Local Projections

Paulo M. M. Rodrigues

Motivation

- Understanding the dynamic effects of economic shocks, such as changes in monetary policy, on macroeconomic variables like output, inflation, and employment.
- Traditional approach: **Vector Autoregressions (VARs)** are widely used to study such dynamic relationships
- However, VARs can be restrictive: they rely on **full-system estimation**, require **strong stability assumptions**, and often **impose linear and time-invariant relationships**.
- These assumptions may not hold in practice—especially when monetary policy operates in an environment with evolving expectations, structural breaks, or nonlinear responses.

Goal: Develop or apply more flexible methods to estimate how a monetary policy shock today influences economic outcomes over time, accounting for potential complexities in the data.

Local Projections (LPs) – The Basic Idea

Jordà (2005): Estimate impulse responses directly using horizon-specific regressions.

$$y_{t+h} = \alpha_h + \theta_h \cdot \text{shock}_t + \gamma_h' X_t + \varepsilon_{t+h}$$

- θ_h: Estimated impulse response at horizon h — captures the effect of the shock h periods ahead.
- X_t: Set of controls, typically **including lags of y_t**, **lags of the shock variable**, and possibly **time or entity fixed effects**.
- **Estimation** is performed separately for each forecast horizon h, allowing for time-varying dynamics and flexible functional forms.
- **Robust to model misspecification**: avoids iterating the system forward like VARs, reducing bias from incorrect dynamic structure.
- Easy to incorporate nonlinearities (e.g., interactions, thresholds) and heterogeneity (e.g., state-dependent responses, panel data).

Key Advantage: LPs offer a simple and robust alternative to VARs for tracing dynamic causal effects of shocks across time horizons.

Local Projections (LPs) – The Basic Idea (Cont.)

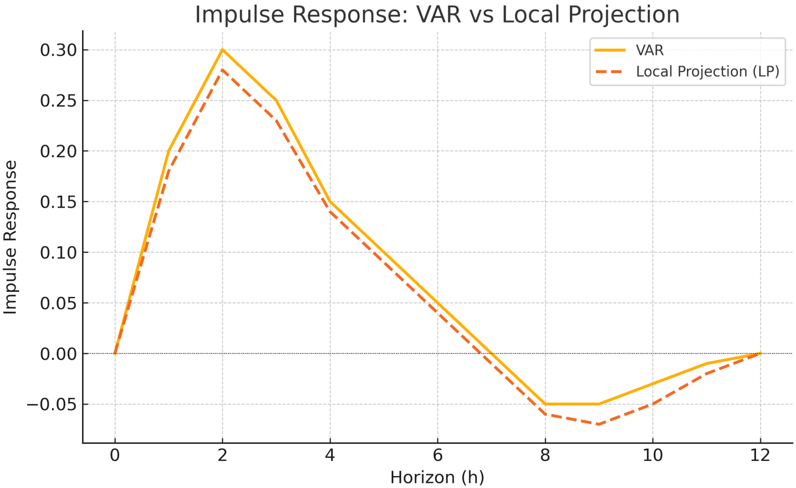

Econometric Reasoning Behind LPs

LPs as a sequence of conditional expectations:

$$\theta_h = \frac{\partial \mathbb{E}[y_{t+h} \mid \text{shock}_t]}{\partial \text{shock}_t}$$

- **Interpretation:** θ_h captures the marginal effect of a shock at time t on the expected value of y at horizon $t+h$
- Estimating impulse responses as derivatives of conditional expectations avoids assumptions about the full dynamic system
- LPs implement this by regressing y_{t+h} on shock_t and controls — effectively estimating $\mathbb{E}[y_{t+h} \mid \text{shock}_t, X_t]$

Econometric Reasoning Behind LPs

- **Non-parametric intuition:** Each horizon h is treated as a separate estimation problem, allowing the response to evolve flexibly over time
- Robust to misspecification of the joint data-generating process (e.g., if dynamics are nonlinear or contain unobserved regime shifts)
- No need to invert a system of equations or rely on recursive structure — LPs directly approximate the conditional mean function at each horizon

Key Insight: LPs reframe dynamic causal inference as a series of predictive problems — estimating how expected outcomes evolve after a shock.

