OESTERREICHISCHE NATIONALBANK
EUROSYSTEM

Working Paper 271

DAG-Based Local Projections

Burkhard Raunig

OeNB Working Paper Series
28 pages

DAG-Based Local Projections

Directed acyclic graphs (DAGs) encode causal assumptions. This paper shows how DAGs can guide the specification of local projections (LPs) for estimating causal responses to policy interventions and shocks. An application to German industrial production reveals substantial negative effects of uncertainty shocks, with estimates varying widely across LP specifications. DAGs help to clarify these differences and assess potential biases arising from violations of the assumed causal structure.

Authors
Burkhard Raunig
Oesterreichische Nationalbank, Research, Vienna, AT

JEL classification
C18, C22, C26, E27

Keywords
Directed acyclic graph, local projection, uncertainty shock

Causal Graphs

Directed acyclic graphs (DAGs) visualize assumed causal relationships. Graphical rules can be applied to DAGs to analyze identification of causal effects, confounding, mediation, and selection bias.

Local Projections

Local projections (LPs) allow for easy calculation of dynamic responses to shocks or policy interventions using simple regressions.

About the study

This study outlines how DAGs can be used to select control variables and instruments in LPs to obtain causal responses. It applies DAG based LPs to empirically analyze the effects of uncertainty shocks on German industrial production.

DAG-Based Local Projections

Burkhard Raunig*

Abstract

Directed acyclic graphs (DAGs) provide a transparent framework for encoding causal structures and identifying causal effects. This paper demonstrates how DAGs help specify local projections (LPs) for estimating causal impulse responses. Examples illustrate how graphical rules can be used to select controls and instruments for identifying overall and path-specific effects. An empirical application to uncertainty shocks reveals substantial differences in the estimated responses of German industrial production across LP designs. The underlying DAGs help explain these differences and diagnose biases arising from violations of assumed causal structures. A DAG-based instrumental-variable LP reveals pronounced negative effects of U.S. uncertainty shocks.

Keywords: Directed acyclic graph; Local projection; Impulse response; Instrumental variable; Uncertainty shocks

JEL codes: C18, C22, C26, E27

*Oesterreichische Nationalbank, Research, Otto-Wagner-Platz 3, A-1090 Vienna, Austria, e-mail: burkhard.raunig@oenb.at. The views expressed in this study do not necessarily reflect the official viewpoint of the Oesterreichische Nationalbank or the Eurosystem.

Non-technical summary

Local projections (LPs) have become popular for estimating impulse responses because they are simple and flexible. However, without explicit causal assumptions, they remain “black boxes” and may fail to identify true causal effects. Directed Acyclic Graphs (DAGs) offer a framework for representing causal relationships, making assumptions transparent and helping researchers choose appropriate control variables and instruments.

This paper demonstrates how DAGs can guide the specification of LPs for estimating impulse responses with a causal interpretation. The examples illustrate how shocks propagate, how to avoid common pitfalls—such as conditioning on colliders—and how to apply formal graphical rules to derive appropriate conditioning sets and develop valid instrumental-variable strategies for identification.

The empirical analysis of the impact of uncertainty shocks on German industrial production underscores the practical relevance of these insights. A DAG-based instrumental-variable LP suggests that U.S. uncertainty shocks have substantial negative effects on German industrial production. The results also show that even small changes in the LP specification can produce markedly different impulse responses, highlighting the importance of correct identification strategies. The underlying DAGs help clarify these differences and allow to assess potential biases that arise when causal assumptions are violated.

1 Introduction

The local projection (LP) method (Jordá, 2005) is a popular single-equation alternative to vector autoregressions (VARs) for estimating impulse responses to shocks and interventions (Plagborg-Møller and Wolf, 2021). Its appeal lies in its simplicity, flexibility, and robustness to certain model misspecifications (Jordà, 2023; Jordà and Taylor, 2025). However, without incorporating causal knowledge, LPs do not necessarily identify causal effects and remain "black boxes" (Olea et al., 2025).

By representing qualitative causal knowledge and making the assumed causal structure of a model transparent, directed acyclic graphs (DAGs) help unpack such black boxes. In particular, formal graphical rules can be applied to a DAG to determine whether causal effects can be identified and estimated from the available data (Pearl, 1995, 2009b; Pearl et al., 2016; Chalak and White, 2011; Elwert, 2013; Didelez, 2018; Hünermund and Bareinboim, 2023).

Despite the growing use of DAGs in fields such as epidemiology, machine learning, statistics, and the social sciences, their application in economics—particularly in macroeconomic time series analysis—remains limited.[1] This paper addresses this gap by demonstrating how DAGs can be used to transparently specify LPs for causal analysis.[2]

More precisely, the response of a variable Y_{t+h} at time $t+h$ to a shock in S_t of size δ at time t is defined as:

$$\mathcal{R}_{S\to Y}(h,\delta) \equiv \mathbb{E}[Y_{t+h} \mid S_t = S_0 + \delta; \mathbf{X}_t] - \mathbb{E}[Y_{t+h} \mid S_t = S_0; \mathbf{X}_t], \quad h = 0, 1, \ldots, H. \tag{1}$$

A LP to estimate this response is:

$$Y_{t+h} = \alpha_h + \beta_h S_t + \gamma_h \mathbf{X}_t + v_{t+h}, \tag{2}$$

where β_h captures the response at time $t+h$. The key question is: which conditioning variables $\mathbf{X}_t$ or instruments $\mathbf{Z}_t$ are required to obtain impulse responses that can be causally interpreted?

Using a series of examples, this paper demonstrates how DAGs and simple graphical rules can help answer this question. Furthermore, an empirical analysis applies DAG-based LPs to a pressing macroeconomic issue, namely the impact of economic uncertainty on economic activity.

[1] Hoover (2001), Demiralp and Hoover (2003), and Hoover (2005) are one of the few applications of causal graphs in macroeconomics.

[2] The paper focuses on identification rather than estimation or inference; for a detailed discussion of LP estimation and inference, see Jordà and Taylor (2025). For more on learning and testing causal models, see Peters et al. (2017) and Assaad et al. (2022).

Specifically, the analysis uses a DAG to develop an instrumental variable (IV) based LP for estimating the impact of uncertainty shocks on German industrial production and compares this approach with alternative LP specifications.

Because DAGs are nonparametric, identification strategies can be derived directly from the graph structure without assuming a specific functional form or statistical model (Pearl, 2009b). For simplicity, this paper assumes linear causal relationships and uses linear LP specifications.[3]

It is also important to note that the DAG framework and the potential outcomes (PO) framework (Rubin, 1974) are logically equivalent ways to express causal assumptions (Galles and Pearl, 1998; Pearl, 2015). While the PO framework encodes causal assumptions using counterfactual independencies, DAGs represent these assumptions graphically.

Each framework has distinct advantages. The PO framework is well-suited for defining estimands and interpreting results, while DAGs are particularly effective for expressing causal assumptions and guiding identification strategies (Pearl, 2009a). By showing how DAGs aid in identifying causal effects in applied time series analysis, this paper therefore complements the work of Rambachan and Shephard (2019), who use the PO framework to define causal effects in time series via treatment paths.

As already mentioned, DAGs are still relatively unknown in economics. Therefore, the paper first introduces some DAG basics in Section 2.1 and graphical rules for identifying causal effects in Section 2.2. Simple DAGs are used to illustrate key causal patterns and core concepts such as confounding, adjustment for confounders, and instrumental variables. Section 2.3 introduces DAGs for time series.

The DAGs in Section 3.1 are based on the simple bi-variate structural process discussed in Jordà and Taylor (2025). They illustrate different types of shocks (Ramey, 2016), such as "narrative" or "generated" shocks (Romer and Romer, 1989) and intervention regimes. The diagrams also depict how these shocks propagate through the system and affect the outcome variable of interest.

When all relevant variables in a multivariate time series model are observed, conditioning on an appropriate set of control variables can yield causal responses. Section 3.2 presents a series of examples showing how DAGs can be used to select suitable sets of control variables for LPs. It also shows how carelessly including controls can undermine identification and produce

[3] When implemented with observed shocks or valid proxies, linear LPs recover weighted averages of causal effects, irrespective of underlying nonlinearities (Kolesár and Plagborg-Møller, 2025). For details on nonlinear LPs, see Cloyne et al. (2023), Gonçalves et al. (2024), and Jordà and Taylor (2025).

misleading results.

