Causal state-dependent local projections

Identification, estimation and inference for policy analysis

Weining Wang and Xiyu Jiao
DG ECFIN short course
Brussels, 23 November 2026
Based on David, Giacomini, Jiao and Wang (8 November 2026)

Course sections

Introduction	Welcome and LP foundations
Section 1	Identification and causal responses
Section 2	Causal state-dependent local projections
Break	
Section 3	Estimation and inference
Section 4	Practical R application
Discussion	Questions and policy applications

Annotated reading list: six core readings and optional references by topic.

Policy questions for this course

Policy application	Empirical question	Course connection
Defence R&D	Does public research crowd in private investment?	Panel LP, scaling, long horizons
Uncertainty	How does investment respond to common and domestic uncertainty?	Shock construction, dependence
Monetary policy	How do labour-supply responses vary by income?	External instruments, heterogeneity
Firm investment	How does financial health shape the response?	Continuous states and splines

Published applications motivate the design. Worked calculations and the R practical develop the tools.

Source: DG ECFIN, QREA 24(3), 2025, Chapters II–III. Also the 2024 uncertainty note and 2026 spillovers report.

A local projection is a regression at each horizon

For each h, estimate a separate regression:

$$Y_{t+h} = \alpha_h + \beta_h X_t + \Gamma_h' W_{t-1} + u_{t+h}.$$

- Shock and information set refer to date t.
- The outcome date changes with h.
- $\{\widehat{\beta}_h\}$ is the response path.

Constructed uncertainty example

h	$\widehat{\beta}_h$	SE
0	-0.10	0.09
1	-0.25	0.12
2	-0.40	0.15
4	-0.20	0.18

If Y is quarterly growth in percentage points, each entry is a growth response per unit shock.

At $h = 2$, a pointwise 95% interval is $-0.40 \pm 1.96(0.15) = [-0.694, -0.106]$.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

Outcome definitions change the estimand

Suppose baseline investment is 100 million per period and the intervention gives $Y_t = 99$, $Y_{t+1} = 98$, $Y_{t+2} = 99$.

Target at $h = 2$	Calculation	Interpretation
Level response	$99 - 100 = -1$	Million at date $t + 2$
Change from $t - 1$	$(99 - 100) - 0 = -1$	Same effect if $t - 1$ is unaffected
Cumulative flow response	$(99 + 98 + 99) - 300 = -4$	Million across three periods

For log levels, $100\beta_h\delta$ approximates a percentage level response. For a share recorded as 0–100, $\beta_h\delta$ is in percentage points.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

LP and VAR estimation share the identification problem

- LP estimates each horizon directly. VAR propagates estimated dynamics
- With unrestricted lags they can recover the same population responses
- Finite-sample bias and precision depend on lag choices and restrictions
- Replacing a VAR with an LP does not make a shock exogenous

Sources: Jordà (2005), Plagborg-Møller and Wolf (2021).

Section 1: Identification and causal responses

Target, assumptions, shock moments and failures of interpretation.

Four common routes to identifying an LP

Design	Economic identifying argument
Conditional exogeneity	Controls remove dependence between treatment and potential outcomes.
Recursive timing	Within-period exclusions isolate a structural innovation.
Observed shock	Narrative or high-frequency variation measures the intended exogenous shock.
External instrument	A relevant proxy is excluded from other determinants of the outcome.

These routes can overlap. The LP estimates the response; the design gives it causal meaning.

Reading: Jordà and Taylor (2025); Stock and Watson (2018).

ECFIN case: public defence R&D

Published design: an annual, unbalanced country panel, 1981–2022, with responses through ten years and a comparison of innovation leaders with moderate innovators.

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \tau_{t,h} + \beta_{h,0} d_{it} + \beta_{h,1} d_{i,t-1} + \gamma_h' W_{i,t-1} + u_{i,t+h}.$$

- d_{it}: public defence R&D as a share of GDP, in percentage points.
- y_{it} here: privately financed business R&D as a share of GDP.
- $\beta_{h,0}$: horizon-h response per unit of current spending, conditional on the specified regressors.

Causal interpretation requires exogeneity of that spending variation after adjustment.

Source: DG ECFIN, QREA 24(3), 2025, Chapter III, pp. 33–36, equation (1), notation simplified.

Defence R&D: the published response paths

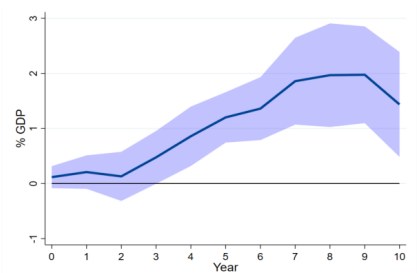

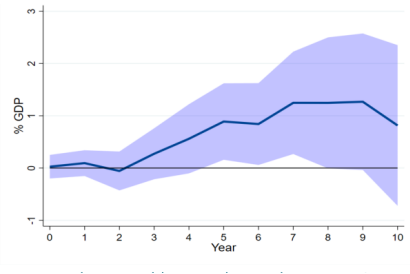

Left: EU sample. Right: OECD sample.
Response to a one-percentage-point-of-GDP spending increase. Shading: 90% intervals.

Interpretation questions
Which horizons are precise? How does the retained sample change? What would a joint claim across years require?

Source: DG ECFIN, QREA 24(3), 2025, p. 36, Graphs III.3–III.4. Cropped, otherwise unchanged. © EU 2025, CC BY 4.0.

Worked calculation: shock size and outcome units

Suppose $\widehat{\beta}_5 = 1.40$ with $\mathrm{se} = 0.50$. Both spending and private R&D are in percentage points of GDP.

$$\widehat{\mathrm{response}}_5 = \delta\widehat{\beta}_5, \qquad \mathrm{se}(\widehat{\mathrm{response}}_5) = |\delta|\,\mathrm{se}(\widehat{\beta}_5).$$

Spending increase	δ	Response	95% interval
1 percentage point of GDP	1.00	1.40	[0.42, 2.38]
0.1 percentage point of GDP	0.10	0.14	[0.042, 0.238]

A 10% rise from a share of 0.05% of GDP is only 0.005 percentage points. Its fitted response is $1.40(0.005) = 0.007$.

Linear rescaling requires the same response relationship over the relevant shock range.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

Defence R&D: an identification exercise

Proposed argument
Security-driven budgets may be less responsive to current domestic activity than ordinary fiscal spending.

Potential omitted variable
A security event may also affect confidence, energy costs and foreign demand.

Participant task
Write the required orthogonality:

$$E[\widetilde{d}_{it}\, u_{i,t+h}] = 0,$$

where $\widetilde{d}$ removes included controls and fixed effects.

Explain why residualisation does not prove this condition.

Checks: announcement timing, pre-shock outcomes, exposure to common events and alternative pre-shock controls. None is a standalone proof of exogeneity.

Source: DG ECFIN, QREA 24(3), 2025, p. 35 motivates the exogeneity discussion. Exercise and diagnostics are course extensions.

Conditional exogeneity: what controls must accomplish

For each treatment value d, require consistency, support and

$$E[Y_{t+h}(d) \mid D_t, W] = E[Y_{t+h}(d) \mid W].$$

With an affine constant response, $E[Y_{t+h}(d) \mid W] = m_h(W) + \theta_h d$. Let $V_t = D_t - E[D_t \mid W]$. Then

$$\theta_h = \frac{E[V_t Y_{t+h}]}{E[V_t^2]}, \qquad E[V_t^2] > 0.$$

- Exogeneity connects observed and potential outcomes.
- Correct adjustment removes the nuisance $m_h(W)$.
- A finite set of linear lag controls need not implement this adjustment.