Comparison with VARs

Feature	VARs	Local Projections (LPs)
Estimation	System of equations	Separate regressions
Flexibility	Less	More
Misspecification	Sensitive	More robust
Long horizon noise	Less noisy (if correct)	More noisy
Ease of extension	More complex	Easier

LP Estimation Strategy

For each horizon $h = 0, 1, 2, \ldots, H$, estimate a separate regression:

1. **Define the outcome:** Set the dependent variable as y_{t+h} — the future value of interest (e.g., output, inflation, etc.).
2. **Specify the shock variable:** z_t — this could be an observed policy shock, a residual from a VAR, or an instrumented proxy.
3. **Include controls:** X_t — typically includes:
    - Lags of y_t and z_t to absorb dynamics;
    - Additional covariates or fixed effects to account for confounding;
    - Time dummies to capture seasonal or cyclical effects.

LP Estimation Strategy (Cont.)

Estimation Equation:

$$y_{t+h} = \alpha_h + \theta_h z_t + \gamma_h' X_t + \varepsilon_{t+h}$$

- Estimate each equation separately via OLS;
- If z_t is endogenous, use **instrumental variables (IV)** to isolate exogenous variation;
- Compute standard errors:
  - Use **HAC (e.g., Newey-West)** to correct for serial correlation;
  - In panel settings, cluster at the appropriate level.
- Stack the estimated θ_h across horizons to construct the full impulse response function.

Result: A flexible, horizon-by-horizon estimation of dynamic responses to a one-time shock at t.

LP Estimation Strategy (Cont.)

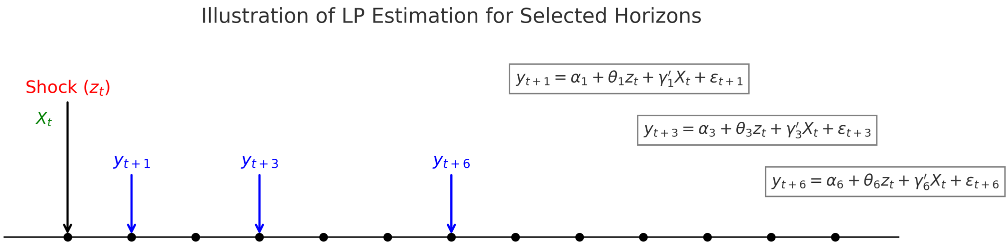

Extensions and Variants

- **Instrumented LPs**: Use external instruments for shocks
- **Nonlinear LPs**: Interactions, threshold effects
- **Panel LPs**: Cross-sectional units (countries, firms)
- **Smooth LPs**: Penalize volatility across horizons

Instrumented Local Projections (IV-LPs)

Goal: Address endogeneity of the shock variable z_t using valid external instruments w_t

$$\text{1st Stage:} \quad z_t = \pi_0 + \pi_1 w_t + \pi_2' X_t + u_t$$

$$\text{2nd Stage:} \quad y_{t+h} = \alpha_h + \theta_h \hat{z}_t + \gamma_h' X_t + \varepsilon_{t+h}$$

- w_t: External instrument satisfying relevance (strongly correlated with z_t) and exogeneity (uncorrelated with ε_{t+h})
- $\hat{z}_t$: Fitted value of z_t from first-stage regression
- Ensures causal interpretation of θ_h when z_t is endogenous
- Common in monetary policy and fiscal policy applications

Estimation: Use 2SLS or GMM at each horizon h

Nonlinear Local Projections

Goal: Allow impulse responses to vary across states or regimes

$$y_{t+h} = \alpha_h + \theta_h^{(1)} z_t \cdot \mathbb{I}(s_t = 1) + \theta_h^{(2)} z_t \cdot \mathbb{I}(s_t = 2) + \gamma_h' X_t + \varepsilon_{t+h}$$

- s_t: Indicator or continuous variable capturing regime/state (e.g., high vs low inflation, expansion vs recession)
- $\theta_h^{(1)}$, $\theta_h^{(2)}$: State-dependent impulse responses
- Can also include smooth interactions: $\theta_h(z_t \cdot f(s_t))$
- Captures nonlinearities and asymmetries in the response

Estimation: Standard OLS or IV with interaction terms

Panel Local Projections

Goal: Estimate impulse responses across multiple cross-sectional units (e.g., countries, regions, firms)

$$y_{i,t+h} = \alpha_{i,h} + \theta_h z_{i,t} + \gamma_h' X_{i,t} + \varepsilon_{i,t+h}$$

- i: Cross-sectional unit; t: time
- Allows for unit-specific fixed effects $\alpha_{i,h}$
- Can pool θ_h across units or estimate heterogeneous responses
- Enhances statistical power, especially with short time series
- Requires clustering standard errors at the unit level