Section 3.3 addresses scenarios in which some variables are unobserved. In such cases, IV strategies can be employed to identify causal impulse responses (IV-LP). Example DAGs clarify the conditions for instrument validity and illustrate how both external and internal instruments can be selected to obtain causal responses within the IV-LP framework, thereby complementing Stock and Watson (2018). Section 3.4 shows how to isolate specific causal pathways with both standard and IV-based LPs.

An example in Kilian et al. (2025) demonstrates that recursive VARs may fail to recover the impact of uncertainty shocks on GDP growth – or even to provider bounds on the true response. Section 4 presents a DAG that illustrates this issue. Based on this DAG, the paper proposes an IV strategy based on two uncertainty indicators to recover the response to uncertainty shocks.

An empirical analysis of uncertainty shocks on German industrial production implements this IV strategy and compares it with alternative LP specifications. The results suggest that U.S. uncertainty shocks can have substantial negative effects on German industrial production. They also reveal that small changes in LP specification can produce large differences in the estimated impulse responses. The underlying DAGs explain these differences and provide interpretive guidance when the actual causal structure diverges from the assumed one.

2 Graphical Modeling

This section introduces DAGs, key causal patterns, d-separation, and the backdoor criterion for identifying causal effects. For more technical treatments specific to time series, see Eichler and Didelez (2007), Peters et al. (2017), and Runge et al. (2023). A comprehensive reference on causal inference using graphical models is Pearl (2009b). For accessible introductions tailored to economists, see Morgan and Winship (2015), Cunningham (2021) and Huntington-Klein (2021).

2.1 DAG Basics

A DAG encodes qualitative causal assumptions about a data generating process through nodes representing variables and arrows indicating causal relationships. Arrowheads indicate causal direction. A solid arrow between two variables indicates that both are observed. A dashed arrow indicates that at least one variable is unobserved. A missing arrow indicates the absence of a causal relationship.

Figure 1 provides some examples. In DAG (a), the solid arrow from X to Y indicates that

X causes Y and that both variables are observed. In this context, X is called a parent and Y is called a child. The DAG also includes the independent disturbances u_X and u_Y. These are usually omitted, as they do not affect the graphical structure of causal relationships. In DAG (b), X and Y share an unobserved common cause W. A dashed bi-directed arc indicates the presence of further unspecified unobserved common causes.

A *path* is a sequence of nodes connected by arrows, regardless of their direction. For example, in DAG (c) $X \leftarrow W \rightarrow Y$ is a path between X and Y. A *directed path*, such as $X \rightarrow W \rightarrow Y$ in DAG (d), has all arrows pointing in the same direction.

As the name suggests, DAGs contain no feedback loops. Therefore, DAGs do not accommodate cyclical models, such as simultaneous equation models, which may represent the equilibrium state of an underlying acyclic causal structure where feedback between variables occurs on a faster time scale than the measurements. Such models can be represented using directed cyclic graphs (DCGs). In this paper, all models will be acyclic. For an accessible introduction to DCGs, see Kenny (1979). A comprehensive formal treatment of DCGs, their properties, and identifiability in causal inference is provided in Bongers et al. (2021).

There are three key *causal structures* for random variables W, X, and Y:

- *Fork (Fig. 1c):* W is a common cause of X and Y ($X \leftarrow W \rightarrow Y$), generating spurious correlation between X and Y. Conditioning on W (i.e., holding W constant) blocks the path and removes the spurious correlation, making X and Y conditionally independent.

- *Chain(Fig. 1d):* W mediates the effect of X on Y ($X \rightarrow W \rightarrow Y$). Conditioning on W again blocks the causal path, eliminating the effect of X on Y and making the variables conditionally independent.

- *Collider (Fig. 1e):* W is a common effect of X and Y ($X \rightarrow W \leftarrow Y$). The path is blocked unless W is conditioned on. In contrast to the first two cases, conditioning on W opens the path and makes X and Y conditionally dependent. This happens because conditioning on a collider creates selection bias.

Conditioning on a collider or on one of its descendants, like V in Fig 1f, introduces spurious associations. DAGs help avoid such biases by making such structures visible (Elwert and Winship, 2014; Cinelli et al., 2024).

Figure 1: Basic DAGs and Causal Patterns

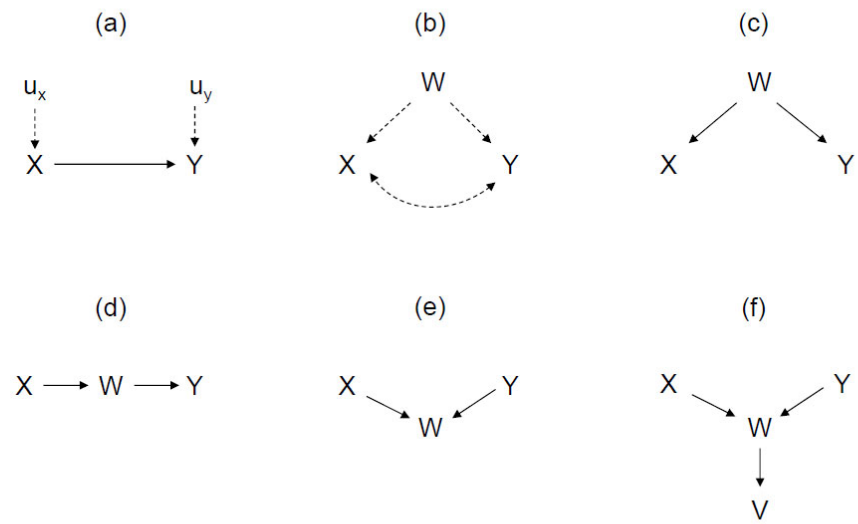

DAG (a) illustrates the causal effect of X on Y along with the independent disturbances u_x and u_y. In DAG (b), W is an unobserved common cause of X and Y, with the dashed bi-directed arc indicating the presence of additional unobserved causes. DAG (c) shows a fork, DAG (d) a chain, and DAG (e) a collider. In DAG (f), V is a descendant of the collider W.

2.2 D-Separation, Backdoor Criterion, Instruments

As a qualitative description of a data generating process, a DAG essentially represents a factorization of the probability distribution of a set of variables, where each variable's probability is conditioned on its parents in the graph. When a probability distribution can be factorized according to a DAG, it is said to be (Markov) *compatible* with the DAG. For example, the joint probability distribution $P(X, Y, W)$ over the variables in DAG (c) can be factorized as $P(W)P(X \mid W)P(Y \mid W)$, and is therefore compatible with that DAG.

D-separation (Pearl (2009b), Theorem 1.2.3): A path in a DAG G can be d-separated (blocked) by conditioning on a set of nodes $\mathbf{C}$ if either (i) the path contains a chain $(... \rightarrow W \rightarrow ...)$ or a fork $(... \leftarrow W \rightarrow ...)$ and the middle node W is in $\mathbf{C}$, or (ii) the path contains a collider $(... \rightarrow W \leftarrow ...)$, and the middle node W or any descendant of W is *not* in $\mathbf{C}$. More generally, if the set $\mathbf{C}$ blocks all paths between two disjoint sets of nodes $\mathbf{A}$ and $\mathbf{B}$ in a DAG G, then $\mathbf{A}$ and $\mathbf{B}$ are d-separated by $\mathbf{C}$, denoted as $(\mathbf{A} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{C})_G$.[4]

D-separation is the bridge between a DAG and compatible probability distributions, because

[4] The "d" in d-separation stands for "directional".

d-separation in a DAG implies conditional independence in all distributions compatible with the DAG. More precisely, $(A \perp\!\!\!\perp B \mid C)_G \Rightarrow (A \perp\!\!\!\perp B \mid C)_P$ in every distribution P compatible with the DAG G (Pearl (2009b), Theorem 1.2.4).

By exploiting d-separation, the backdoor criterion is a key tool for identifying the causal effect of a variable X on another variable Y.[5] It helps to find suitable variables to block all spurious paths between X and Y.