Recursive timing is an identifying restriction

For reduced-form innovations ordered as outcome, then policy,

$$u_{y,t} = b_{11}\varepsilon_{1t}, \qquad u_{d,t} = b_{21}\varepsilon_{1t} + b_{22}\varepsilon_{2t}.$$

- The target policy shock has no contemporaneous effect on the first variable.
- Residualising $u_{d,t}$ on $u_{y,t}$ isolates $b_{22}\varepsilon_{2t}$ under this model.
- Within-quarter exclusions can be much stronger than within-minute exclusions.
- The resulting shock can enter an LP; the identifying assumption remains the timing restriction.

Reading: Plagborg-Møller and Wolf (2021).

ECFIN case: the cost of uncertainty

Published approach

The 2024 note uses a VAR and Cholesky identification. A one-standard-deviation innovation is about 50 EPU index points.

Its estimated annual GDP-growth drag is 0.45 percentage points in the full-sample analysis.

LP version for the course

Keep the same identified shock $\widehat{\varepsilon}_t$ and a consistent sample:

$$Y_{t+h} = \alpha_h + \beta_h \widehat{\varepsilon}_t + \Gamma_h' W_{t-1} + u_{t+h}.$$

Compare paths and intervals. Differences can reflect lag restrictions and finite-sample estimation.

Causal content still comes from shock identification. The note's alternative uncertainty paths are conditional VARX scenarios.

Source: DG ECFIN, The cost of uncertainty, 15 November 2024.

Uncertainty: an innovation versus a policy scenario

Single innovation

Set one structural shock at date t. Later endogenous variables follow the estimated dynamics.

$$\Delta Y_{t+h} = \psi_h \delta.$$

Specified uncertainty path

Prescribe the index at future dates. Recover innovations consistent with that path and the dynamic model.

$$\Delta Y_{t+h} = \sum_{j=0}^{h} \psi_{h-j} \delta_j.$$

In a linear, invariant system, the second expression aggregates shock responses. An EPU-index path is generally different from a shock sequence.

Worked example later

The practical compares a single shock with repeated shocks using an explicitly supplied response kernel.

Source: DG ECFIN, The cost of uncertainty, 15 November 2024.

Observed shocks: measurement and economic content

- Narrative measures use policy records and the decision-maker's information.
- High-frequency surprises narrow the event window but may mix policy and other news.
- Predictability from pre-announcement data motivates careful information adjustment.
- A noisy proxy used directly need not have the same scale as the structural shock.

If $Y = \theta\varepsilon + u$ and $X = \varepsilon + \nu$ with mutually uncorrelated components,

$$\beta^{OLS} = \theta \frac{\mathrm{Var}(\varepsilon)}{\mathrm{Var}(\varepsilon) + \mathrm{Var}(\nu)}.$$

Examples: Romer and Romer (2004); Gertler and Karadi (2015); Bauer and Swanson (2023).

LP-IV: relevance, exclusion and normalisation

With instrument Q_t, residualised instrument $\widetilde{Q}_t$, and $Y_{t+h} = \theta_h D_t + m_h(W) + u_{t+h}$,

$$E[\widetilde{Q}_t u_{t+h}] = 0, \quad E[\widetilde{Q}_t D_t] \neq 0 \quad \Longrightarrow \quad \theta_h = \frac{E[\widetilde{Q}_t Y_{t+h}]}{E[\widetilde{Q}_t D_t]}.$$

- Exclusion must hold for the horizon-h error, including relevant dated shocks.
- A proxy design often identifies the response per unit impact movement in D_t.
- LP-IV can avoid invertibility under strong instrument conditions; this is not unconditional.
- Weak instruments require appropriate confidence sets; a valid first stage does not prove exclusion.

Reading: Stock and Watson (2018).

ECFIN case: monetary policy and labour supply

Published application: monthly labour-flow measures from EU-SILC and high-frequency monetary-policy instruments in an IV-VAR.

LP-IV teaching extension

D_t is an interest-rate change in percentage points and Q_t an announcement instrument. With matched controls:

$$\theta_h = \frac{E[\widetilde{Q}_t Y_{t+h}]}{E[\widetilde{Q}_t D_t]}.$$

What must be defended

The instrument moves the intended policy variable and is excluded from other determinants of the future outcome.

Announcements can also reveal central-bank information about the economy.

Source: DG ECFIN, QREA 24(3), 2025, Chapter II and Box II.1. LP-IV is a teaching extension, not the published estimator.

LP-IV: a numerical normalisation example

After removing identical controls, suppose the sample moments are

$$\widehat{E}[\widetilde{Q}D] = 0.08, \qquad \widehat{E}[\widetilde{Q}Y_{t+4}] = -0.016.$$

Then

$$\widehat{\theta}_4 = -0.016/0.08 = -0.20, \qquad \text{response to 25 b.p. tightening} = 0.25(-0.20) = -0.05.$$

- If Y is a transition probability in percentage points, the response is -0.05 percentage points.
- Multiplying Q by 100 changes both moments by 100 and leaves the ratio unchanged.
- An imprecise denominator makes conventional normal intervals unreliable. Report first-stage information and use weak-IV-robust inference when needed.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

Literature map: from an LP coefficient to our target

- **Estimation:** Jordà (2005); LP–VAR equivalence in Plagborg-Møller and Wolf (2021).
- **Causal objects:** Rambachan and Shephard (2025 revision).
- **Nonlinear responses:** Kolesár and Plagborg-Møller (2025): causal weighted marginal effects under suitable observed-shock or proxy assumptions.
- **State dependence:** Gonçalves et al. (2024): finite versus infinitesimal responses depend on the state mechanism.
- **Our course paper:** explicit conditions for $\mathrm{IRF}_h(z;\delta) = \delta g_h(z)$ with an unrestricted past-information nuisance.

Next: separate causal exogeneity, innovation orthogonality and an affine response in the shock.

Potential outcomes: the intervention and the population

Let $Y_{i,t+h}(x)$ be the outcome when the date-t shock is set to x.

$$\mathrm{IRF}_h(z;\delta) = E[Y_{i,t+h}(X_t+\delta) - Y_{i,t+h}(X_t) \mid Z_{i,t-1} = z].$$

- Shift the realised innovation by δ and average within the initial-state group.
- Later outcomes and states evolve under the intervention.
- Consistency: the observed outcome equals $Y_{i,t+h}(X_t)$.
- Both X_t and $X_t + \delta$ must lie in the relevant conditional support.

Source: manuscript, Section 3.1. A fixed comparison of x and $x + \delta$ gives the same slope under the affine condition below.

The causal comparison

- Specify outcome, horizon, shock units and initial state
- Change the date-t innovation; let later outcomes evolve
- Condition on the initial state and allow subsequent states to evolve

$$\mathrm{IRF}_h(z;\delta) = \delta g_h(z), \qquad \Delta_h = g_h(z_A) - g_h(z_B)$$

Get the timing right

Example: an ECB announcement and quarterly firm investment

When	Available information	Role
Before announcement	Last published balance sheet and financial-state measure	Initial state and controls
Announcement window	Surprise change in market rates	Candidate shock or instrument
After announcement	New credit conditions and revised confidence	Potential transmission channels
Future quarter $t + h$	Investment outcome	LP dependent variable

A variable dated $t - 1$ but published after the announcement may use unavailable information. A post-announcement credit spread may absorb part of the total effect.

