $\text{DE}(u) := \{v \in V \mid v \text{ is a descendant of } u\}$.

¹⁹This means $(k, s) < (v, g)$ if and only if $k < v$ or $(k = v \text{ and } s < g)$.

²⁰More precisely, we assume that all solutions induce the same distribution.

A.3 Graph marginalization

Graph marginalization has been studied extensively (see, e.g., Verma and Pearl, 1991). The main idea of graph marginalization is to simplify the graph while keeping the same d-separation statements over the observed nodes. In the context of this paper, we marginalize out the unobserved error term, which causes the endogeneity of prices and quantities (see Figure 11). A formal definition of a marginalized time graph is given in Thams et al. (2022), which, too, consider graphs with infinitely many nodes.

Definition A.1. Consider a full time graph \mathcal{G}_{FT} over nodes $V = \{(t, s) \mid t \in \mathbb{Z}, s \in \{1, \dots, d\}\}$. Let $M = \{(t_1, i_1), \dots, (t_m, i_m)\}$ be a finite collection of nodes in \mathcal{G}_{FT} . The *marginalized time graph*, \mathcal{G}_M , is the graph over nodes M where for all $i, j \in M$ there is:

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

B Details of the two-stage-least-squares estimator

Let $T \in \mathbb{N}$ and let $\mathcal{B}, \mathcal{I}, \mathcal{X}$ and Y satisfy (CIV1) to (CIV3). Assume that we have observations at time points $\{1, \dots, T\}$. Let $\{(t_1, s_1), \dots, (t_n, s_n)\} = \mathcal{B}$, where (t_i, s_i) are in lexicographic order. Without loss of generality, assume that $T \in \{t_1, \dots, t_n\}$ and let $k := \max\{|t_i - t_j| \mid i, j \in \{1, \dots, n\}\}$. Define

$$\mathbf{B} := \begin{bmatrix} X_{t_1}^{s_1} & \dots & X_{t_n}^{s_n} \\ X_{t_1-1}^{s_1} & \dots & X_{t_n-1}^{s_n} \\ \vdots & & \vdots \\ X_{t_1-T+k+1}^{s_1} & \dots & X_{t_n-T+k+1}^{s_n} \end{bmatrix}$$

and define \mathbf{X}, \mathbf{Y} and \mathbf{I} analogously. Define $r_{\mathbf{I}}$ as the residuals of regressing \mathbf{I} on \mathbf{B} with ordinary least squares and define $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ analogously. Define $P_{\mathbf{I}} := r_{\mathbf{I}}(r_{\mathbf{I}}^\top r_{\mathbf{I}})^{-1} r_{\mathbf{I}}^\top$ and

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

This is the closed-form solution to (10). By Thams et al. (2022, Theorem 5) $\hat{\beta}$ is consistent for $T \rightarrow \infty$.

C Additional information on Section 4

C.1 Additional autocorrelation and error plots for conditional IV

In Section 4, we show that the absolute percentage error of the naive IV estimator $\text{CIV}(W_t | P_t \rightarrow D_t | \emptyset)$, which neglects time dependencies, depends on the structural autocorrelation present in both the instrument and the demand equation (see Figure 6 for empirical results for Model I).

Figure 12 shows the corresponding plot for Model II, where the autocorrelation is induced by the autocorrelated error term with β^{B1} , and an estimator that conditions on the past of the dependent variable: $\text{CIV}(W_t|P_t \rightarrow D_t|D_{t-1})$ (estimator #3). The results for Model II are consistent with the general findings that the error depends on the underlying structural autocorrelation of both the instrument and the demand.