Backdoor Criterion: Given a pair of variables (X, Y) in a DAG, a set of variables $\mathbf{C}$ satisfies the backdoor criterion if:

- No variable in $\mathbf{C}$ is a descendant of X, and
- $\mathbf{C}$ blocks all paths from X to Y that contain an arrow into X.

In DAG (a) in Figure 2, X has both a *direct effect* on Y ($X \to Y$) and an *indirect effect* through V ($X \to V \to Y$). Assuming linearity, the *total effect* of X on Y is the sum of these two effects. The variable W acts as a confounder, introducing spurious correlation via the backdoor path $X \leftarrow W \to Y$.

In this DAG, conditioning on W blocks the backdoor path $X \leftarrow W \to Y$, as indicated by the missing arrows in DAG (b). Therefore, the conditioning set $\mathbf{C} = W$ satisfies the backdoor criterion and the total effect of X on Y is identified. Conditioning on both W and V blocks the backdoor path $X \leftarrow W \to Y$ and the indirect path $X \to V \to Y$, isolating the direct effect $X \to Y$, as shown in DAG (c).

If some variables are unobserved, the backdoor criterion remains valid, but it may not lead to identification if confounders are among the unobserved variables. In such cases, instruments help identify causal effects (Chalak and White, 2011).

A variable Z is a valid *instrument* for estimating the effect of X on Y if Z is correlated with X (relevance) and influences Y only through X (exclusion). This implies that the instrument must be exogenous (i.e., uncorrelated with the disturbance term). In a time-series context, Z must also be uncorrelated with future and past disturbance terms (lead-lag exogeneity). For a single instrument the IV estimator is

$$\beta = \frac{E(YZ)}{E(XZ)}. \tag{3}$$

In DAG (d) in Figure 2, the causal effect of X on Y is confounded by unobserved variables, as indicated by the dashed bi-directed arc, implying that X is correlated with the error term

[5] The name refers to paths between X and Y that point into X and thus enter X via the "back door".

and thus endogenous. In this example, Z can serve as a valid instrument because it directly influences X and affects Y only through X, thereby enabling consistent estimation of the causal effect of X on Y.

In DAG (e), Z is not a valid instrument because the path $Z \to V \to Y$ violates the exclusion condition. However, as shown in DAG (f), conditioning on V turns Z into a valid instrument by blocking the path $Z \to V \to Y$ (Brito and Pearl, 2002). The conditional IV estimator for the causal effect of X on Y is

$$\beta = \frac{E(YZ \mid V)}{E(XZ \mid V)}. \tag{4}$$

Figure 2: Conditioning and Instruments

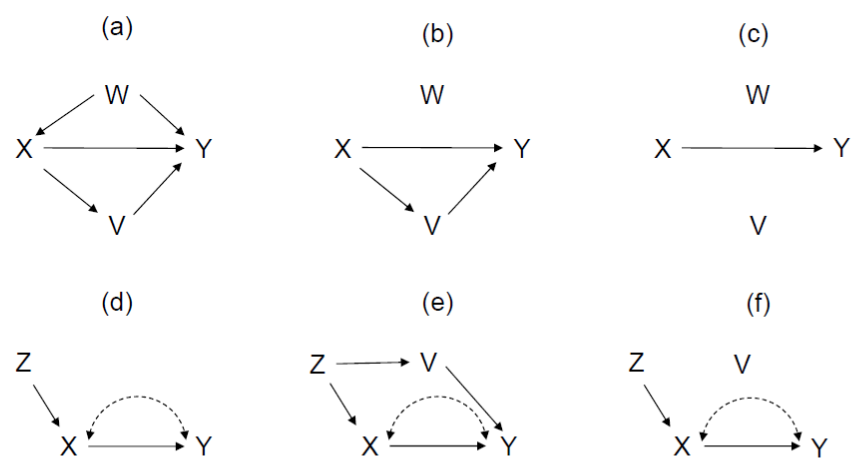

DAG (a) shows the causal effect of X on Y, along with a confounder W and a mediator V. DAG (b) illustrates the result of conditioning on W. DAG (c) shows the result of conditioning on both W and V. DAG (d) introduces a valid instrument Z. In DAG (e), Z is an invalid instrument. DAG (f) demonstrates how Z becomes a valid instrument after conditioning on V.

2.3 DAGs for Time Series

We consider stationary multivariate time series processes that may include relationships between lagged and current variables (e.g., $X_{t-j} \to Y_t$), as well as instantaneous relationships (e.g., $X_t \to Y_t$) but no cycles. Some variables may not be observed. It is assumed that all relationships remain constant in direction throughout time.

Since time series often contain many observations or are assumed to begin in the infinite past, drawing a complete DAG is usually not possible. Instead, a so-called "window DAG" can be used to schematically represent causal relationships in multivariate time series (Assaad et al., 2022).

As an example, consider the system described in Peters et al. (2017), Ch. 10:

$$X_t = a_1 X_{t-1} + a_2 Z_{t-1} + u_{1,t}$$
$$Y_t = b_1 Y_{t-1} + b_2 Z_{t-2} + u_{2,t},$$

where Z_t is unobserved. Figure 3 shows the corresponding DAG. Solid arrows indicate the causal links between the observed variables X_t and Y_t, while dashed arrows represent the causal links emanating from the unobserved variable Z_t. The independent disturbances are omitted for clarity.

The DAG for this system highlights an interesting feature. Due to the unblocked backdoor path $X_t \leftarrow Z_{t-1} \rightarrow Y_{t+1}$ between X_t and Y_{t+1}, a naive application of Granger causality tests would incorrectly suggest that X_t causes Y_{t+1} and that Y_t does not cause X_{t+1}. For the same reason, a naive application of LPs would incorrectly suggest that Y responds to shocks in X.

Figure 3: Window DAG for Time Series

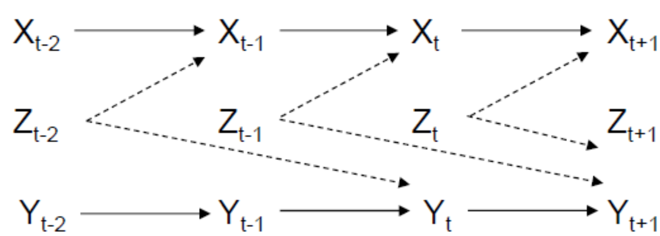

The DAG shows a trivariate time series (X_t, Y_t, Z_t) where Z_t is unobserved.

3 DAGs for Local Projections

This section first describes different types of shocks and their propagation, and then shows how DAGs help to find appropriate conditioning sets and instruments for estimating causal responses with LPs.

3.1 Shocks and Their Propagation

The variable S_t in the LP

$$Y_{t+h} = \alpha_h + \beta_h S_t + \gamma_h \mathbf{X}_t + v_{t+h}$$

can represent different types of shocks.

S_t may represent an exogenous shock. If this shock is constructed from historical records, it is often called a “narrative” shock (Romer and Romer, 1989). When the shock is truly exogenous, the response can be estimated using a LP without any control variables.

S_t may represent an “external” shock $\hat{u}_{j,t} = f(g_{j,t}, \hat{\varphi}_j)$, generated with variables $g_{j,t}$ that lie outside the VAR system. Since $\hat{u}_{j,t}$ is a generated regressor (Pagan, 1984), the LP must include the variables $g_{j,t}$ for valid inference (Breitung and Brüggemann, 2023). External shocks may also serve as instruments in IV-LPs (Stock and Watson, 2018).

Finally, S_t may represent a variable within the VAR. If the shock in S_t can be expressed as a linear combination of the system variables, the shock is often called “fundamental” or “internal”. If S_t is a system variable, the LP must typically include a suitable set of control variables or S_t must be instrumented.

Some examples. Following Jordà and Taylor (2025), let us consider the structural bi-variate process,

$$\begin{pmatrix} 1 & -\beta \\ 0 & 1 \end{pmatrix} \begin{pmatrix} Y_t \\ S_t \end{pmatrix} = \begin{pmatrix} \phi_{yy} & \phi_{ys} \\ \phi_{sy} & \phi_{ss} \end{pmatrix} \begin{pmatrix} Y_{t-1} \\ S_{t-1} \end{pmatrix} + \begin{pmatrix} u_{y,t} \\ u_{s,t} \end{pmatrix}, \tag{5}$$

where $u_{y,t}$ and $u_{s,t}$ are uncorrelated zero-mean disturbances. The DAGs in Figure 4 describe the shocks and their propagation that arise in four distinct configurations of ϕ-coefficients.