From causal assumptions to a response function

For every admissible x, causal mean exogeneity requires

$$E[Y_{i,t+h}(x) \mid \mathcal{F}_{t-1}, X_t] = E[Y_{i,t+h}(x) \mid \mathcal{F}_{t-1}].$$

Consistency and the observed conditional-mean restriction imply

$$E[Y_{i,t+h}(x) \mid \mathcal{F}_{t-1}] = r_{i,h}(\mathcal{F}_{t-1}) + g_h(Z_{i,t-1})x.$$

Subtract the counterfactual means and average over initial-state groups:

$$\mathrm{IRF}_h(z;\delta) = \delta g_h(z).$$

Roles of the assumptions

Exogeneity connects the observed and counterfactual means. The affine restriction makes the slope independent of shock size and starting value.

The identification conditions

- Causal exogeneity links shocks to potential outcomes
- An unpredictable shock removes the past-information nuisance
- The conditional mean is affine in the shock
- The state captures the relevant response coefficient; support is positive

$$E[X_t \mid \mathcal{F}_{t-1}] = 0, \qquad E[Y_{i,t+h} \mid \mathcal{F}_{t-1}, X_t] = g_h(Z_{i,t-1})X_t + r_{i,h}(\mathcal{F}_{t-1})$$

The shock moment: a derivation

Write $Y = g_h(Z)X + r(\mathcal{F}) + u$, where $E[u \mid \mathcal{F}, X] = 0$, Z is $\mathcal{F}$-measurable and $E[X \mid \mathcal{F}] = 0$.

$$
\begin{aligned}
E[XY \mid Z] &= g_h(Z)E[X^2 \mid Z] + E[Xr(\mathcal{F}) \mid Z] + E[Xu \mid Z], \\
E[Xr(\mathcal{F}) \mid Z] &= E[r(\mathcal{F})E[X \mid \mathcal{F}] \mid Z] = 0, \\
E[Xu \mid Z] &= E[XE[u \mid \mathcal{F}, X] \mid Z] = 0.
\end{aligned}
$$

If $E[X^2 \mid Z = z] > 0$, the ratio of moments identifies $g_h(z)$. Statistical identification alone does not establish the counterfactual interpretation.

Why the past-information nuisance can be left unrestricted

Under A1–A3, for any predetermined basis $\phi_J(Z)$ and controls W,

$$E[X_t \phi_J(Z_{i,t-1}) r_{i,h}(\mathcal{F}_{t-1})] = 0, \qquad E[X_t \phi_J(Z_{i,t-1}) W_{i,t-1}^\top] = 0.$$

- The m.d.s. condition removes arbitrary nonlinear functions of the past from the shock moment.
- Predetermined controls can improve precision without spanning the entire nuisance.
- A shock residual orthogonal to a few selected lags need not satisfy $E[X_t \mid \mathcal{F}_{t-1}] = 0$.
- Sample transformations, including within-unit demeaning, require a separate justification.

Source: manuscript, Section 3.1, especially footnote 5.

Shock construction is an identification decision

- A policy change need not be an exogenous policy shock
- Assess anticipation, information effects and omitted news
- Residualizing a series does not establish causal exogeneity
- An instrument requires an IV design and IV inference

Serial correlation creates a nuisance term

Suppose $X_t = \rho_x X_{t-1} + \eta_t$ and

$$Y_{t+h} = g_h X_t + \gamma_h X_{t-1} + u_{t+h}, \qquad E[X_t u_{t+h}] = 0.$$

For stationary, mean-zero X_t, a regression omitting X_{t-1} has slope

$$\frac{E[X_t Y_{t+h}]}{E[X_t^2]} = g_h + \gamma_h \rho_x.$$

- Serial correlation creates bias here through the omitted lag.
- Correct lag controls or correctly specified residualisation can remove this statistical term.
- Neither operation proves that η_t is a causally exogenous innovation.

A nonlinear shock response changes the estimand

Consider $Y = g(z)X + c(z)X^3 + u$, with $X \sim N(0,1)$ independent of the state and causally exogenous.

$$\beta^{LP}(z) = \frac{E[XY \mid Z = z]}{E[X^2 \mid Z = z]} = g(z) + 3c(z).$$

The paper's finite-shift target is

$$E[Y(X + \delta) - Y(X) \mid Z = z] = \delta\{g(z) + 3c(z)\} + c(z)\delta^3.$$

- Multiplying the LP slope by δ misses $c(z)\delta^3$.
- For a fixed comparison from 0 to δ, the response is $g(z)\delta + c(z)\delta^3$.
- Flexible dependence on z does not remove nonlinearity in X.

Controls and fixed effects need a purpose

Choice	What it absorbs	Implication
Country fixed effects	Persistent country differences	Within-country variation identifies slopes
Year fixed effects	Common movements across countries	Country-specific spending remains usable
Date effects with common X_t	The common shock regressor	Absolute response level is unidentified
Post-shock credit conditions	A possible transmission channel	Changes the total-effect target

Dynamic fixed-effect estimators can have finite-T bias. Choosing fixed effects and choosing a dependence-robust covariance are separate decisions.

Discuss your identification design

- Name your outcome, shock, initial state and horizon
- Identify one plausible threat to causal interpretation
- Decide whether the current aim is descriptive or causal

Section 2: Causal state-dependent local projections

Propagation, endogenous initial states and functional-form restrictions.

Worked propagation example

Assume $Y_{it} = \rho Y_{i,t-1} + a(Z_{i,t-1})X_t + \varepsilon_{it}$, and future innovation terms have zero conditional mean under the intervention.

$$E[Y_{it}(x+\delta) - Y_{it}(x) \mid \mathcal{F}_{t-1}] = a(Z_{i,t-1})\delta,$$
$$E[Y_{i,t+1}(x+\delta) - Y_{i,t+1}(x) \mid \mathcal{F}_{t-1}] = \rho a(Z_{i,t-1})\delta,$$
$$g_h(z) = \rho^h a(z).$$

With $\rho = .6$ and $a(z) = 1 - z^2$, $g_0(0) = 1$, $g_0(.75) = .4375$, and their impact contrast is .5625.

Scope

Stable linear propagation and the future-innovation restrictions are substantive assumptions.

An endogenous state can still be admissible

- Initial conditions may depend on earlier shocks
- Stable propagation preserves linearity in the current shock
- Changing propagation or shock nonlinearities need more care

$$Y_{it} = a(Z_{i,t-1})X_t + \rho Y_{i,t-1} + \varepsilon_{it} \quad \Longrightarrow \quad g_h(z) = \rho^h a(z)$$

When state transitions alter propagation

If $Y_{t+1} = \rho(Z_t)Y_t + \varepsilon_{t+1}$ and the date-t intervention changes Z_t, then

$$Y_{t+1}(x+\delta) - Y_{t+1}(x) = \rho(Z_t(x+\delta))Y_t(x+\delta) - \rho(Z_t(x))Y_t(x).$$

- A single constant propagation coefficient generally cannot factor out.
- The response can depend on the starting shock value and the shock size.
- Conditioning on realised Z_t can condition on a consequence of treatment.
- Initial-state heterogeneity and interventions on the state are distinct causal questions.

Three DGPs and the scope of the method

Dynamic structure	What the flexible LP can recover
Linear shock exposure, fixed propagation	$g_h(z)$ under A1–A3 and support conditions
Shock-dependent propagation	Generally a projection; the affine representation can fail
Nonlinear shock exposure	Generally a projection; the finite-shift response depends on X_t and δ

A

fixed state alone does not establish A3. First-order HA/HANK solutions can satisfy A3 within their linear approximation. A nonlinear economy requires checking the approximation over the relevant shock range.