In Section 4, we further argue that knowing the observed level of autocorrelation of demand, e.g., obtained by analyzing the (partial) autocorrelation function, does not suffice to correct the estimation bias: the observed autocorrelation of the demand function is associated with different levels of error for the naive IV estimator $\text{CIV}(W_t|P_t \rightarrow D_t|\emptyset)$ (see Figure 7). Figure 13 extends the same analysis to estimator #3, which conditions on past demand. As expected from the theory, if the estimation model matches the data generation process, the estimator is unbiased regardless of the observed level of autocorrelation. However, under model misspecification, the same observed level of autocorrelation can be associated with percentage errors ranging from $-300 \text{ MW}/(\text{EUR}/\text{MWh})$ to $200 \text{ MW}/(\text{EUR}/\text{MWh})$.

C.2 Additional indicator plots for Model II and Model III

In Section 4, we have compared the statistical performance of estimators #2, #8 and #4 on data generated from Model I. Figure 14 and Figure 15 compare the performances on data generated by Models II and III, respectively. As expected from theory, all estimators maintain coverage regardless of sample size. Estimator #8 tends to have a smaller absolute percentage error and smaller confidence intervals.

C.3 Divergence between CIV estimates for simulated data

In Section 3, we have argued that non-overlapping confidence intervals of presumably valid estimators should lead to model rejection (see also Figure 5 for data from Model I). Figures 16 and 17 show the corresponding patterns of CIV estimators for Model II and Model III, respectively. The table to the right of the figures shows which estimators are valid according to the CIV criteria.

For Model II (Figure 16), the estimators with a checkmark in column II are valid, and the confidence intervals of the corresponding estimators overlap. As done by estimator #3, conditioning on lagged demand generally induces a bias as the estimator conditions on a collider. For Model II, the estimator #7 is a special case: Based on the CIV criteria alone, the estimator is invalid. The estimated effect $P_t \rightarrow D_t$ is nevertheless close to the real effect. The pattern of estimators we observed in Figure 17 is also consistent with the expectation given the Model III assumption. If one were to assume Model III, the researcher does not see much evidence of rejecting that model assumption. However, if, for example, she were to incorrectly assume Model I, there is evidence that this assumption is false (see, e.g., estimators #2 and #3).

D Additional information on Section 5

D.1 Data

We now describe in more detail the data used in the application. The main variables are electricity load (MW), electricity price (€/MWh), and, for the instrument, wind generation (MW, including both onshore and offshore generation), all obtained from ENTSO-E (2023), the European Network of Transmission System Operators for Electricity. ENTSO-E’s measurement and reporting of variables is bidding zone-specific. Since October 1, 2018, the bidding zone includes the territory of Germany and Luxembourg. Previously, it also included the territory of Austria. The data used in our analysis spans from November 1st, 2017, to December 31st, 2020.

In addition to the main variables, several control series are incorporated into the analysis. If the data exists at a daily frequency, we assume the same value for every hour of the day.

- at hourly frequency:
 - solar PV power generation (from ENTSO-E (2023))
 - surface temperature data for Germany from NASA’s MERRA-2 (Global Modeling And Assimilation Office and Pawson, 2015), transformed according to Eurostat’s methodology (Eurostat, 2023).
- at daily frequency, all from Investing.com (2023):
 - price of emissions allowances (EUA yearly futures)
 - coal prices (API 2 CIF ARA ARGUS-McCloskey futures)
 - natural gas price (Dutch TTF futures)

Public holidays and school holidays are two variables constructed as indices based on the fraction of the German population affected by a holiday on a given day, using data from the Federal Statistical Office of Germany (Statistisches Bundesamt Deutschland) for the population by state (Destatistis, 2023), public holiday dates from Python’s ‘holidays’ package (Montel and Yakovets, 2023), and school holiday dates from Feiertagskalender (2023). However, holidays common to the whole country and resulting in a value of $\text{holiday}_t = 1$ (e.g., Easter, Christmas, and the Day of German Unity on October 3rd) are dropped from the dataset. The largest observed value strictly smaller than one equals 0.644 (three days) and 0.568 (four days).

We also include the following covariates: sunlight, indicating whether at a given hour of the year, it is day (1) or night (0) at the geographic center of Germany (at 51°09’ N, 10°26’ E) using Python’s ‘suntime’ package (Stopa, 2019); hour of the week (one-hot encoding, 167 variables), month of the year (one-hot encoding, 11 variables), and year (one-hot encoding, four variables).