Estimation: Panel OLS/IV with fixed effects and clustered SEs

Smooth Local Projections

Goal: Impose smoothness across horizons to reduce volatility in θ_h estimates

$$\min_{\{\theta_h\}} \sum_{h=0}^{H} \left(\text{RSS}_h + \lambda \cdot \Delta^2 \theta_h \right)$$

- Adds penalty term for roughness in impulse responses (e.g., second-difference of θ_h)
- Encourages more stable and interpretable IRFs, especially in noisy or small samples
- λ: Tuning parameter controlling the smoothness penalty
- Can be viewed as a regularized estimator or a form of shrinkage

Estimation: Penalized least squares or Bayesian priors over θ_h

Instrumented Local Projections (IV-LPs)

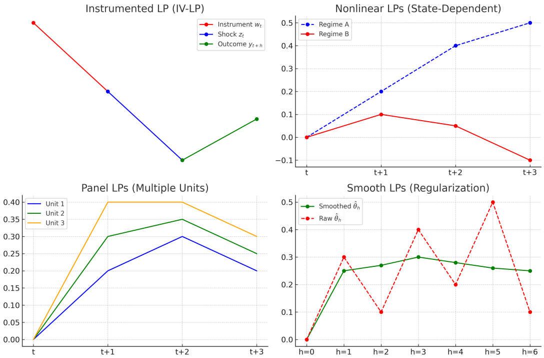

Top-left: Instrumented LP

- Use external instruments w_t to isolate exogenous variation in the shock z_t
- Two-stage estimation:
  1. Stage 1: Predict z_t using w_t
  2. Stage 2: Use predicted $\hat{z}_t$ in LP regressions
- Addresses endogeneity concerns common in macroeconomic applications
- Requires instruments to satisfy relevance and exogeneity

Nonlinear Local Projections

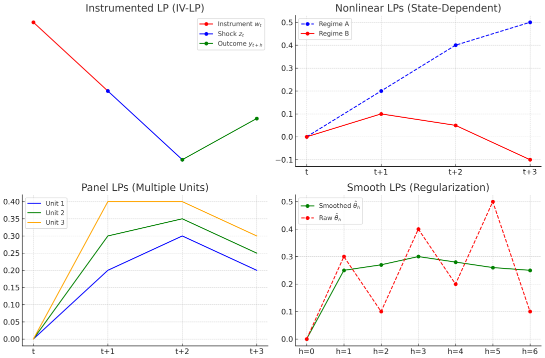

Top-right: Nonlinear LP (State-Dependent IRFs)

- IRFs vary by regime or state (e.g., recession vs expansion)
- Introduce interaction terms: $z_t \cdot \mathbb{I}(s_t = \text{regime})$
- Can capture asymmetries or threshold effects in responses
- Useful for studying fiscal multipliers, monetary policy under ZLB, etc.

Smooth Local Projections

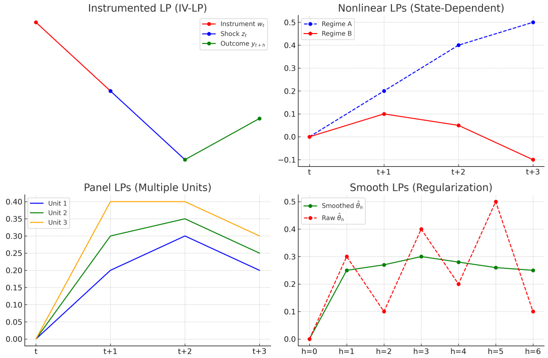

Bottom-right: Smooth LP (Regularized IRFs)

- LP estimates can be noisy across horizons
- Smooth LPs impose a penalty on roughness (e.g., $\Delta^2\theta_h$)
- Encourages coherent and interpretable IRFs
- Can be implemented via penalized regression or Bayesian priors
- Especially useful in small samples

Empirical Illustration

Goal: Estimate the dynamic effects of a monetary policy shock on output and inflation.