Source: manuscript, Sections 2 and 3.2. Scalar average-effect interpretations have their own assumptions and do not recover the full state function.

Coarsening a state changes the weights

Within a reported state group, two equally common latent types have slopes 1 and 3, and shock variances 1 and 4. Assume conditional shock means are zero.

$$\text{Average causal slope} = \tfrac{1}{2}(1+3) = 2, \qquad \text{LP slope} = \frac{\frac{1}{2}(1)(1) + \frac{1}{2}(4)(3)}{\frac{1}{2}(1) + \frac{1}{2}(4)} = 2.6.$$

More generally, with latent slope a and variance v,

$$\beta^{LP}(z) - E[a \mid z] = \frac{\mathrm{Cov}(v, a \mid z)}{E[v \mid z]}.$$

A coarse recession/expansion classification can therefore change the target, even when every latent-type response is causal.

The linear interaction is a benchmark

- Estimate the common shock and its interaction with the state
- A small interaction does not establish homogeneous effects
- Compare the fitted function with a flexible alternative

$$Y_{i,t+h} = \alpha_h X_t + \beta_h X_t Z_{i,t-1} + W_{i,t-1}^{\top} \gamma_h + u_{i,t+h}$$

A linear interaction averages slopes

Under the assumptions of Proposition 2,

$$\beta_h^{\text{linearSD}} = \int_{\mathcal{Z}} \omega_h(z) g_h'(z)\, dz, \qquad \omega_h(z) \geq 0, \quad \int_{\mathcal{Z}} \omega_h(z)\, dz = 1.$$

- This coefficient summarizes local slopes of the response curve.
- A positive coefficient does not establish that the curve rises everywhere.
- Opposing slopes can cancel even when state contrasts are large.
- Recovering a response function requires estimating its shape.

Source: manuscript, Proposition 2 and Appendix A. Conditional homoskedasticity gives $\beta_h = \text{Cov}(Z, g_h(Z))/\text{Var}(Z)$.

Why a zero linear interaction can hide heterogeneity

Let $Z \sim U[-1,1]$, with an independent mean-zero shock of constant variance, and $g(z) = 1 - z^2$. Project $g(Z)$ onto $(1, Z)$:

$$\beta = \frac{\mathrm{Cov}(Z, 1 - Z^2)}{\mathrm{Var}(Z)} = 0, \qquad \alpha = E[1 - Z^2] = \frac{2}{3}.$$

- Symmetry cancels the positive and negative slopes.
- The response is 1 at the centre and 0 at the boundaries.
- Testing one linear interaction does not test general homogeneity.
- The result uses these specific distribution and weighting assumptions.

Why the linear interaction can miss heterogeneity

- A symmetric hump can have an almost zero linear slope
- A zero interaction coefficient need not imply equal responses
- This figure shows a population illustration, not ECFIN estimates

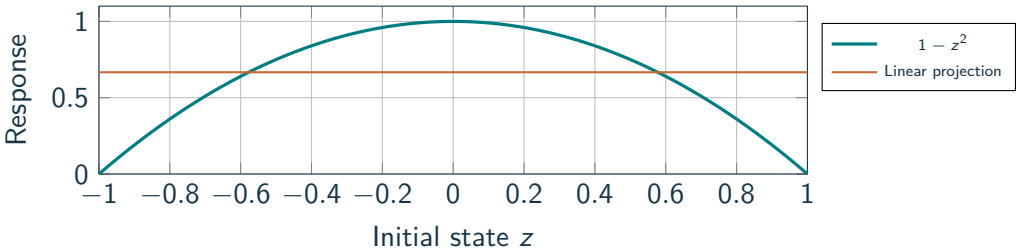

Common shocks and response normalization

With unrestricted date effects,

$$Y_{i,t+h} = \alpha_{t,h} + X_t g_h(Z_{i,t-1}) + \cdots .$$

For any constant c, the same fitted values follow from

$$g_h^*(z) = g_h(z) + c, \qquad \alpha_{t,h}^* = \alpha_{t,h} - cX_t.$$

- The common level needs a normalization or additional identifying structure.
- A contrast $g_h(z_A) - g_h(z_B)$ is invariant to that level shift.
- A constant is already in the span of the full spline basis. Adding a separate X_t column causes collinearity.

Defence R&D: group heterogeneity and initial states

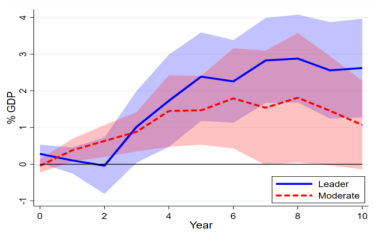

Published leader/moderate-innovator comparison.
Shading: 90% intervals.

Different causal questions
A fixed group describes heterogeneity across country types. A lagged state describes conditions before a particular shock.

A 2024 classification is not automatically a predetermined state for earlier observations.

Course extension: construct a pre-shock innovation measure $Z_{i,t-1}$ and estimate $g_h(Z_{i,t-1})$. This changes the estimand and requires new assumptions.

Source: DG ECFIN, QREA 24(3), 2025, pp. 33, 36, Graph III.5. Cropped, otherwise unchanged. © EU 2025, CC BY 4.0.

A continuous innovation-state specification

A proposed extension of the defence application is

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + \tau_{t,h} + d_{it} g_h(Z_{i,t-1}) + \gamma_h' W_{i,t-1} + u_{i,t+h}.$$

- Define Z using information available before the spending decision.
- Begin with a linear interaction, then compare a small spline basis.
- Check residual spending variation at low and high Z.
- Report a preselected contrast, such as the 75th minus 25th percentile.

Spending levels are not automatically martingale-difference shocks. The course paper's unrestricted-nuisance argument needs an additional justification here.

Proposed teaching extension. No new empirical estimates are asserted.

Application: monetary policy and firm investment

$$k_{i,t+h} - k_{i,t-1} = \alpha_{sth} + g_h(Z_{i,t-1})X_t + \rho_h \Delta k_{i,t-1} + W_{i,t-1}^{\top}\gamma_h + u_{i,t+h}.$$

- Quarterly Compustat capital stock and high-frequency monetary surprises.
- Initial state: standardized distance to default (D2D). Higher values mean safer firms.
- Figures show cumulative investment responses to a 25 b.p. surprise rate cut.
- Sector-time-horizon effects absorb the common response. Interpret the plotted curves under the paper's level normalization.

Source: manuscript, Section 6, equation (27). This is separate from the synthetic R exercise.

Firm responses across the financial-state distribution

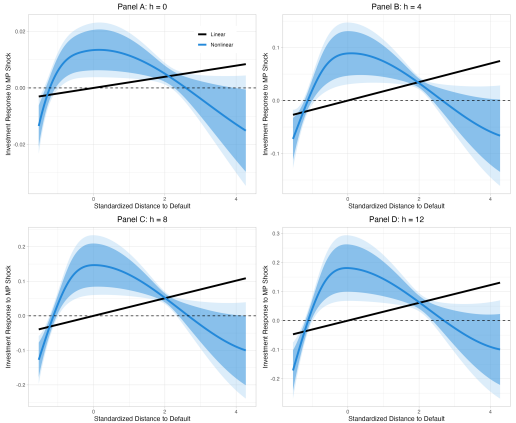

Source: manuscript, Figure 3. Differential cumulative responses to a 25 b.p. cut at horizons 0, 4, 8 and 12 quarters. Dark/light shading: 68%/90% state-uniform bands, separately at each horizon.