Figure 18 shows the autocorrelation and partial autocorrelation plots of the instrument and the electricity load after controlling for the control variables. We observe a significant autocorrelation, which is most pronounced for lags 1 and 2 but extends to up to 50 lags, corresponding to more than three days.

D.2 Sensitivity analysis for the price elasticity of German electricity demand

In Section 5, we assume a homogeneous elasticity across hours. However, structurally different consumers may be active across different hours of the day. In particular, one may expect different consumers during peak and off-peak hours. They could either have different elasticities, or they could also have different dynamic behavior. We split the data into two periods and perform the same analysis as in Section 5 on each of the splits: on-peak in Figure 19 and off-peak in Figure 20. The magnitudes of the elasticity estimates (log-log) and the demand response (linear) are comparable. However, the pattern of the estimates is different in the two different splits. This suggests that nighttime consumption has different dynamics than daytime consumption. Given the observed pattern of estimators, we further have to reject all three proposed simple models. This calls for carefully considering the model dynamics if one wants to delve deeper into how electricity consumers respond to a high-frequency price signal.

ELASTICITIES WITH GRAPHS

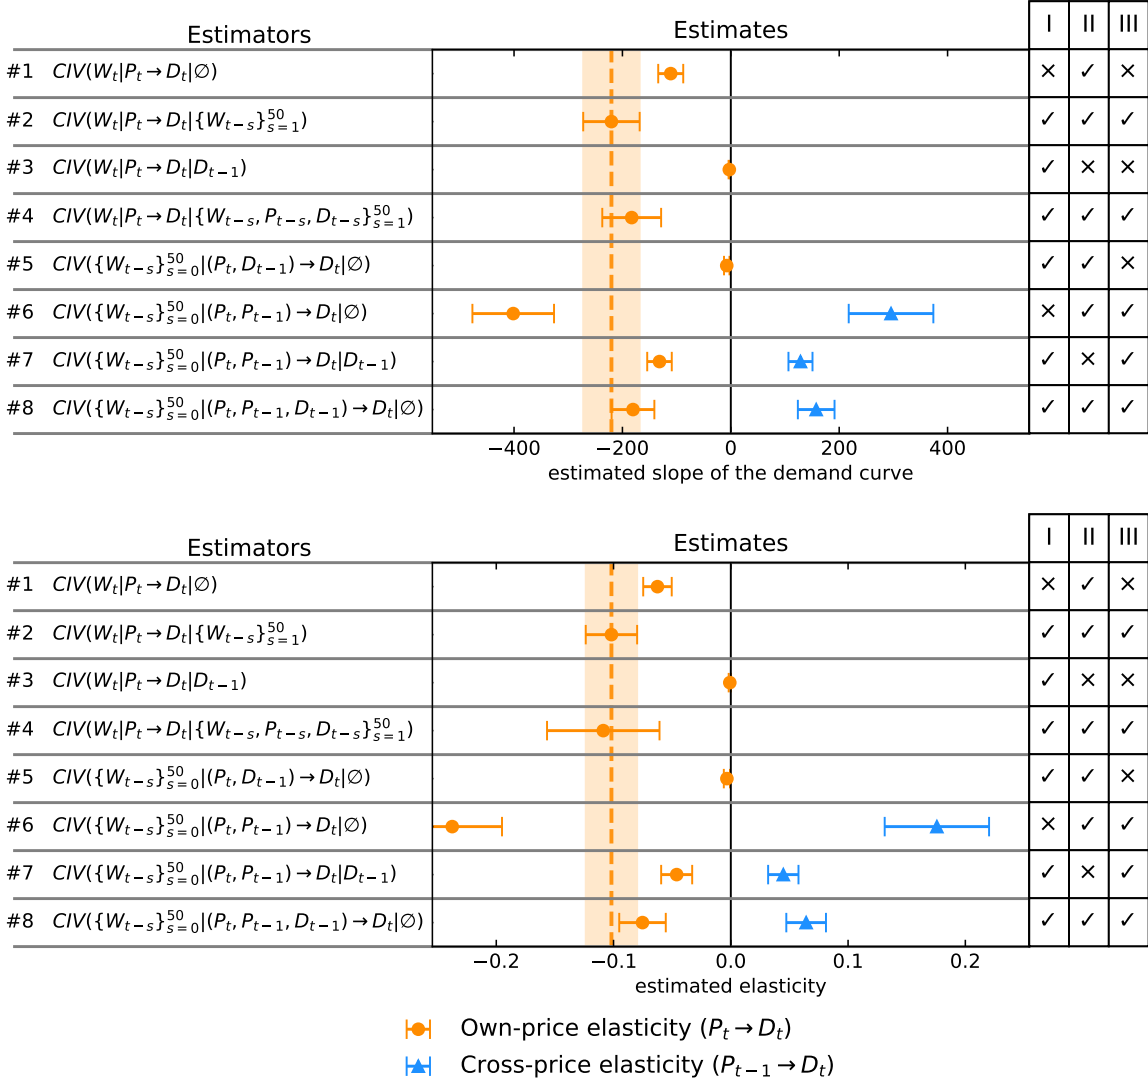