- How does a rise in the interest rate affect real GDP and inflation over time?
- We focus on quarterly U.S. macroeconomic data: 1985Q1–2022Q4
- Estimate Impulse Response Functions (IRFs) using Local Projections

Data Overview

Key Variables (quarterly):

- y_t: Log real GDP
- π_t: Inflation (log-difference of CPI or GDP deflator)
- i_t: Federal Funds Rate
- z_t: Monetary policy shock (e.g., high-frequency surprise or Taylor rule residual)

Controls:

- Lags of y_t, i_t, π_t
- Time fixed effects (quarter dummies)

LP Estimation Equation

Jordà (2005) LP setup: For each horizon $h = 0, 1, \ldots, H$ estimate:

$$y_{t+h} = \alpha_h + \theta_h z_t + \gamma_h' X_t + \varepsilon_{t+h}$$

- y_{t+h}: Output or inflation at horizon h
- z_t: Monetary policy shock at time t
- X_t: Controls (lags of endogenous vars + time FE)
- Estimate each h separately via OLS

Estimation Procedure

- Forecast horizon: $H = 12$ (3 years)
- Separate regressions for each h
- Use Newey-West HAC standard errors (4 lags)
- Estimate for both dependent variables:
  - y_{t+h} = real GDP (log)
  - π_{t+h} = inflation (log-diff)
- Stack $\hat{\theta}_h$ to obtain IRFs

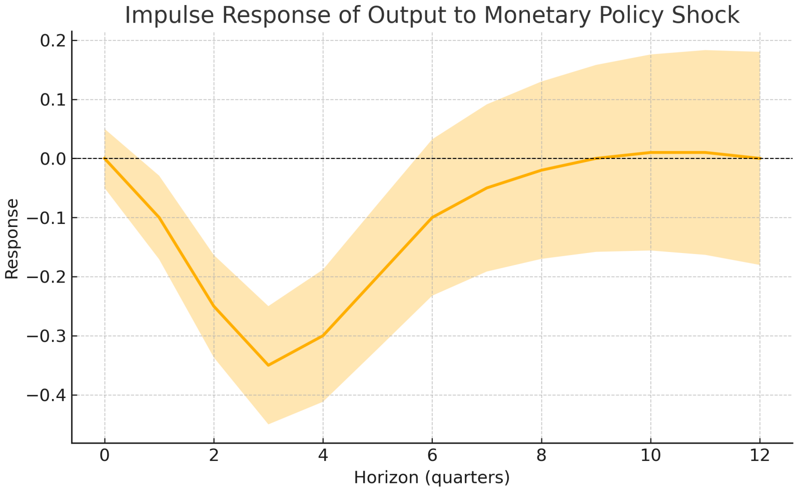
Impulse Response of Output to Monetary Policy Shock
Response
Horizon (quarters)
0.2
0.1
0.0
−0.1
−0.2
−0.3
−0.4
0
2
4
6
8
10
12

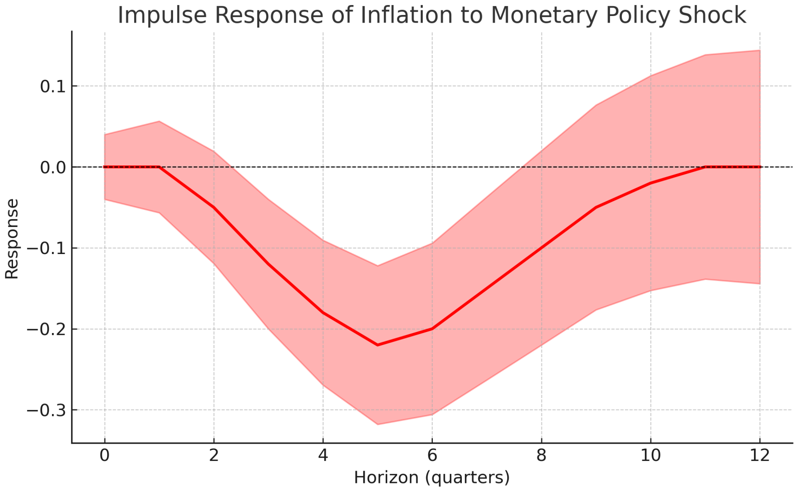
Impulse Response of Inflation to Monetary Policy Shock
0.1
0.0
−0.1
−0.2
−0.3
Response
0
2
4
6
8
10
12
Horizon (quarters)

Impulse Response

Interpretation:

- A contractionary monetary policy shock reduces output for several quarters
- Inflation responds with a lag — consistent with sticky prices
- Confidence bands widen at longer horizons

Inference and Robustness

Inference:

- HAC standard errors (Newey-West) correct for serial correlation
- Uniform confidence bands: Bonferroni, bootstrap, or delta method

Robustness Checks:

- Vary lag length in X_t
- Use alternative shock measures (e.g., narrative shocks)
- Check stability across pre- and post-crisis subsamples

Impulse Responses to Multiple Shocks

- LP easily generalizes to include multiple shocks:

$$y_{t+h} = \alpha_h + \beta_{1h} z_{1t} + \beta_{2h} z_{2t} + \gamma_h' X_t + \varepsilon_{t+h}$$

- Example: responses to monetary and fiscal shocks.
- Separate β_{jh} traces out the IRF to each shock.
- Plot multiple IRFs on the same figure for comparison.