Reading the empirical response curve

- The flexible estimate peaks around mean D2D and falls toward both tails.
- Responsiveness increases as firms become safer in the left tail, then decreases to the right of the peak.
- A linear interaction imposes a single slope across both regions.
- Relative to the earlier empirical specification, controls and fixed-effect treatment also differ. The full change is not attributable solely to the spline.

Discussion: which pair of states would you compare, and what covariance would you need?

Source: manuscript, Section 6.1. A response pattern does not uniquely identify its economic mechanism.

Section 2 discussion

- Which variable is the initial state in your project?
- Could it be affected by the current intervention?
- Would coarsening the state alter the weighting?
- Which restriction makes the response affine in the shock?

ECFIN 2026 uncertainty: Germany and France

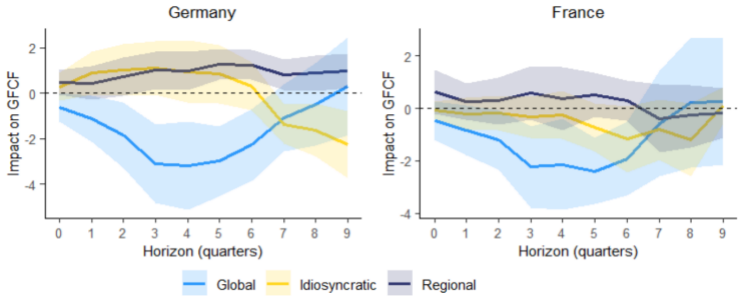

Discussion: compare the timing of investment-growth responses to global and domestic EPU components. Testing a difference requires their joint covariance.

Source: DG ECFIN, Institutional Paper 346, September 2026, p. 33, Graph 38, upper panels. Cropped, otherwise unchanged. © EU 2026, CC BY 4.0.

ECFIN 2026 uncertainty: Spain and Belgium

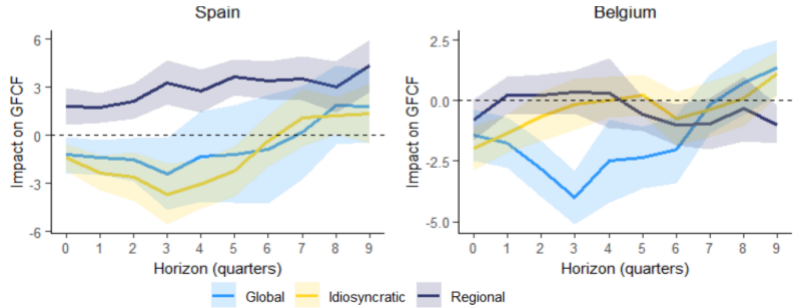

Discussion: do the signs support one common narrative? The published LP treats the shock as exogenous. Component decomposition alone does not establish that assumption.

Source: DG ECFIN, Institutional Paper 346, September 2026, p. 33, Graph 38, lower panels. Cropped, otherwise unchanged. © EU 2026, CC BY 4.0.

Common-component removal and shock identification

For survey uncertainty, the report uses a GDP-weighted leave-one-country-out aggregate. In simplified notation:

$$C_{-i,t} = \sum_{j \neq i} w_{ij} U_{jt}, \qquad U_{it} = a_i + b_i C_{-i,t} + v_{it}, \qquad \sum_{j \neq i} w_{ij} = 1.$$

What the projection supplies
Residual orthogonality to included regressors, and fitted/residual components.

What still needs justification
Unpredictability, causal exogeneity, and residual regional factors that affect investment.

These are generated regressors. A resampling scheme intended to capture their estimation error must repeat the decomposition.

Source: DG ECFIN, Institutional Paper 346, September 2026, pp. 29–30, footnotes 22–23. Causal and inference discussion: course extension.

Initial-state means and quantile responses

Object	Conditioning or target	Interpretation
State-dependent mean	$E[Y_{t+h}(x) \mid Z_{t-1} = z]$	Mean effect for an initial state
Group-specific mean	$E[Y_{t+h}(x) \mid G_i = g]$	Mean effect for a country or income group
Quantile response	$Q_\tau(Y_{t+h}(x) \mid W)$	Change in a distributional quantile

A lower-tail outcome response is not automatically the mean response when the economy starts in recession. Quantile causal interpretation needs its own assumptions.

Source: DG ECFIN, Institutional Paper 346, September 2026, p. 32 motivates the comparison. Distinctions and notation: teaching explanations.

Break

Next section: Estimation and inference.

Section 3: Estimation and inference

Weining Wang and Xiyu Jiao

A flexible response with cubic B-splines

- Interact every basis column with the shock
- Keep the same basis when predicting at new state values
- Choose J from 4 to 20 using an AIC criterion

$$\widehat{g}_h(z) = \phi_J(z)^\top \widehat{b}_h, \qquad D_{it} = (X_t \phi_J(Z_{i,t-1})^\top, W_{i,t-1}^\top)^\top$$

```
B <- splines::bs(z, degree=3, knots=knots,
                 intercept=TRUE)
D <- cbind(B*x, W)
```

Spline estimation as a partially residualised regression

For a fixed horizon, stack $Q_{it} = X_t \phi_J(Z_{i,t-1})$ and controls W_{it} into matrices Q, W.

$$M_W = I - W(W'W)^{-1}W', \qquad \widehat{b}_h = (Q' M_W Q)^{-1} Q' M_W Y_h.$$

$$\widehat{g}_h(z) = \phi_J(z)' \widehat{b}_h.$$

- Full column rank is necessary. Use QR methods numerically.
- Under A1–A3, m.d.s. orthogonality removes the past-information nuisance. Controls may improve precision. If these conditions fail, adding selected controls is not an automatic repair.
- Retain the training knots and boundary knots for prediction.

Basis dimension: fit, variance and bias

The teaching implementation compares $J = 4, \ldots, 20$ using

$$\mathrm{AIC}(J) = n_h \log(\mathrm{RSS}_J / n_h) + 2p_J,$$

where p_J counts all fitted coefficients.

- A larger basis can reduce approximation error and increase variance.
- This criterion is a teaching selection rule, not a dependence-adjusted optimality theorem.
- For inference about the true function, approximation bias must be negligible or explicitly corrected.
- Re-estimate with a few fixed dimensions and report sensitivity.

The dependence structure determines standard errors

- Time series: assess dependence of regression scores across dates
- Country panels: assess within-country and cross-country dependence
- Common shocks: many firms do not create many independent shock dates
- Lag-augmented time-series LPs can permit heteroskedasticity-robust inference under their assumptions

Source: Montiel Olea and Plagborg-Møller (2021). Their result does not automatically extend to this flexible panel estimator.

Long horizons change the usable sample

Consider a balanced annual panel from 1981 through 2022: 42 dates. Use $y_{i,t+h} - y_{i,t-1}$ with one lag of regressors.

Horizon h	First t	Last t	Maximum dates per country
0	1982	2022	41
5	1982	2017	36
10	1982	2012	31

- Unbalanced panels and missing controls reduce these counts further.
- Different horizons can describe different country-date compositions.
- A common sample removes one source of change, at a cost in precision.

Constructed sample-count exercise using the published date range, not the study's estimation counts.

A common shock limits panel information

Suppose the score has a common component:

$$q_{it} = a_t + e_{it}, \qquad \mathrm{Var}(a_t) = 1, \quad \mathrm{Var}(e_{it}) = 4,$$

with a_t and e_{it} uncorrelated. For independent e_{it} across $N = 100$ units:

$$\mathrm{Var}\left(\frac{1}{N}\sum_i q_{it}\right) = 1 + \frac{4}{100} = 1.04.$$

Treating units as independent would give $(1 + 4)/100 = 0.05$.