In DAG (a), all ϕ coefficients are zero. Consequently, the shock $S_t = u_{s,t}$ is exogenous and does not propagate to future values of Y_{t+h}. Since there are no backdoor paths, a simple regression of Y_t on S_t identifies β_0 at $h = 0$, and $\beta_h = 0$ for $h > 0$. If $u_{s,t}$ can be perfectly measured, regressing Y_t on $u_{s,t}$ yields equivalent (scaled) results.

DAG (b) shows the case where $\phi_{ss} \neq 0$, while $\phi_{yy} = \phi_{ys} = \phi_{sy} = 0$. Examples include a series of interventions, or a self-propagating shock to stock returns. As before, there are no open backdoor paths between S_t and Y_{t+h}. Therefore, a simple regression of Y_{t+h} on S_t estimates the overall response of Y_{t+h}.

In DAG (c), $\phi_{yy} \neq 0$, while $\phi_{ys} = \phi_{sy} = \phi_{ss} = 0$. Here, too, an exogenous intervention or shock affects both current and future outcomes, but through a different causal mechanism. The response originates from a single change in S_t and propagates solely through the outcomes. Since there are no open backdoor paths, a regression of Y_{t+h} on S_t again provides a valid estimate of the total response of Y_{t+h} to the shock in S_t.

In DAG (d), all ϕ-coefficients are nonzero. Consequently, a shock to S_t affects both the current and future values of S and Y through direct and indirect channels, thus generating dynamic feedback. The LP specification depends on whether the original shock $u_{s,t}$ or S_t is used.

If the shock $u_{s,t}$ can be perfectly measured, Y_{t+h} can be regressed on $u_{s,t}$ without any controls because S_t acts as a collider between S_{t-1} and Y_{t-1}, blocking all backdoor paths. When the shock is generated, the LP must include the variables used to construct it to ensure valid inference.

If S_t is used in the LP, the backdoor paths $S_t \leftarrow S_{t-1} \rightarrow Y_{t-1} \rightarrow Y_t$, $S_t \leftarrow S_{t-1} \rightarrow Y_t$, and $S_t \leftarrow Y_{t-1} \rightarrow Y_t$ must be blocked. $\mathbf{X}_t = (S_{t-1}, Y_{t-1})$ blocks these paths and satisfies the backdoor criterion, thereby identifying the causal response at horizon h.

Figure 4: Shock Propagation

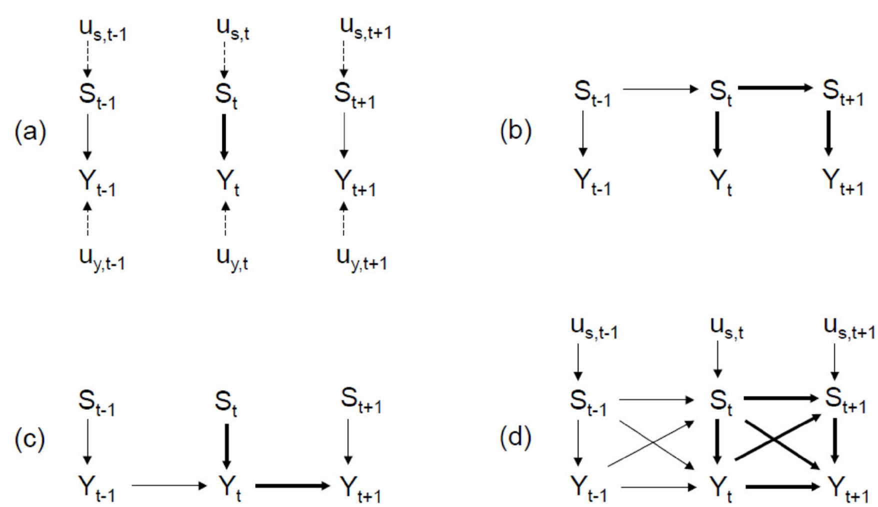

The DAGs shows different shocks and their propagation. The responses are represented by bold arrows.

3.2 Causal Local Projections using Conditioning

As we have just seen, conditioning on an appropriate set of variables can identify causal responses when all relevant variables are observed. The DAGs in Figure 5 provide additional examples. In the diagrams, S is the shocked variable and Y is the response variable. Bold arrows indicate how the shock propagates through the system.

In DAG (a), both S and Y depend on their own lags and the lag of the other variable. No contemporaneous causal link exists between X_t and Y_t. Applying d-separation reveals five conditioning sets that satisfy the backdoor criterion: (S_{t-1}, Y_{t-1}, Y_t), (S_{t-1}, Y_{t-1}), (S_{t-1}, Y_t), (Y_{t-1}, Y_t), and (Y_t). Each of these sets can be used to get valid causal responses of Y to shocks in S. These sets are valid for any horizon h. For instance, at $h = 2$, the full response includes the paths $S_t \to S_{t+1} \to Y_{t+2}$ and $S_t \to Y_{t+1} \to Y_{t+2}$, and is identified by each of the listed control sets.

In DAG (b), the system is recursive with S_t contemporaneously causing Y_t – a common setup in empirical applications. Two sets of controls satisfy the backdoor criterion and yield valid causal responses: (S_{t-1}, Y_{t-1}) and (Y_{t-1}). Crucially, one must *not* condition on Y_t, because doing so blocks the path $S_t \to Y_t \to Y_{t+1}$ and thereby prevents the identification of the full causal effect.

In DAG (c) the structure is reversed: Y_t contemporaneously causes S_t. The following conditioning sets satisfy the backdoor criterion: (S_{t-1}, Y_{t-1}, Y_t), (S_{t-1}, Y_t), and (Y_{t-1}, Y_t). In contrast to the previous case, it is essential to include Y_t as a control in $\mathbf{X}_t$ to block the backdoor path $S_t \leftarrow Y_t \to Y_{t+1}$ that would otherwise confound the estimation of the causal response.

DAG (d) shows a tri-variate system in which S_t affects Y_{t+1} indirectly through a mediator W_t. Conditioning on S_{t-1} is sufficient to identify the causal effect of a shock in S_t on future Y_{t+h}. Although larger conditioning sets are possible, these sets must *not* include the mediator W_t, because doing so would block the causal path $S_t \to W_t \to Y_{t+1}$ and open the spurious path $S_t \to W_t \leftarrow Y_t \to Y_{t+1}$, thereby undermining identification.

Figure 5: Conditioning

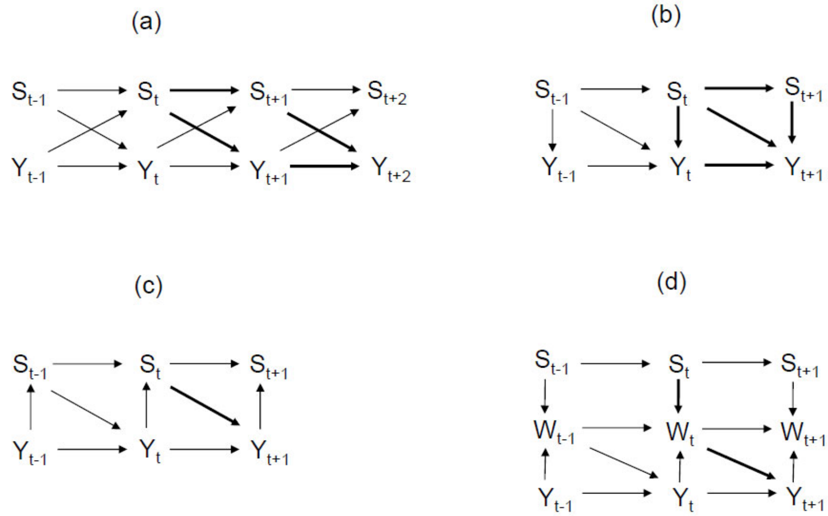

The DAGs provide examples where conditioning yields causal responses. S_t is the shocked variable. Bold arrows indicate the response.