Figure 9: The own-price elasticity of aggregated German electricity demand, obtained by CIV estimators (top: linear specification; bottom: log-log specification). The highlighted area corresponds to the confidence interval of estimator #2. The tables on the right are reproductions of Table 1, showing the validity of each strategy for the structural causal models described in Section 2.

ELASTICITIES WITH GRAPHS




Figure 10: Illustrations of conceptually different interventions in a market at equilibrium. Left: Market at equilibrium, i.e., the state without any intervention. Middle: Hypothetical intervention on the equilibrium price, fixing the equilibrium price at P_{fix} . The demand elasticity is sufficient to calculate the effect of the intervention on the quantity demanded. However, the market does not clear because marginal cost is less than marginal willingness to pay. Right: Factual intervention X_{shift} that shifts the supply curve outward, leading to a new market equilibrium.




Figure 11: Left: Snapshot of a full-time graph at time point t . Right: Marginalized version (see Definition A.1), at time t . The resulting graphical object includes both directed (\leftarrow or \rightarrow) and bidirected edges (\leftrightarrow).

ELASTICITIES WITH GRAPHS

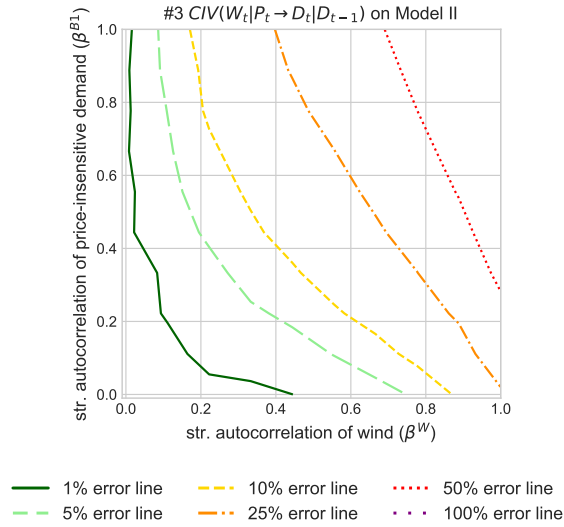


Figure 12: Isolines showing the average absolute percentage error of the point estimator as a function of the structural autocorrelation of wind (instrument) and the structural autocorrelation β^{B1} of the price-insensitive demand (dependent variable) for estimator #3. For the simulations, we divide both axes into ten equidistantly spaced autocorrelation coefficients between 0 and 1 each, based on which we run 20 simulations with five years of data ($T = 43,800$).