- More countries remove idiosyncratic noise but leave the common component.
- Serial dependence in a_t also matters across dates.
- Sum scores within date, then allow temporal dependence in those sums.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

Why precision can change across the state space

The paper lets the effective-score dependence scale grow with N:

$$\chi_{h,J,N} = 1 \vee \lambda_{\max}(\Omega_{h,J,N}), \qquad \|\widehat{b}_h - b_h\|_2 = O_p(J^{1/2-\kappa}) + O_p\left(\sqrt{\frac{J\chi_{h,J,N}}{NT}}\right).$$

- Weak score dependence: stochastic coefficient error of order $\sqrt{J/(NT)}$.
- Strong aggregate components, $\chi \asymp N$: order $\sqrt{J/T}$.
- Aggregate components can load differently at different states. One common rate need not describe every point on the curve.
- Use the estimated state-specific standard error when studentizing.

Source: manuscript, Theorem 1. Joint growth, moment, approximation and dependence restrictions apply. This is a coefficient rate, not a universal pointwise rate.

Overlapping outcomes induce serial dependence

For $Y_t = \rho Y_{t-1} + \epsilon_t$, the h-step forecast error contains

$$e_{t,h} = \sum_{j=1}^{h} \rho^{h-j} \epsilon_{t+j}.$$

Adjacent forecast errors share future innovations. In a panel, common components can also correlate observations within a date.

- Treating all unit-date rows as independent can understate uncertainty.
- Clustering only within dates does not allow serial dependence between dates.
- Date-aggregated HAC addresses both in the teaching design, under suitable time-series regularity.

HAC covariance: the full construction

With full regressor vector D_{it} and residual $\widehat{u}_{i,t+h}$, form date sums $s_t = \sum_i D_{it}\widehat{u}_{i,t+h}$.

$$\widehat{\Omega} = \sum_t s_t s_t' + \sum_{\ell=1}^{L}\left(1 - \frac{\ell}{L+1}\right)\sum_{t=\ell+1}^{T_h}(s_t s_{t-\ell}' + s_{t-\ell}s_t').$$

$$\widehat{V} = (D'D)^{-1}\widehat{\Omega}(D'D)^{-1}.$$

- These are sums throughout; no additional division by n_h belongs in this formula.
- Extract the spline coefficient block for function inference.
- Order scores by calendar date and assess bandwidth sensitivity.

The inference question determines the uncertainty measure

Policy claim	Required uncertainty calculation
Investment falls four quarters after the shock	Pointwise interval at the prespecified horizon
High-debt countries respond more strongly	Interval for the high-minus-low contrast
The response is negative throughout this range	Simultaneous band over that state range
At least one of 20 country responses differs	Joint test or multiplicity adjustment

Searching across horizons, countries or cut-offs and then reporting an unadjusted interval changes its coverage interpretation.

A contrast needs the covariance between fitted responses

For $\Delta = g(z_A) - g(z_B)$, set $a = \phi(z_A) - \phi(z_B)$.

$$\widehat{\Delta} = a'\widehat{b}, \quad \widehat{\mathrm{Var}}(\widehat{\Delta}) = a'\widehat{V}_b a.$$

Equivalently,

$$\mathrm{Var}(\widehat{g}_A - \widehat{g}_B) = V_A + V_B - 2C_{AB}.$$

If both standard errors are .10 and the covariance is .008, the contrast standard error is $\sqrt{.02 - .016} \approx .063$.

Interpretation

Overlapping individual confidence intervals do not determine whether a contrast is significant.

Worked contrast: significance of a difference

At a prespecified horizon, suppose $\widehat{g}_A = -0.50$ and $\widehat{g}_B = -0.20$, each with standard error 0.20.

$$\widehat{\Delta} = -0.30, \qquad \mathrm{se}(\widehat{\Delta}) = \sqrt{0.20^2 + 0.20^2 - 2C_{AB}}.$$

Assumed covariance	Contrast SE	95% interval
$C_{AB} = 0$	0.283	$[-0.854, 0.254]$
$C_{AB} = 0.030$	0.141	$[-0.577, -0.023]$

The separate response estimates and their individual intervals are identical. Only the joint covariance changes inference about the difference.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

Empirical contrasts answer a specific question

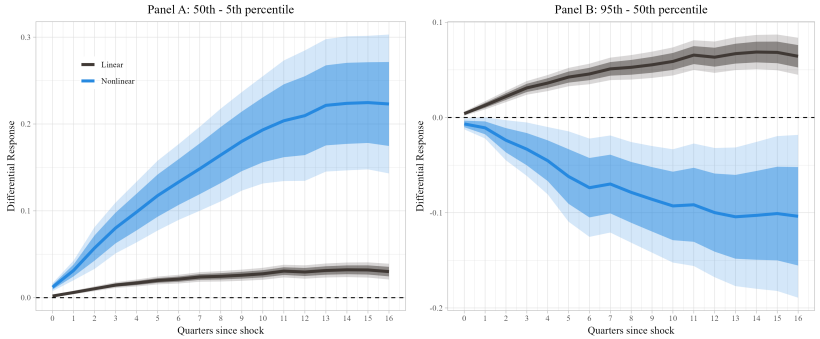

Left: median minus 5th-percentile D2D. Right: 95th percentile minus median.

Source: manuscript, Figure 4. Responses to a 25 b.p. cut. Shading gives 68% and 90% pointwise intervals for each prespecified state contrast at each horizon, not a simultaneous band across horizons.

Pointwise intervals and simultaneous bands

- Pointwise: one preselected state at one horizon
- Simultaneous band: the whole specified state grid at one horizon
- Studentize Gaussian draws by each state's standard error
- A state band is not a band over every horizon

$$c_{.95} = q_{.95}\left[\max_{z \in \mathcal{Z}_G} \frac{|\phi(z)^\top \widehat{V}_b^{1/2} \xi|}{\widehat{\sigma}(z)}\right], \qquad \widehat{g}(z) \pm c_{.95}\widehat{\sigma}(z)$$

```
curve <- infer_curve(impact,
  grid=seq(-.8,.8,length.out=161))
```

Gaussian simulation of a simultaneous state band

At one fixed horizon and a prespecified grid $\mathcal{Z}_G$:

1. Compute $\widehat{\sigma}(z)^2 = \phi(z)' \widehat{V}_b \phi(z)$.
2. Draw $\xi^{(b)} \sim N(0, I_J)$ and set $G^{(b)}(z) = \phi(z)' \widehat{V}_b^{1/2} \xi^{(b)}$.
3. Record $M^{(b)} = \max_{z \in \mathcal{Z}_G} |G^{(b)}(z)| / \widehat{\sigma}(z)$.
4. Use the .95 quantile of these maxima as $c_{.95}$.

$$\widehat{g}(z) \pm c_{.95} \widehat{\sigma}(z).$$

A finite grid approximates a continuum. Coverage requires covariance consistency, distributional approximation and adequate bias control.

A confidence band also needs bias control

- Spline approximation error must be small relative to sampling uncertainty
- AIC selects a fit. It does not by itself guarantee confidence-band coverage
- Inspect sensitivity to basis size, support and the state range
- The practical illustrates asymptotic inference, not a finite-sample coverage guarantee

Simulation evidence: narrow bands can still undercover

Paper simulation, $N = 500$, $T = 200$, $h = 4$, AIC selection, nominal 95% state-uniform bands:

Specification	Coverage	Average width
Without intermediate shock terms	0.89	1.30
With intermediate shock terms	0.83	0.05

- Extra shock interactions can remove common error variation in this DGP.
- The large precision gain does not deliver nominal coverage here.
- The paper retains the unaugmented baseline because additional terms increase dimension and can be unstable.