3.3 Causal Local Projections with Instruments

Backdoor paths between the intervention variable and the outcome variable often cannot be blocked when key variables in the system are unobserved. This typically occurs when important variables must be omitted because they are difficult to measure or the system is too complex.

Provided that valid instruments are available, causal responses can be estimated using an IV-LP. With a single instrument Z_t for S_t, an estimator of the response at horizons h is

$$\beta_h = \frac{E(Y_{t+h}Z_t)}{E(S_tZ_t)}. \tag{6}$$

Figure 5 shows two DAGs for systems that require an IV-LP approach. In DAG (a), unobserved variables cause both S_t and Y_{t+1}, as indicated by the bi-directed arch $S_t \leftarrow\text{----}\rightarrow Y_{t+1}$. These backdoor paths between S_t and Y_{t+1} cannot be blocked. To recover the impulse responses, one can use either an external instrument Z_t from outside the system or an internal instrument such as S_{t-1}.

As already mentioned, for time series an external instrument must meet the following criteria: it must be relevant (i.e., correlated with S_t) and satisfy the exclusion restriction (i.e., affect the outcome only through S_t), possibly conditional on a proper set of control variables. This implies that the instrument must be contemporaneously exogenous and lead-lag exogenous (Stock and Watson, 2018).

In DAG (a) of Figure 6, Z_t satisfies these conditions, and the IV estimator (6) yields the causal response of Y_{t+h} to a shock in S_t. The IV-LP correctly identifies the response at all horizons h. For example, at $h = 2$, the path $S_t \rightarrow S_{t+1} \leftarrow\text{-}\text{-}\text{-}\rightarrow Y_{t+1}$ remains blocked because S_{t+1} is a collider. No backdoor paths enter S_t via S_{t-1}, since Y_{t+1} is a collider relative to Y_t, and S_t does not cause Y_t. Only the paths $S_t \rightarrow Y_{t+1} \rightarrow Y_{t+2}$ and $S_t \rightarrow S_{t+1} \rightarrow Y_{t+2}$ – which the instrument Z_t isolates – remain open. These two paths represent the total response of Y_{t+2} to a shock in S_t.

Looking at DAG (a) reveals that conditioning on Y_t turns S_{t-1} into a valid instrument, as it isolates the path $S_{t-1} \rightarrow S_t \rightarrow Y_{t+1}$ while blocking all other paths from S_{t-1} to Y_{t+1}. The conditional IV-LP, which uses S_{t-1} as an instrument for S_t, is given by:

$$\beta_h = \frac{E(Y_{t+h}S_{t-1} \mid Y_t)}{E(S_t S_{t-1} \mid Y_t)}. \tag{7}$$

As before, the conditional IV-LP identifies the causal responses at all horizons $h > 0$.

DAG (b) in Figure 6 shows a system where omitted variables affect both current and future values of S, and S_t also directly influences Y_t. The immediate response at $h = 0$ can be estimated using a standard LP with either S_{t-1} or (S_{t-1}, Y_{t-1}) as controls. However, for $h > 0$, an IV-LP is required, because the spurious correlations between current and future S cannot be blocked. In this setting, S_{t-1} is not a valid internal instrument because it directly affects Y_t, violating the exclusion restriction. The external instrument Z_t remains valid, however, and the IV estimator in equation (6) yields causal responses at all horizons.

Figure 6: Instrumental Variables

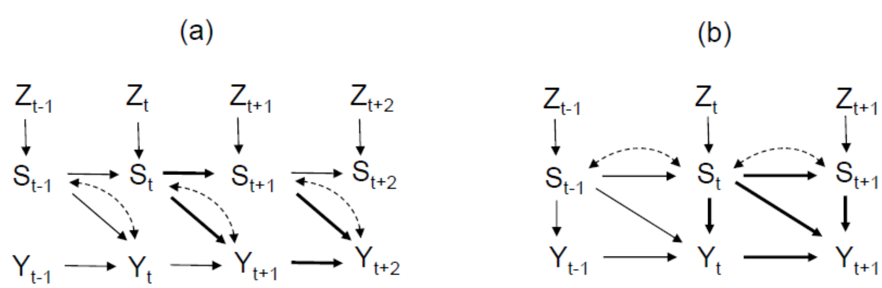

The DAGs shows time series where IV strategies yield causal responses. S_t is the shocked variable. Bold arrows indicate responses.

3.4 Isolating Pathways

Sometimes certain components of a full response are of special interest. For example, one might wish to compare the direct effect of a financial stimulus on consumption with its indirect multiplier effect. To make such comparisons, it is necessary to estimate both the total and the direct response to the intervention. LPs for direct responses can be specified by applying d-separation to the DAG that encodes the assumed causal structure. Such LPs typically contain leads of certain variables as controls.

For example, in traditional LP settings – such as in DAG (a) in Figure 5 – one must, in addition to one of the identifying sets of controls, also condition on Y_{t+1} to obtain an estimate of the direct effect $S_t \rightarrow S_{t+1} \rightarrow Y_{t+2}$ of the interventions at times t and $t+1$ on Y_{t+2}. In DAG (b) one must in addition condition on Y_t and S_{t+1} to estimate the direct effect of the intervention S_t on Y_{t+1}.

Isolating specific pathways works analogously in IV-LP settings such as in Figure 6. In DAG (a), one must condition on Y_{t+1} to estimate the effect of changes in S_t and S_{t+1} on Y_{t+2}. Likewise, in DAG (b), one must condition on Y_t and S_{t+1} to obtain the direct effect of a change in S_t on Y_{t+1}.

4 Uncertainty Shocks

Kilian et al. (2025) present an example where recursive VARs neither capture the true response nor provide valid bounds for the impact of uncertainty shocks on GDP growth. We now consider a variant of this example where uncertainty and GDP growth are contemporaneously causally unrelated but are both influenced by uncertainty shocks and shocks to GDP growth. The corresponding DAG highlights that in this causal structure the LP must use either the true uncertainty shock or valid instruments for the shock to fully uncover the causal response.

The first part of this section discusses the example and presents an IV strategy to recover the true response. The second part uses the proposed IV strategy to empirically examine the impact of uncertainty shocks on German industrial production.

4.1 Example of the Impact of Uncertainty Shocks on GDP Growth

Consider two reduced form shocks $u_{gdp,t}$ and $u_{unc,t}$ affecting GDP growth and a measure of uncertainty. Using the same numbers as in Kilian et al. (2025), these shocks are linked to the structural shocks $w_{other,t}$ and $w_{unc,t}$ via the following matrix B:

$$\begin{bmatrix} u_{\text{gdp, t}} \\ u_{\text{unc, t}} \end{bmatrix} = \begin{bmatrix} 1 & 0.5 \\ -0.75 & 1.5 \end{bmatrix} \begin{bmatrix} w_{\text{other, t}} \\ w_{\text{unc, t}} \end{bmatrix}, \tag{8}$$

$$\Sigma = BB^{'} = \begin{bmatrix} 1.25 & 0 \\ 0 & 2.8125 \end{bmatrix}, \tag{9}$$

$$\text{chol}(\Sigma) = \begin{bmatrix} 1.118 & 0 \\ 0 & 1.6771 \end{bmatrix}. \tag{10}$$

Thus, an uncertainty shock, $w_{\text{unc},t}$, increases GDP growth on impact by 0.5.

The original example in Kilian et al. (2025) assumes a simultaneous relationship between a measure of uncertainty and GDP growth. In contrast, this example directly models the impacts of the shocks.[6] Both shocks jointly affect GDP growth, denoted as GDP_t, and the uncertainty measure UNC_t. There is no simultaneous relationship between GDP_t and UNC_t because these variables are not observed in real time. However, their measurements for time t may influence uncertainty and GDP growth in the subsequent period.

The DAG in Figure 7 shows the causal structure of the example. At time t, both UNC_t and GDP_t are jointly influenced by the uncertainty shock $w_{\text{unc},t}$ and the GDP growth shock $w_{\text{other},t}$. There is no recursive structure because there is no contemporaneous causal relation between UNC_t and GDP_t. Furthermore, the variance-covariance matrix Σ of the reduced-form errors is diagonal. Therefore, as in the original example by Kilian et al. (2025), no recursive VAR can correctly recover or bound the true response of GDP growth to an uncertainty shock.