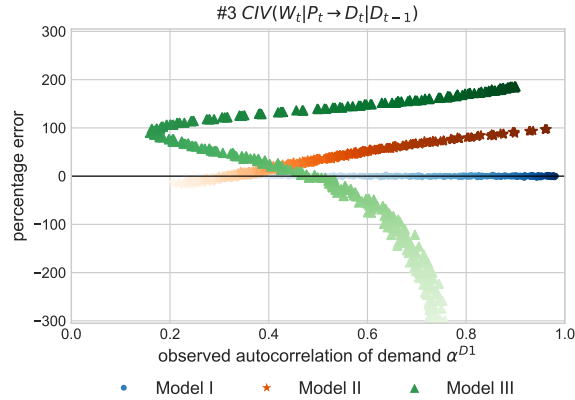


Figure 13: The percentage error of the CIV estimator #3 is shown as a function of the observed autocorrelation of demand α^{D1} . In each simulation we vary the structural dependencies: In Models I and II, the strengths of the structural autocorrelations, β^{D1} and β^{B1} , are both varied between -0.25 and 0.99 . In Model III, β^{P1} is varied between -250 MW/(EUR/MWh) and 250 MW/(EUR/MWh). The color intensity increases in tandem with the aforementioned ranges. For each model, 400 simulations were conducted, with a sample period of five years ($T = 43,800$).

ELASTICITIES WITH GRAPHS

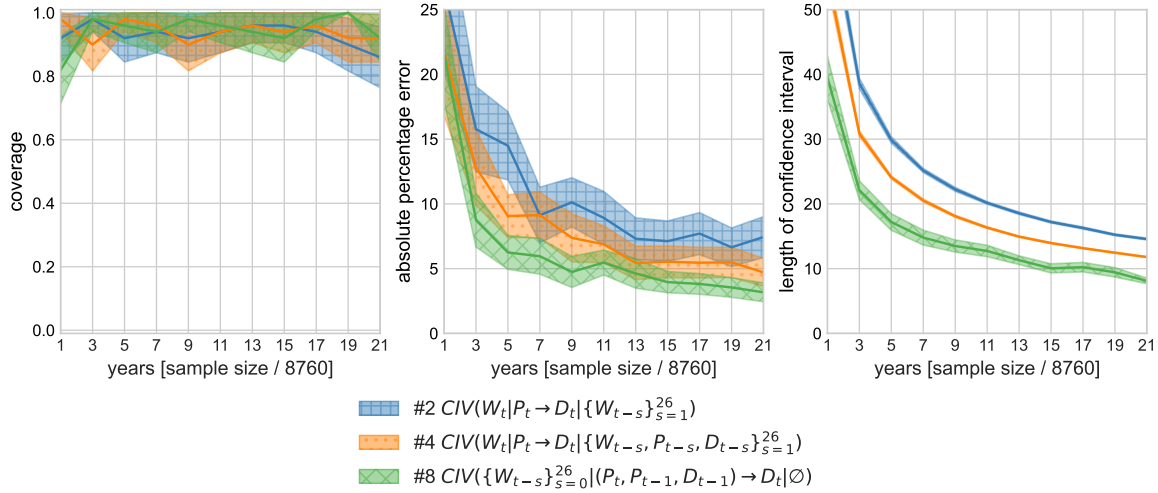


Figure 14: The performance of valid CIV estimators based on the indicators (a) coverage, (b) percentage error, and (c) length of the confidence interval for Model II at different sample sizes. The shaded areas represent the 95% confidence interval for the indicators. For each sample size, 50 simulations were run. For data generation, the parameters of the demand equation (5) are set to $\beta^P = -100 \text{ MW}/(\text{EUR}/\text{MWh})$ and $\beta^{B1} = 0.9$.