Source: manuscript, Table 3. These are Monte Carlo results, separate from the course simulation.

Diagnostics before interpretation

Diagnostic	What to inspect	Possible response
Shock support	Variation within each state	Trim unsupported regions
Influential episodes	Leave-one-episode-out estimates	Report crisis dependence
Long horizons	Dates and countries retained	Compare a common sample
Specification	Lags, basis size and controls	Show a sensitivity table
Uncertainty	Score dependence and bandwidth	Recompute intervals consistently

Stability across specifications supports robustness within those designs. It does not establish an untested exclusion restriction.

Section 4: Practical R application

Work through the participant handout.

Practical R application: exercise steps

Exercise step	Participant task
Data checks	Inspect timing, keys and state support
Estimation	Estimate the impact response and plot it
State contrasts	Contrast two states and compare uncertainty
Interpretation	Sensitivity and a short interpretation

Use the participant handout and `exercises/participant_template.R`. The setup file generates synthetic data and defines the functions without running the full horizon loop. Optional extensions follow after the core tasks.

The simulated panel

- 100 units observed over 240 shock dates
- A common shock; a pre-shock financial-state index
- Known response: a hump in the state, with stable decay
- Synthetic units: no empirical policy calibration

$$g_h(z) = 0.6^h(1 - z^2)$$

Setup and files

```
source("code/R/exercise_setup.R")
impact <- estimate_horizon(0, Jset=4:10)
```

Run from the extracted course folder. All data in this exercise are simulated.

R lab: check timing and the sample

- Check unique unit-date keys and missing periods
- Verify that z is pre-shock and x is common within date
- Inspect state support and the number of shock dates

```
stopifnot(!anyDuplicated(d[c("id","t")]))
quantile(d$z, c(.05,.5,.95))
length(unique(d$t))
```

R lab: estimate the response function

- Run the simulated example and fit the impact response
- Inspect the AIC table and selected basis dimension
- Plot the response only over the chosen data-supported range

```
impact <- estimate_horizon(0, Jset=4:10)
impact$AIC
impact$J
infer_contrast(impact)
```

R exercise: the linear benchmark

- Compare the impact response under a linear interaction and a spline
- The true impact function is $1 - z^2$
- Use the benchmark to diagnose functional-form restrictions

```
m <- lm(y ~ x + I(x*z) + ylag, data=d)
zgrid <- seq(-.8,.8,length.out=161)
glin <- coef(m)["x"] + coef(m)["I(x * z)"]*zgrid
lines(zgrid, glin, col="orange")
```

R lab: inference for a state contrast

- Choose the states before looking for a large difference
- Use the covariance between the two fitted responses
- Report the estimate, interval, horizon and shock units

$$a = \phi(z_A) - \phi(z_B), \qquad \widehat{\Delta} = a^\top \widehat{b}, \qquad \mathrm{se}(\widehat{\Delta}) = \sqrt{a^\top \widehat{V}_b a}$$

```
infer_contrast(impact, zA=0, zB=.75)
```

Worked policy path: single and repeated shocks

Let the effect of one unit shock on the *level* of Y at lag h be

$$(\psi_0, \psi_1, \psi_2, \psi_3, \psi_4) = (0.2,\ 0.5,\ 0.8,\ 0.6,\ 0.3).$$

Date/horizon	0	1	2	3	4
Single: $\delta = (.1, 0, 0, 0, 0)$	.02	.05	.08	.06	.03
Repeated: $\delta = (.1, .1, .1, 0, 0)$	.02	.07	.15	.19	.17

At $h = 3$, the repeated-shock response is $.1\psi_3 + .1\psi_2 + .1\psi_1 = .06 + .08 + .05 = .19$.

Assumptions: linear superposition, unchanged propagation and a specified shock sequence. A maintained spending *level* requires mapping that level path to innovations.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

R calculation: response to a shock sequence

```
psi <- c(.2, .5, .8, .6, .3)
delta <- c(.1, .1, .1, 0, 0)
response <- vapply(0:4, function(h) {
  j <- 0:h
  sum(psi[h-j+1] * delta[j+1])
}, numeric(1))
round(response, 2)
# 0.02 0.07 0.15 0.19 0.17
```

- ψ_h is a level-response kernel. Do not cumulate it a second time.
- A growth-response kernel requires conversion to the level target.
- Kernel estimation uncertainty needs cross-horizon covariance.

Standalone exercise: `code/ECFIN_worked_examples.R`. Constructed inputs.

Exercise: translating an ECFIN study into an LP

Choose the defence or uncertainty application. Complete:

Design element	Required decision
Outcome	Level, growth, cumulative change, or cumulative flow?
Treatment	Spending level, index movement, innovation, or instrument?
State	Which measure is observed before the shock?
Adjustment	Which pre-shock controls and fixed effects?
Inference	Which dependence, contrast and horizon range?
Policy statement	What units, shock size and causal qualification?

Deliver one equation and a short interpretation. Identify the most vulnerable assumption and the first diagnostic you would run.

An interpretation with the right scope

Reporting template

At horizon ____, the estimated response per unit shock at state ____ is _____. The interval is _____ and covers a _____ claim.

- State the shock normalization and outcome units.
- Report the state contrast directly, with its own uncertainty.
- In this practical, describe a synthetic demonstration.
- For real data, add the identification argument and the main threat to it.

Optional R extension: an endogenous initial state

Keep the same outcome equation, but let

$$Z_{it} = \tanh(0.75Z_{i,t-1} + 0.25X_t + 0.3v_{it}).$$

- The date-t shock changes the next state. The regressor still uses $Z_{i,t-1}$.
- Independent future shocks and fixed outcome propagation preserve $g_h(z) = 0.6^h(1 - z^2)$.
- This exercise tests the distinction between endogenous evolution and post-shock conditioning.

```
endog <- simulate_panel(state="endogenous")
efit <- estimate_horizon(0, Jset=4:10, data=endog)
infer_contrast(efit)
```

Discussion and Q&A

Applications to DG ECFIN policy analysis and participants' research questions.

Before using an LP result for policy

- Can I defend the shock and initial-state timing?
- Does the response-function interpretation fit this setting?
- Does inference match dependence and the scope of the claim?
- Have I reported support, units and meaningful sensitivity?

Open questions and research clinic

- Describe your data, outcome, shock and initial state
- What would make the result useful for your project?
- Identify the main identification or inference difficulty
- Agree one concrete next analysis or diagnostic

Optional extensions

Additional derivations and diagnostics for follow-up reading.

Worked answer: an uncertainty-state extension

Let X_t be a common uncertainty innovation and $Z_{i,t-1}$ pre-shock leverage:

$$Y_{i,t+h} = \alpha_{i,h} + \tau_{t,h} + X_t g_h(Z_{i,t-1}) + \Gamma_h' W_{i,t-1} + u_{i,t+h}.$$

- With unrestricted date effects, report the differential response $\delta\{g_h(z_H) - g_h(z_L)\}$ unless additional structure identifies the level.
- Defend causal exogeneity and the affine response restriction.
- Check variation across the supported state range and the rank of the interaction design.
- Sum scores by date and allow serial dependence under an appropriate asymptotic framework.

Proposed course design. This is not a re-estimation of the ECFIN studies.