The DAG also reveals another problem that is otherwise harder to see: to fully recover the response to an uncertainty shock, one should use the uncertainty shock $w_{\text{unc},t}$ itself rather than UNC_t, the variable that measures the uncertainty, since the latter has no contemporaneous impact on GDP_t. How can this be done when the shock cannot be directly observed?

A solution is to use two indicators of the uncertainty shock, $I_{1,t} = \delta_0 + \delta_1 w_{\text{unc},t} + \epsilon_{1,t}$ and $I_{2,t} = \rho_0 + \rho_1 w_{\text{unc},t} + \epsilon_{2,t}$, and instrument one with the other in an IV-LP. For this strategy to work, the covariances $Cov(w_{\text{unc},t}, \epsilon_{1,t})$, $Cov(w_{\text{unc},t}, \epsilon_{2,t})$, and $Cov(\epsilon_{1,t}, \epsilon_{2,t})$ must be zero, and $\epsilon_{1,t}$ and $\epsilon_{2,t}$ must be uncorrelated with uncertainty and GDP growth (cf. Wooldridge (2010)).

[6] In SVAR terminology, the original example specifies an A-model that focuses on the instantaneous relationships between the observed variables, whereas our example considers a B-model that directly focuses on the unexpected shocks to the variables. See, Ch. 9 in Lütkepohl (2005) for further details.

Under these conditions, running the IV-LP

$$Y_{t+h} = \alpha_h + \beta_h I_{1,t} + v_{t+h}, \tag{11}$$

using $I_{2,t}$ as an instrument for $I_{1,t}$ yields the full causal response to the uncertainty shock.

Whether this identification strategy works in practice depends on the ability to obtain two measurements of the uncertainty shock with uncorrelated measurement errors. In this context, it is important not to use a domestic uncertainty indicator. As Figure 7 shows, the error in this indicator includes the shock to GDP growth and therefore correlates with GDP growth.

A simulated VAR(1) model illustrates the proposed IV strategy. The simulation uses the structural impact matrix B from equation (7) and a coefficient matrix A_1 on the lagged variables such that both GDP growth and uncertainty are moderately persistent (see, Appendix 1). The simulation generates a sample of 1,000 observations, using the last 500 for analysis. The responses are estimated using six different methods: a LP with the true structural shock, an IV-LP using two noisy indicators of the structural uncertainty shock, an IV-LP instrumenting the uncertainty variable, a standard LP using the uncertainty variable, and two recursive VARs with different orderings.

As can be seen in Figure 8, only the LP using the true structural uncertainty shock and the IV-LP with two indicators of the shock accurately recover the causal response. Both methods correctly estimate the contemporaneous impact ($h = 0$) on GDP growth, which is 0.5 in the simulated model. All the other methods produce incorrect results. The IV-LP that directly instruments the uncertainty variable underestimates the response, while the standard LP and the recursive VARs estimate it as zero on impact. Moreover, as emphasized in Kilian et al. (2025), the recursive VARs do not provide valid bounds for the true response.

Figure 7: DAG of Example for Uncertainty Shock on GDP Growth

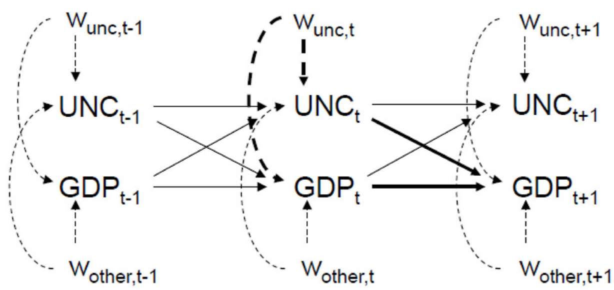

The DAG shows the causal structure of an uncertainty shock on GDP growth in the example in Kilian et al. (2025).

Figure 8: Estimated Responses to an Uncertainty Shock in a Simulated SVAR

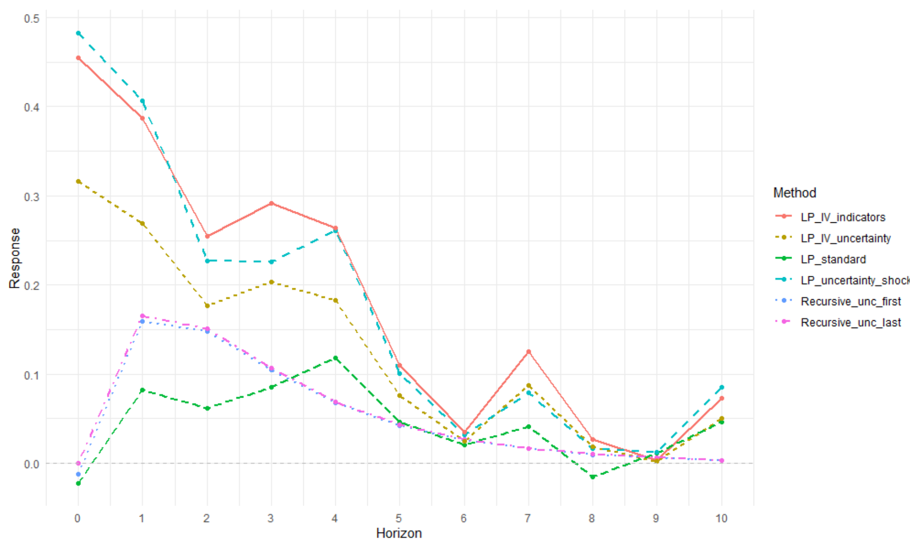

The figure shows the response of GDP growth to an uncertainty shock in a simulated SVAR, estimated using six methods: an LP with the true structural uncertainty shock (`LP_structural_shock`), an IV-LP with indicators of the structural shock (`LP_IV_indicators`), an IV-LP instrumenting the uncertainty variable (`LP_IV_uncertainty`), a standard LP using the uncertainty variable (`LP_standard`), and two recursive VARs with uncertainty ordered first (`Recursive_unc_first`) and last (`Recursive_unc_last`).

4.2 The Impact of Uncertainty Shocks on German Industrial Production

This empirical analysis examines the impact of uncertainty shocks on German industrial production. By estimating four different LPs, it demonstrates that slightly different specifications can produce very different impulse responses. The underlying DAGs help explain these differences.

The results obtained with an IV-LP using two uncertainty indicators suggest that the impact of U.S. uncertainty shocks on German industrial production could be substantial.

The first LP estimates the response of German industrial production to U.S. uncertainty shocks using the IV strategy outlined above. DAG (a) in Figure 9 depicts the assumed causal structure. The bold arrows indicate the paths that are isolated by the IV strategy. The IV-LP is given by

$$ind_{de,t+h} = \alpha_h + \beta_h epu_{us,t} + \gamma_h \mathbf{X}_t + v_{t+h}, \tag{12}$$

where $ind_{de,t+h}$ is the log of German industrial production at time $t+h$. The two indicators for U.S. uncertainty shocks are the unexpected variation in the log of the news based U.S. economic policy uncertainty index ($epu_{us,t}$) of Baker et al. (2016) and the log of the CBOE Volatility Index (vix_t). The later serves as an instrument for $epu_{us,t}$. It is assumed that the measurement errors in the indicators are uncorrelated with German Industrial Production and German EPU. Additionally, it is assumed that uncertainty originating in Germany does not influence U.S. uncertainty. The vector of control variables is defined as $\mathbf{X}_t = (vix_{t-1}, epu_{us,t-1}, epu_{de,t-1}, ind_{de,t-1})$, since this set of controls blocks all backdoor paths and turns vix_t into a valid instrument. This vector of controls also isolates the unpredicted variation in the uncertainty indicators, such that the conditional correlation between the indicators captures an unexpected uncertainty shock.