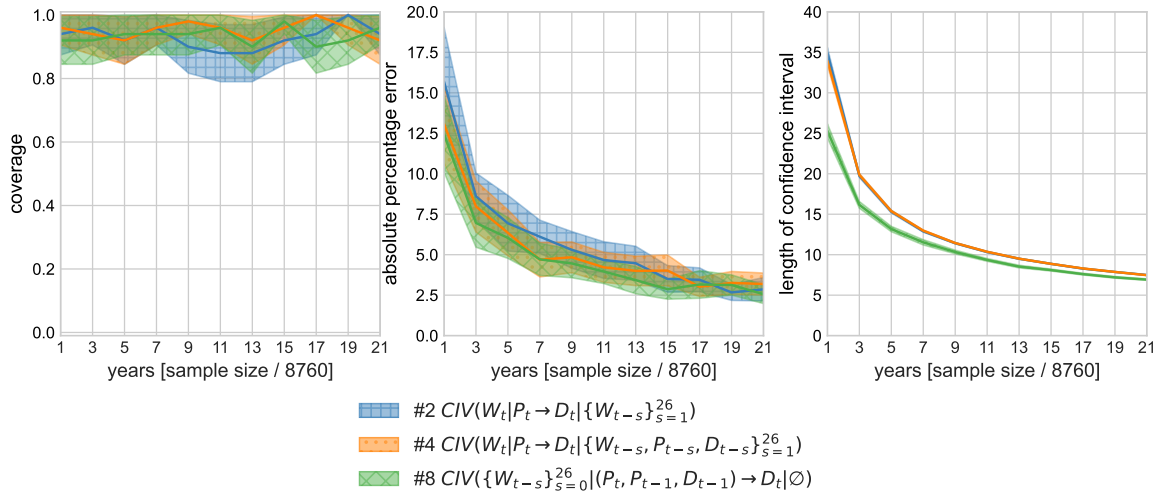


Figure 15: The performance of valid CIV estimators based on the indicators (a) coverage, (b) percentage error, and (c) length of the confidence interval for Model III at different sample sizes. The shaded areas represent the 95% confidence interval for the indicators. For each sample size, 50 simulations were run. For data generation, the parameters of the demand equation (8) are set to $\beta^P = -100 \text{ MW}/(\text{EUR}/\text{MWh})$ and $\beta^{P1} = +50 \text{ MW}/(\text{EUR}/\text{MWh})$.

ELASTICITIES WITH GRAPHS




Figure 16: CIV estimates of the demand response $\hat{\beta}^P$ when the simulated data does not exhibit structural autocorrelation of demand (Model II, see Figure 3 (left)). The slope of the price-sensitive demand β^P is set to -100 MW/(EUR/MWh), and the autocorrelation coefficient of the non-sensitive demand is set to $\beta^{B1} = 0.9$.

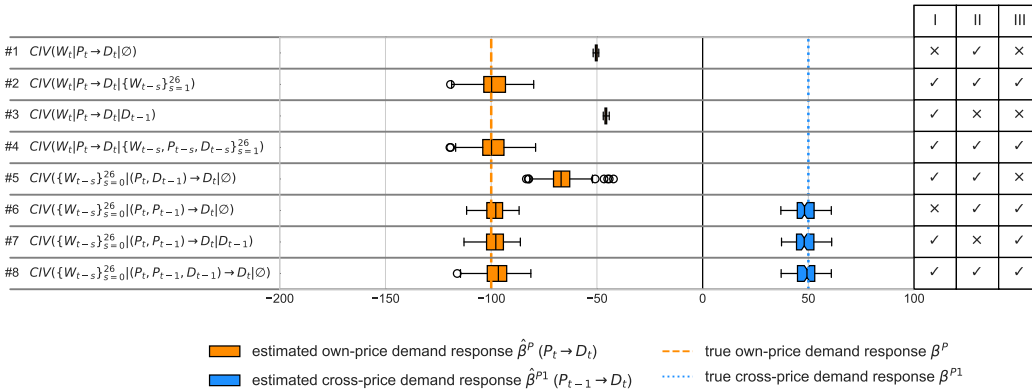


Figure 17: CIV estimates of the demand response $\hat{\beta}^P$ when the simulated data represents a cross-price response (Model III, see Figure 3 (right)). The true effect is $\beta^P = -100$ MW/(EUR/MWh), and the cross-price elasticity is $\beta^{P1} = +50$ MW/(EUR/MWh), representing substitution. According to theory (see Table 1), the estimators in column III should be consistent, and the simulation results support this. In this model, the validity of the estimators extends to both effects: A valid strategy recovers both the own-price response $\hat{\beta}^P$ and the cross-price response $\hat{\beta}^{P1}$.