Worked answer: a defensible policy statement

Constructed example

At five years, a 0.1-percentage-point-of-GDP rise in public defence R&D is associated with a 0.14-percentage-point rise in the private-R&D share. The pointwise 95% interval is $[0.042, 0.238]$.

- This compares the outcome at the stated horizon.
- It does not add private R&D spending over all five years.
- Causal language requires exogeneity of spending variation after adjustment.
- A sustained programme needs a mapping from policy to the shock sequence.

Teaching illustration. Numerical values are constructed, not DG ECFIN estimates.

Replication decisions for the published applications

- Match definitions, units, country coverage and data vintages before modifying the specification.
- The 2024 uncertainty webpage gives inconsistent sample endpoints in its text, chart note and footnote 7. Resolve the exact sample using replication files or the authors.
- For defence R&D, distinguish a fixed later classification from an innovation state known before each spending decision.
- For decomposed uncertainty, retain component construction, normalisation and lag selection in full-procedure resampling.

The numerical exercises reproduce teaching calculations. They do not claim to replicate published estimates.

Source: DG ECFIN, The cost of uncertainty, 15 November 2024.

Horizon construction must respect the calendar

- A lead of four means four periods, even when observations are missing
- Match unit and date keys instead of shifting rows blindly
- Report observations and shock dates separately at each horizon
- A common sample helps assess composition sensitivity

```
key <- paste(d$id, d$t, sep=":")
h <- 4
yh <- d$y[match(paste(d$id,d$t+h,sep=":"), key)]
sum(!is.na(yh))
```

Responses across horizons

- Rebuild the outcome and estimation sample at each horizon
- Distinguish levels, changes and cumulative outcomes
- Pointwise horizon intervals do not cover the whole path jointly

```
fits <- lapply(0:8, estimate_horizon)
```

Replicate first, then assess sensitivity

- Match the published target, sample, transformations and units
- Match controls, horizons, tuning and covariance conventions
- Reproduce a reported number before changing the design
- Then test sensitivity to choices that matter for the claim

```
estimate_horizon(4, Jset=7)
estimate_horizon(4, bandwidth=8)
```

Appendix: weighted and unweighted state targets

- The ordinary subpopulation response averages causal slopes
- The LP additionally weights by conditional shock variance
- Their difference is a conditional covariance divided by mean variance
- Constant conditional variance is sufficient for equality

$$\frac{\mathrm{IRF}_h(z;\delta)}{\delta} = E[\widetilde{g}_h(s) \mid Z = z],$$

$$g_h^{\mathrm{LP}}(z) - E[\widetilde{g}_h(s) \mid Z = z] = \frac{\mathrm{Cov}(v(s,z), \widetilde{g}_h(s) \mid Z = z)}{E[v(s,z) \mid Z = z]}, \quad v(s,z) = E[X^2 \mid s, Z = z].$$

Appendix: finite shifts with nonlinear shock exposure

For $f_z(x) = g(z)x + c(z)x^3$ and $X \sim N(0,1)$,

$$E[f_z(X+\delta) - f_z(X)] = g(z)\delta + c(z)\{3\delta E[X^2] + 3\delta^2 E[X] + \delta^3\}.$$

Thus the average derivative equals $g(z) + 3c(z)$, while the average finite shift also contains $c(z)\delta^3$.

```
nonlinear_shift_demo(z=0, delta=1, cubic=.25)
```

The output compares the LP slope times the shift with the finite-shift target. At these values the exact discrepancy is 0.25.

Appendix: FWL and the limits of linear controls

Suppose $Y = \theta D + m(W) + u$, $E[u \mid D, W] = 0$, and $V^L = D - P_W D$, where P_W is the linear projection including an intercept.

$$\beta^{OLS} = \theta + \frac{E[V^L m(W)]}{E[(V^L)^2]}.$$

- Correct conditional exogeneity does not guarantee a correct linear adjustment.
- The extra term vanishes if $m(W)$ is in the control span, or if $P_W D = E[D \mid W]$.
- Exact two-step FWL equivalence requires the same OLS sample and controls.
- Changing samples, controls or estimation methods can break that equivalence.

Further reading: Lloyd and Manuel (2024).

Appendix: why a nonlinear LP can still be causal

For a causally exogenous continuous shock X and smooth response $\Psi_h(x) = E[Y_{t+h}(x)]$,

$$\beta_h^{LP} = \int \omega(x)\Psi_h'(x)\,dx, \qquad \omega(x) = \frac{E[(X - EX)1\{X \geq x\}]}{\mathrm{Var}(X)}.$$

Under interval support and suitable integrability, $\omega(x) \geq 0$ and $\int \omega(x)dx = 1$.

- The coefficient averages marginal causal effects with specific weights.
- With Gaussian shocks and regularity, it equals the average derivative.
- It need not equal a derivative at a chosen point or a finite-shift effect.

Observed-shock illustration of Kolesár and Plagborg-Møller (2025). Proxy and conditional-control designs require additional assumptions.

DG ECFIN sources for the worked examples

- Calò, S., R. Kramer and F. Tomasone (2025). *The impact of public defence R&D spending on private R&D spending*. QREA 24(3), Chapter III, pp. 31–40. Official PDF.
- Bloise, F., C. Cantore, M. Licchetta and G. Ragusa (2025). *Monetary policy and labour supply in the euro area*. QREA 24(3), Chapter II, pp. 19–30. Same official PDF.
- European Commission (2024). *The cost of uncertainty – new estimates*. Autumn 2024 Economic Forecast, 15 November. Official webpage.
- Aruqaj, E., A. Bodea, M. Salish and M. Výškrabka (2026). *Spillovers in the EU. Trade, Valuation Changes, and Uncertainty*. Institutional Paper 346, Chapter 3, pp. 27–35. Official PDF.

References and course sources

- David, J. M., R. Giacomini, X. Jiao and W. Wang (2026). *Causal State-Dependent Local Projections in Micro-Macro Data*. Manuscript dated 8 November 2026. Core source for assumptions, estimation, inference and firm results.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161–182.
- Plagborg-Møller, M. and C. K. Wolf (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2), 955–980.
- Montiel Olea, J. L. and M. Plagborg-Møller (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4), 1789–1823.

Course code and synthetic exercises are teaching illustrations. They do not reproduce the empirical dataset or establish finite-sample coverage.

References: identification and nonlinear responses

- Jordà and Taylor (2025). Local Projections. *JEL* 63(1), 59–110.
- Stock and Watson (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal* 128, 917–948.
- Rambachan and Shephard (2025 revision). When do common time series estimands have nonparametric causal meaning? arXiv:1903.01637v4.
- Kolesár and Plagborg-Møller (2025). Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly. *JBES* 43(4), 737–754.
- Kitagawa, Wang and Xu (2025). Nonlinearity in Dynamic Causal Effects: Making the Bad into the Good, and the Good into the Great? *JBES* 43(4), 770–777.
- Gonçalves, Herrera, Kilian and Pesavento (2024). State-dependent local projections. *JoE* 244(2), 105702.

Reading list: priorities, annotations and links. Companion notes: assumptions and derivations.

References: shock construction and controls

- Romer and Romer (2004). A New Measure of Monetary Shocks: Derivation and Implications. *AER* 94(4), 1055–1084.
- Gertler and Karadi (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *AEJ: Macro* 7(1), 44–76.
- Bauer and Swanson (2023). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual*.
- Lloyd and Manuel (2024). Controls, not shocks: estimating dynamic causal effects in macroeconomics. Bank of England Staff Working Paper 1079.

Compare how each design uses information and shock measures before copying its regression specification.