Bootstrap Algorithm for LP Inference

Consider the local projection regression for horizon h:

$$y_{t+h} = \alpha_h + \beta_h z_t + \gamma_h' X_t + \varepsilon_{t+h},$$

where z_t is the shock of interest, X_t is a vector of control variables, and β_h is the impulse response at horizon h.

Step-by-Step Bootstrap Procedure

1. **Estimate** the LP regressions for each horizon $h = 0, 1, \ldots, H$ and store the residuals $\hat{\varepsilon}_{t+h}$.
2. **Resample** residuals:
   - *Residual Bootstrap:* Sample with replacement from $\hat{\varepsilon}_{t+h}$.
   - *Moving Block Bootstrap:* Resample blocks to preserve time dependence.
   - *Wild Bootstrap:* Multiply residuals by random variables (e.g., Rademacher distribution).
3. **Generate bootstrap samples:**
   - Construct new dependent variable $y_{t+h}^{*(b)} = \hat{\alpha}_h + \hat{\beta}_h z_t + \hat{\gamma}_h' X_t + \varepsilon_{t+h}^{*(b)}$
4. **Re-estimate** the LP model on the bootstrap sample and collect $\hat{\beta}_h^{*(b)}$.
5. **Repeat** steps 2–4 for B bootstrap replications.
6. **Construct confidence intervals**:
   - Use empirical percentiles (e.g., 5th and 95th for 90% CI).
   - Alternatively, use standard deviation of $\hat{\beta}_h^{*(b)}$ to compute normal approximation.

Remarks

- The choice of bootstrap method depends on the properties of the data, especially serial correlation.
- Block bootstraps are preferable when residuals exhibit strong time dependence.
- Wild bootstrap is particularly useful for heteroskedastic errors.

References

- Jordà, Ò. (2005). *Estimation and Inference of Impulse Responses by Local Projections*. AER.
- Ramey, V. (2016). *Macroeconomic Shocks and Their Propagation*. Handbook of Macroeconomics.
- Barnichon Brownlees (2019). *Impulse Response Estimation by Smooth Local Projections*. RESTAT.

Quiz Questions

Quiz: Core Concepts

1. What is the econometric interpretation of the impulse response θ_h in a Local Projection?
   A. The expected value of y_t
   B. The partial derivative of $\mathbb{E}[y_{t+h} \mid z_t]$ with respect to z_t
   C. The coefficient on the lag of y_t
   D. The correlation between z_t and y_{t+h}
2. True or False: LPs estimate the full joint distribution of the data-generating process.
3. Why do we run a separate regression for each horizon h in LPs?

Quiz: Estimation Strategy

4. What is typically included in the control vector X_t?
   A. Only future values of y_t
   B. Lags of y_t, z_t, and other covariates
   C. Unrelated variables to reduce variance
   D. None of the above
5. What is the purpose of using HAC standard errors in LPs?

Quiz: Extensions

6. Which conditions must a valid instrument w_t satisfy in IV-LPs?
7. Which LP extension best captures asymmetric responses in different macroeconomic regimes?
8. Why do we cluster standard errors in panel LPs?
9. What is the benefit of using smooth LPs?

Quiz: Inference and Confidence Intervals

10. Why are standard errors in LPs often estimated using Newey-West or other HAC methods?
    A. Because errors are heteroskedastic and serially correlated across t
    B. To reduce bias in small samples
    C. Because standard OLS errors are always invalid
    D. To correct for panel fixed effects
11. What is a key drawback of using separate confidence intervals for each horizon h?
    A. They overfit the model
    B. They do not account for the joint uncertainty over the entire IRF
    C. They assume perfect foresight
    D. They cannot be computed using OLS
12. What are some approaches to constructing uniform confidence bands for impulse responses?
    A. Bonferroni correction
    B. Bootstrap-based joint inference
    C. Delta method with multiple-testing adjustment
    D. All of the above

Quiz: Applied Reasoning

13. You're studying monetary shocks during periods of high and low inflation. Which LP extension should you use and why?
14. What could happen if your instrument in IV-LPs is weak?