In the second IV-LP,

$$ind_{de,t+h} = \alpha_h + \beta_h epu_{de,t} + \gamma_h \mathbf{X}_t + v_{t+h}, \tag{13}$$

uses the same control variables as above, but now the log of the German EPU is instrumented by vix_t. This LP demonstrates that the German EPU should not be instrumented in the causal structure described in DAG (a) because the error in German EPU is correlated with the error in German industrial production via the path $EPU_{DE,t} \leftarrow w_{local,t} \rightarrow IND_{DE,t}$. This violates the requirement that the errors of both indicators are uncorrelated with German industrial production.

The third and fourth LPs are based on OLS and assume the recursive causal structure depicted in DAG (b) of Figure 9. In this structure, U.S. uncertainty–measured by the U.S. EPU–directly influences both German EPU and German industrial production. Meanwhile, German EPU affects only German industrial production and has no impact on U.S. EPU.

The third LP treats $epu_{de,t}$ as the shock variable and includes the control vector $\mathbf{X}_t = (epu_{us,t}, epu_{us,t-1}, epu_{de,t-1}, ind_{de,t-1})$ to block all backdoor paths. Under the assumed causal structure, the LP estimates the response to a shock in German EPU.

The fourth LP inserts $epu_{us,t}$ as the shock variable. The control vector is defined as $\mathbf{X}_t = (epu_{us,t-1}, epu_{de,t-1}, ind_{de,t-1})$. Under the assumed causal structure, this LP estimates the response of German industrial production to a shock in U.S. EPU.

The dataset consists of monthly observations from January 1993 to March 2025. All uncertainty indicators come from the FRED database of the Federal Reserve Bank of St. Louis. German industrial production comes from the Federal Statistical Office of Germany. Table 1 presents summary statistics for reference.

Figure 10 shows the impulse responses obtained with the different LP specifications, along with 90% confidence intervals based on autocorrelation and heteroskedasticity robust standard errors (Newey and West, 1987). Since all variables are in logarithmic form, the estimates reflect (approximate) percentage changes in industrial production in response to a 1% change in uncertainty.

All estimated responses indicate a negative impact of uncertainty shocks on German industrial production. However, the results vary considerably across the LP specifications. The IV-LP for U.S. uncertainty that uses the two U.S. uncertainty indicators yields the strongest response. Industrial production falls sharply – by about -0.16% in the month after the shock and then by about -0.12% in the following months. In comparison, the IV-LP using the German EPU shows a peak response of about -0.11% in the month following the shock and approximately -0.08% thereafter. The two OLS-based LPs yield much smaller responses: the response to German EPU is virtually zero, while the response to U.S. EPU hovers around -0.01%.

From the perspective of the causal structure described by DAG (a), the differences in the responses can be explained as follows: The smaller response in the IV-LP using German EPU arises because instrumenting German EPU is incorrect–its measurement error contains the shock $w_{local,t}$ and correlates with German industrial production. If this correlation is negative, the estimated response is biased toward zero because the measurement error moves in the opposite direction of the true uncertainty shock, making the measured uncertainty shock appear smaller than it actually is. For the same reason, and due to endogeneity, the response from the third LP is also biased towards zero. Similarly, the small response in the fourth LP may result from endogeneity and measurement error in U.S. EPU.

DAGs not only communicate assumed causal structures but also aid in assessing potential bias when some assumptions are violated. For example, the assumption concerning the macroeconomic shock affecting German industrial production might be violated. The DAG assumes that this shock is local, while the U.S. uncertainty shock is considered global. However, if the macro shock is global, we need to include an arrow pointing to U.S. uncertainty. In that case, the measurement errors in U.S. uncertainty would also be correlated with German industrial production. As discussed above, such a negative correlation would introduce a downward bias in the estimated response. Consequently, the true response to a U.S. uncertainty shock could be even stronger than what is estimated using the IV-LP approach.

Figure 9: DAGs for Uncertainty Shock on German Industrial Production

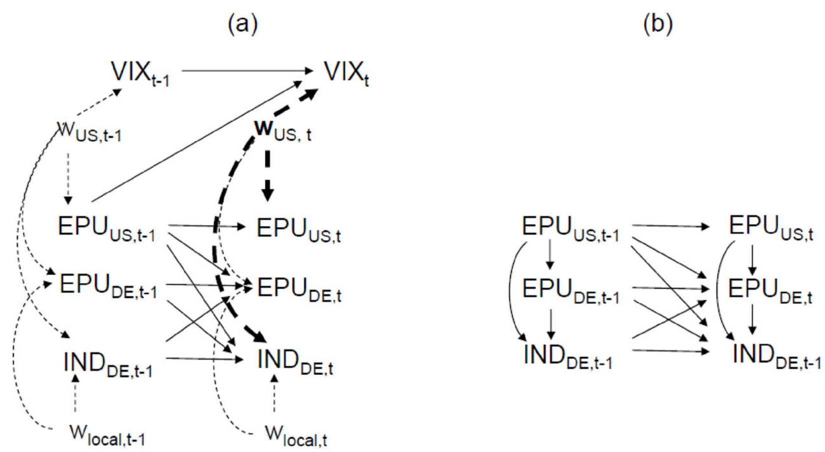

The DAG shows the DAGs for two causal structures for uncertainty shocks on German Industrial Production.

Figure 10: Impulse Responses of German Industrial Production to EPU Shocks

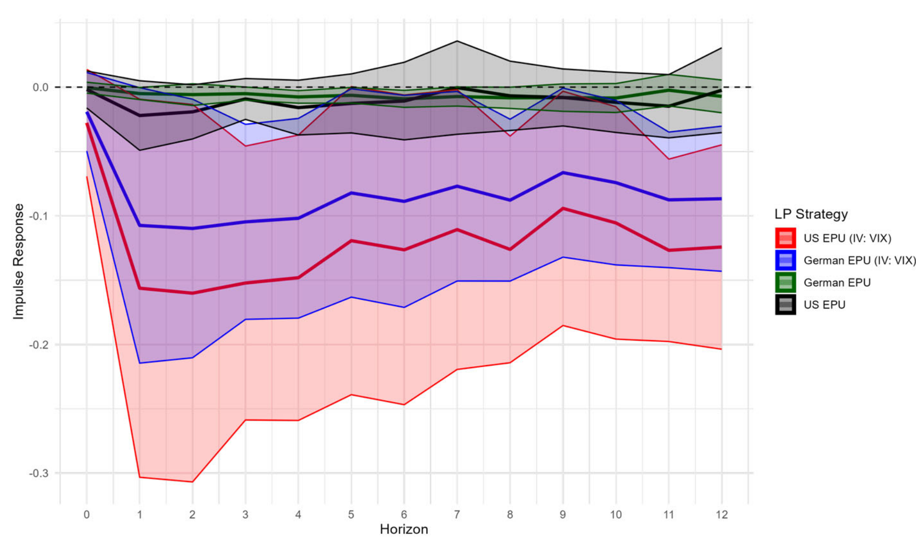

5 Conclusions

Without incorporating causal assumptions, impulse responses often lack a clear causal interpretation. This paper demonstrated how DAGs provide a transparent framework for embedding causal assumptions and selecting control variables and instruments to identify causal impulse responses within the LP approach.

A series of DAGs illustrated the propagation of shocks, how to avoid pitfalls–such as conditioning on colliders–and how to derive conditioning and instrumental variables strategies to isolate both total and pathway-specific responses to policy interventions and shocks.

A DAG for uncertainty shocks highlighted the limitations of traditional recursive VARs discussed in Kilian et al. (2025) and revealed an instrumental variable strategy for identifying the causal response to an uncertainty shock–insights that would be more difficult to obtain without expressing the assumed causal structure with a DAG.

The empirical analysis of the effects of uncertainty shocks on German industrial production underscores the practical relevance of these insights, as the proposed IV approach suggests that U.S. uncertainty shocks have a substantial negative impact on German industrial production. The results further demonstrate that even small changes in the LP specification based on different causal assumptions can lead to very different impulse responses. The underlying DAGs help explain these differences and to assess potential bias when causal assumptions are violated.

Beyond methodological rigor, DAGs also help improve empirical practice. As demonstrated in this study, they can be used not only to justify identification choices but also to communicate assumptions more effectively, thereby enhancing reproducibility and transparency. This approach thus aligns with the growing focus on credibility and robustness in applied econometrics and macroeconomic research.