ELASTICITIES WITH GRAPHS

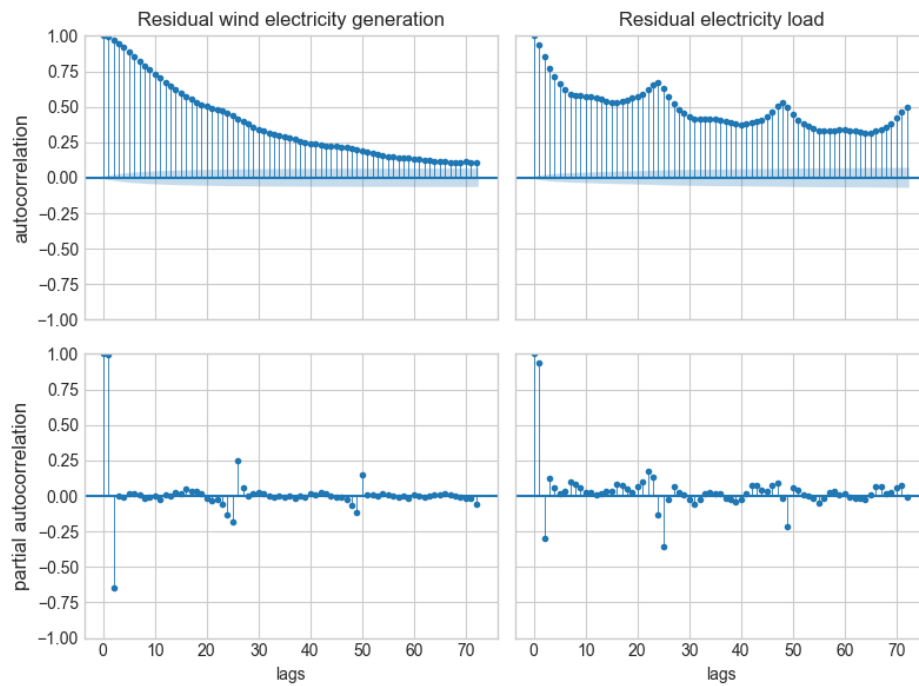


Figure 18: Autocorrelation function (above) and partial autocorrelation function (below) of the residuals of the instrumental time series wind power generation (left), and the residuals of the dependent variable German electricity demand (right). In all four cases, we first regress the original data on the covariates and then obtain the autocorrelation of the residual time series.

ELASTICITIES WITH GRAPHS




Figure 19: Own-price elasticity of on-peak German electricity demand, obtained from the CIV estimators given in Table 1 in linear specification (top) and log-log specification (bottom). The highlighted area corresponds to the confidence interval of estimator #2. The tables on the right are reproductions of Table 1, showing the validity of each strategy under different structural model assumptions. On-peak hours refer to the time between 8:00 and 19:59.