Empirical Illustration: Fiscal Shock

- **Question:** What is the effect of a government spending shock on output?
- **Data:** Quarterly US data (GDP, government spending, interest rates), 1960–2020.
- **Shock:** Structural spending shock from a narrative approach (e.g., Ramey 2011).
- **Specification:**

$$y_{t+h} = \alpha_h + \beta_h \text{Shock}_t + \gamma_h' X_t + \varepsilon_{t+h}$$

- y_{t+h}: log real GDP at horizon h
- Controls: lags of GDP and shock

Code Example (in Python)

```
import statsmodels.api as sm
import numpy as np

H = 20  # horizon
betas = []
for h in range(H+1):
    y_lead = y.shift(-h)  # dependent variable at t+h
    X_reg = sm.add_constant(pd.concat([shock, controls], axis=1))
    model = sm.OLS(y_lead, X_reg, missing='drop').fit(cov_type='HAC', cov_kwds={'maxlags':4}
    betas.append(model.params['shock'])
```

Quiz on Bias-Variance Trade-off

- Why do VAR estimates generally have lower standard errors than LP?
- What risks are associated with VAR's extrapolation approach?

Control Variables and Lag Length

- Controls ensure valid identification by removing predictable components of the shocks.
- Controls increase estimation efficiency by reducing residual variance, thus improving standard errors.
- Lag augmentation helps LP remain robust to dynamic misspecification by explicitly capturing important past dynamics.
- Lagged controls simplify confidence interval construction by ensuring residual errors are less correlated.

The End

Quiz: Core Concepts

1. What is the econometric interpretation of the impulse response θ_h in a Local Projection?
   A. The expected value of y_t
   B. The partial derivative of $\mathbb{E}[y_{t+h} \mid z_t]$ with respect to z_t
   C. The coefficient on the lag of y_t
   D. The correlation between z_t and y_{t+h}

   Answer: B
2. True or False: LPs estimate the full joint distribution of the data-generating process.
   Answer: False
3. Why do we run a separate regression for each horizon h in LPs?
   Answer: To flexibly estimate dynamic effects without imposing a parametric structure on time dynamics.

Quiz: Estimation Strategy

4. What is typically included in the control vector X_t?
   A. Only future values of y_t
   B. Lags of y_t, z_t, and other covariates
   C. Unrelated variables to reduce variance
   D. None of the above

   Answer: B
5. What is the purpose of using HAC standard errors in LPs?
   Answer: To account for autocorrelation and heteroskedasticity in residuals across forecast horizons.

Quiz: Extensions

6. Which conditions must a valid instrument w_t satisfy in IV-LPs?
   Answer: It must be relevant (correlated with z_t) and exogenous (uncorrelated with ε_{t+h}).
7. Which LP extension best captures asymmetric responses in different macroeconomic regimes?
   Answer: Nonlinear LPs
8. Why do we cluster standard errors in panel LPs?
   Answer: To correct for within-unit autocorrelation and maintain valid inference.
9. What is the benefit of using smooth LPs?
   Answer: They stabilize IRFs across horizons and reduce noise by penalizing roughness in θ_h.

Quiz: Inference and Confidence Intervals

10. Why are standard errors in LPs often estimated using Newey-West or other HAC methods?
    A. Because errors are heteroskedastic and serially correlated across t
    B. To reduce bias in small samples
    C. Because standard OLS errors are always invalid
    D. To correct for panel fixed effects

    Answer: A
11. What is a key drawback of using separate confidence intervals for each horizon h?
    A. They overfit the model
    B. They do not account for the joint uncertainty over the entire IRF
    C. They assume perfect foresight
    D. They cannot be computed using OLS

    Answer: B

Quiz: Applied Reasoning

12. What are some approaches to constructing uniform confidence bands for impulse responses?
    A. Bonferroni correction
    B. Bootstrap-based joint inference
    C. Delta method with multiple-testing adjustment
    D. All of the above

    Answer: D
13. You're studying monetary shocks during periods of high and low inflation. Which LP extension should you use and why?
    Answer: Nonlinear LPs, because they allow impulse responses to vary across regimes using interaction terms.
14. What could happen if your instrument in IV-LPs is weak?
    Answer: The first-stage regression will poorly identify z_t, leading to biased and inconsistent estimates of θ_h.