Appendix 1: Simulated VAR Model

The simulated VAR(1) model is

$$y_t = A_1 y_{t-1} + u_t$$

$$u_t = Bw_t$$

$$A_1 = \begin{bmatrix} 0.5 & 0.1 \\ 0.2 & 0.4 \end{bmatrix}$$

$$B = \begin{bmatrix} 1 & 0.5 \\ -0.75 & 1.5 \end{bmatrix}$$

$$w_t \sim \mathcal{N}(0, I).$$

Thus the equations for output and uncertainty are:

$$\begin{aligned} \text{output}_t &= 0.5 \cdot \text{output}_{t-1} + 0.1 \cdot \text{uncertainty}_{t-1} + w_{1,t} + 0.5 w_{2,t} \\ \text{uncertainty}_t &= 0.2 \cdot \text{output}_{t-1} + 0.4 \cdot \text{uncertainty}_{t-1} - 0.75 w_{1,t} + 1.5 w_{2,t}. \end{aligned}$$

The two noisy indicators of the uncertainty shock $w_{2,t}$ are:

$$\begin{aligned} \text{indicator1}_t &= w_{2,t} + \eta_{1,t}, \quad \eta_{1,t} \sim \mathcal{N}(0, 0.25) \\ \text{indicator2}_t &= w_{2,t} + \eta_{2,t}, \quad \eta_{2,t} \sim \mathcal{N}(0, 0.25). \end{aligned}$$

Appendix 2: Summary Statistics

Table 1: Descriptive Statistics for U.S. and German Variables

Variable	Nobs	Min	Max	Mean	Median	SD	Skewness	Kurtosis
VIX	387	10.13	62.67	19.51	17.69	7.68	2.01	9.56
EPU_{US}	387	57.20	350.46	116.36	105.70	43.83	1.60	6.80
EPU_{DE}	387	28.43	1,095.93	193.22	124.33	189.44	2.41	8.73
IND_{DE}	387	69.80	111.10	92.32	94.10	10.97	-0.31	1.83

Notes: This table reports summary statistics for monthly observations from January 1993 to March 2025. VIX is the CBOE Volatility Index, EPU_{US} and EPU_{DE} are the Economic Policy Uncertainty indices for the United States and Germany, respectively, and IND_{DE} is German industrial production. The statistics include the number of observations (Nobs), minimum and maximum values, mean, median, standard deviation (SD), skewness, and kurtosis.

References

Assaad, C. K., Devijver, E., and Gaussier, E. (2022). Survey and evaluation of causal discovery methods for time series. *Journal of Artificial Intelligence Research*, 73:767–819.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Bongers, S., Forré, P., Peters, J., and Mooij, J. M. (2021). Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915.

Breitung, J. and Brüggemann, R. (2023). Projection estimators for structural impulse responses. *Oxford Bulletin of Economics and Statistics*, 85(6):1320–1340.

Brito, C. and Pearl, J. (2002). Generalized instrumental variables. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, pages 85–93.

Chalak, K. and White, H. (2011). An extended class of instrumental variables for the estimation of causal effects. *Canadian Journal of Economics*, 44(1):1–51.

Cinelli, C., Forney, A., and Pearl, J. (2024). A crash course in good and bad controls. *Sociological Methods & Research*, 53(3):1071–1104.

Cloyne, J., Jordà, Ò., and Taylor, A. M. (2023). State-dependent local projections: Understanding impulse response heterogeneity. Technical report, National Bureau of Economic Research.

Cunningham, S. (2021). *Causal Inference: The Mixtape*. Yale University Press.

Demiralp, S. and Hoover, K. D. (2003). Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and Statistics*, 65(Supplement):745–767.

Didelez, V. (2018). Causal concepts and graphical models. In Maathuis, M., Drton, M., Lauritzen, S., and Wainwright, M., editors, *Handbook of Graphical Models*, pages 353–380. CRC Press.

Eichler, M. and Didelez, V. (2007). Causal reasoning in graphical time series models. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, pages 109–116.

Elwert, F. (2013). Graphical causal models. In Morgan, S. L., editor, *Handbook of Causal Analysis for Social Research*, pages 245–273. Springer, Dordrecht.

Elwert, F. and Winship, C. (2014). Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology*, 40(1):31–53.

Galles, D. and Pearl, J. (1998). An axiomatic characterization of causal counterfactuals. *Foundations of Science*, 3(1):151–182.

Gonçalves, S., Herrera, A. M., Kilian, L., and Pesavento, E. (2024). State-dependent local projections. *Journal of Econometrics*, 244(2):105702.

Hoover, K. D. (2001). *Causality in Macroeconomics*. Cambridge University Press, Cambridge.

Hoover, K. D. (2005). Automatic inference of the contemporaneous causal order of a system of equations. *Econometric Theory*, 21(1):69–77.

Huntington-Klein, N. (2021). *The Effect: An Introduction to Research Design and Causality*. Chapman and Hall/CRC.

Hünermund, P. and Bareinboim, E. (2023). Causal inference and data fusion in econometrics. *The Econometrics Journal*, 28(1):41–82.

Jordá, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Jordà, Ò. and Taylor, A. M. (2025). Local projections. *Journal of Economic Literature*, 63(1):59–110.

Jordà, (2023). Local projections for applied economics. *Annual Review of Economics*, 15:607–631.

Kenny, D. A. (1979). *Correlation and Causality*. Wiley-Interscience, New York.

Kilian, L., Plante, M. D., and Richter, A. W. (2025). Macroeconomic responses to uncertainty shocks: the perils of recursive orderings. *Journal of Applied Econometrics*, 40:395–410.

Kolesár, M. and Plagborg-Møller, M. (2025). Dynamic causal effects in a nonlinear world: the good, the bad, and the ugly. *Journal of Business & Economic Statistics*, 43(4):737–754.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Springer, Berlin, Heidelberg.

Morgan, S. L. and Winship, C. (2015). *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press, 2 edition.

Newey, Whitney, K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Olea, J. L. M., Plagborg-Møller, M., and Wolf, C. K. (2025). Local projections or VARs? A primer for macroeconomists. *arXiv preprint*, arXiv:2503.17144.

Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, 25(1):221–247.

Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.

Pearl, J. (2009a). Causal inference in statistics: An overview. *Statistics Surveys*, 3:96 – 146.

Pearl, J. (2009b). *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2 edition.

Pearl, J. (2015). Trygve Haavelmo and the emergence of causal calculus. *Econometric Theory*, 31(1):152–179.

Pearl, J., Glymour, M., and Jewell, N. P. (2016). *Causal Inference in Statistics: A Primer*. Wiley, Hoboken, NJ.

Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press.

Plagborg-Møller, M. and Wolf, C. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Rambachan, A. and Shephard, N. (2019). Econometric analysis of potential outcomes time series: instruments, shocks, linearity and the causal response function. *arXiv preprint arXiv:1903.01637*.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.

Romer, C. D. and Romer, D. H. (1989). Does monetary policy matter? A new test in the spirit of Friedman and Schwartz. *NBER Macroeconomics Annual*, 4:121–170.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688.

Runge, J., Gerhardus, A., Varando, G., Eyring, V., and Camps-Valls, G. (2023). Causal inference for time series. *Nature Reviews Earth & Environment*, 4:487–505.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. The MIT Press, Cambridge, Massachusetts, 2 edition.

The Working Paper series of the Oesterreichische Nationalbank is designed to disseminate and to provide a platform for discussion of either work of the staff of the OeNB economists or outside contributors on topics which are of special interest to the OeNB. To ensure the high quality of their content, the contributions are subjected to an international refereeing process. The opinions are strictly those of the authors and do in no way commit the OeNB.

The Working Papers are also available on our website (http://www.oenb.at) and they are indexed in RePEc (http://repec.org/).

Publisher and editor

Oesterreichische Nationalbank
Otto-Wagner-Platz 3, 1090 Vienna, Austria
PO Box 61, 1011 Vienna, Austria
www.oenb.at
oenb.info@oenb.at
Phone (+43-1) 40420-6666

Editor Martin Summer

Cover Design Information Management and Services Division

DVR 0031577

ISSN 2310-533X (Online)