ELASTICITIES WITH GRAPHS

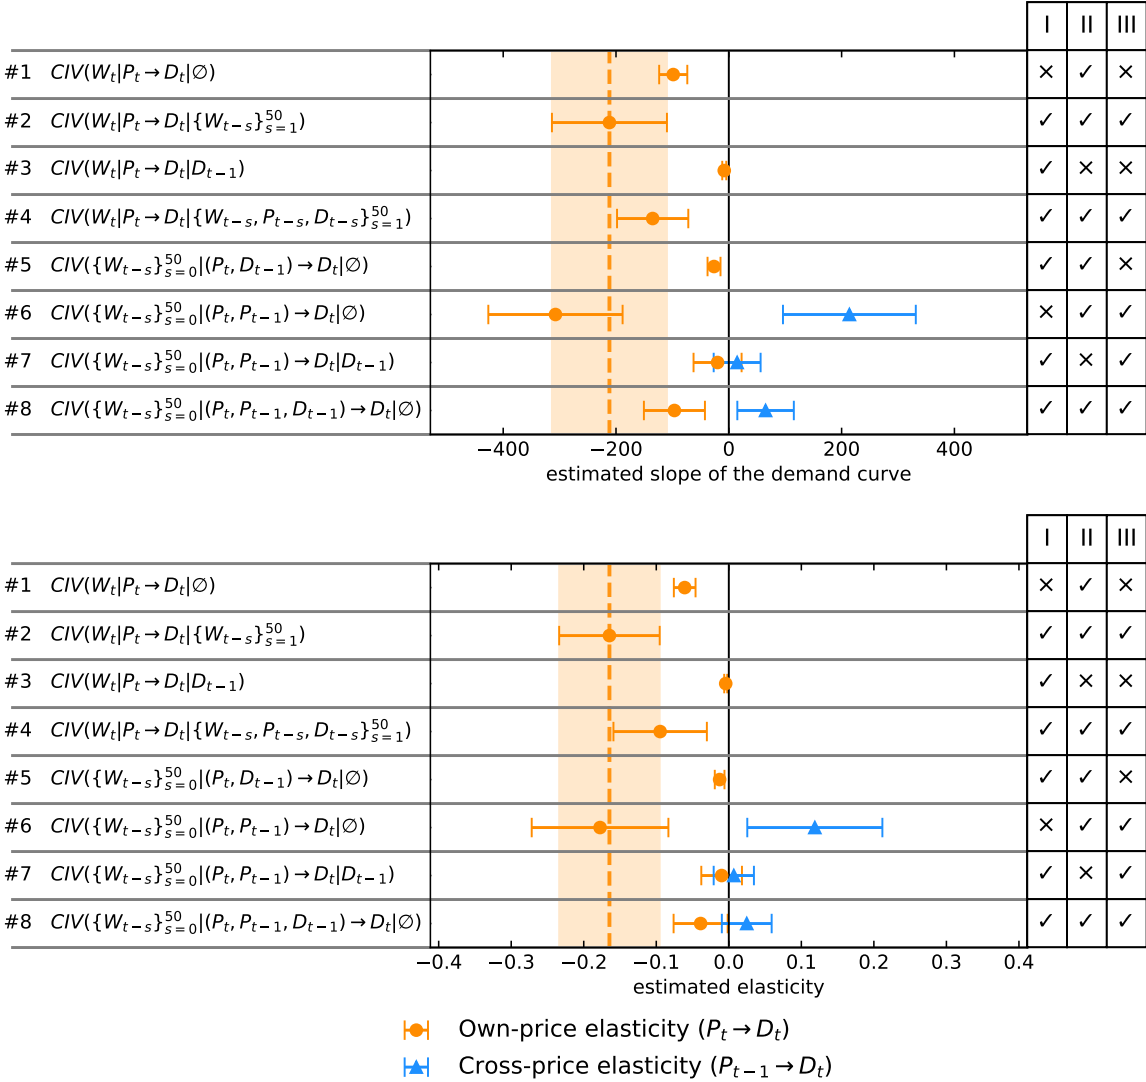


Figure 20: Own-price elasticity of off-peak German electricity demand, obtained from the CIV estimators given in Table 1 in the linear specification (top) and log-log specification (bottom). The highlighted area corresponds to the 95% confidence interval of estimator #2. The tables on the right are reproductions of Table 1, showing the validity of each strategy under different structural model assumptions. Off-peak hours refer to the time between 20:00 and 7:59.