BLP

Bayesian Local Projections

EViews Add-in Documentation

Dulguun Lkhagvadorj
Virginia Tech, Department of Economics

Version 1.0 · June 2026

Contents

---

1 Overview

`BLP` estimates impulse response functions to a structural shock using the Bayesian Local Projection (BLP) method of Ferreira, Miranda-Agrippino, and Ricco (2023), ported from the authors' original MATLAB toolbox into native EViews matrix-language code. It brings Bayesian shrinkage — with the prior's tightness chosen automatically at every horizon by maximizing the marginal likelihood — to the local-projection framework of Jordà (2005).

Key features:

- Horizon-by-horizon local projections $h = 0, \ldots, H$, each with its own Normal-Inverse-Wishart (Minnesota-style) prior whose overall tightness λ_h is re-optimized at that horizon via marginal-likelihood maximization (a 1-D golden-section search, replacing the original toolbox's `csminwel` optimizer).
- Random-Walk prior mean (coefficients centered on the identity on the first own lag), matching the toolbox's `RW` prior-type case.
- Cholesky structural identification, with the shocked variable's own on-impact response normalized to the requested shock size.
- Closed-form Normal-Inverse-Wishart posterior simulation for the $h = 0, 1$ error bands; analytical Newey–West (HAC) sandwich/delta-method bands for $h > 1$ — matching the simulation/analytical split used in the original toolbox.

2 Installation

2.1 Add-ins manager registration

Copy the `BLP` folder into your EViews Add-ins directory (`Add-ins` → `Manage Add-ins` will show you the path), then load it via:

`Add-ins` → `Manage Add-ins` → `Load`

- **Installer file:** `blp_addin.prg`
- **Command keyword:** `blp`
- **Type:** Global
- **Menu entry:** "Bayesian Local Projections..."

Alternatively, for the current session only, open `blp_addin.prg` in EViews and click **Run**.

2.2 Requirements

- EViews 10 or later.
- An active workfile containing the series for the system (any frequency), with the shock variable among them.

3 Syntax

3.1 GUI dialog

Selecting `Add-ins` → `Bayesian Local Projections...` (or typing `blp` at the command line) opens a dialog asking for:

- the series in the system (space-separated names, e.g. `rgdp rcon ffr`);
- the shock variable (must be one of the names above);
- number of lags, maximum horizon, error-band coverage (%);
- number of posterior draws used for the $h = 0, 1$ bands;
- the size of the shock;
- an output name prefix.

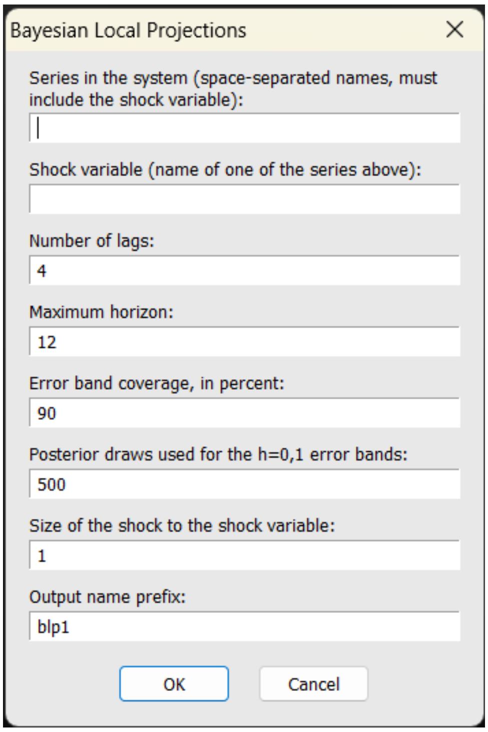

Figure 1: GUI dialog

4 Estimation Details

4.1 Model

For horizons $h = 0,1$, the add-in estimates a Bayesian VAR(p),

$$y_t = a + \sum_{j=1}^{p} B_j y_{t-j} + \varepsilon_t, \tag{1}$$

and reads the impulse responses off its posterior mean. For horizons $h \geq 2$, it estimates, independently at each h, the local projection

$$y_{t+h} = a_h + \sum_{j=1}^{p} B_j^h y_{t-j} + u_{t+h}^h. \tag{2}$$

In both cases the coefficient matrices carry a Normal-Inverse-Wishart (Minnesota-style) prior centered on a random walk: $B_1 \sim N(I, \cdot)$, all other lags and the constant centered at zero. The prior variance of the coefficient on lag j of variable i is

$$\omega_j^{(i)} = \frac{\lambda^2}{j^2\, \psi_i}, \tag{3}$$

where ψ_i is a residual-variance scale (the univariate AR(1) residual variance at the VAR step; a Newey–West-corrected own-lag residual variance at each local-projection horizon) and λ is the overall tightness.

4.2 Hyperparameter optimization

λ is chosen independently at every horizon by maximizing the log marginal likelihood (closed-form, following Giannone, Lenza, and Primiceri's 2015 dummy-observation BVAR, combined with a Gamma (mode $= 0.4$) hyperprior on λ whose dispersion widens with horizon). Because this is the *only* hyperparameter left free under the toolbox's own default options (lag-decay α fixed at 2, residual-scale ψ fixed rather than estimated, no sum-of-coefficients or cointegration dummies), the optimization reduces to a

1-D search, implemented here as a golden-section search over $\lambda \in [10^{-4}, 5]$ — a faithful, numerically simpler replacement for the original toolbox's general-purpose `csminwel` optimizer.

4.3 Structural identification

The contemporaneous impact matrix B_0 is the Cholesky factor of the (posterior-mean) residual covariance matrix, with each row rescaled so its own diagonal entry equals 1. With the shock variable ordered last in the system, this is the standard recursive identification: the shock has no contemporaneous effect on the other variables, and the shocked variable's own on-impact response equals the calibrated shock size exactly.

4.4 Error bands

- $h = 0, 1$**:** i.i.d. Monte Carlo draws from the exact (conjugate) Normal-Inverse-Wishart posterior — an Inverse-Wishart draw for Σ via Bartlett decomposition, a matrix-Normal draw for $B \mid \Sigma$, and a fresh Cholesky identification at each draw. Bands are the empirical quantiles of the resulting draws.
- $h > 1$**:** analytical Newey–West (HAC) sandwich covariance of the horizon-h local-projection coefficients, mapped through B_0 via the delta method — no simulation at these horizons, matching the original toolbox.

5 Examples

First, run the smoke test: Open **blp_test.prg** and Run it. It builds its own 3-variable simulated workfile and runs the full engine end to end. A dialog reports PASS/FAIL. Do this before pointing the add-in at your own data, and if it fails, report the exact error text and line number – that pinpoints the fix needed.

5.1 Real-data replication

blp_demo_JPTQ.prg replicates the Leonardo N. Ferreira, Silvia Miranda-Agrippino, Giovanni Ricco (2023) on FRED quarterly data (1965Q1–2019Q4): a Cholesky-identified one-percentage-point FFR shock, 5 lags, 20 horizons, 90% bands.

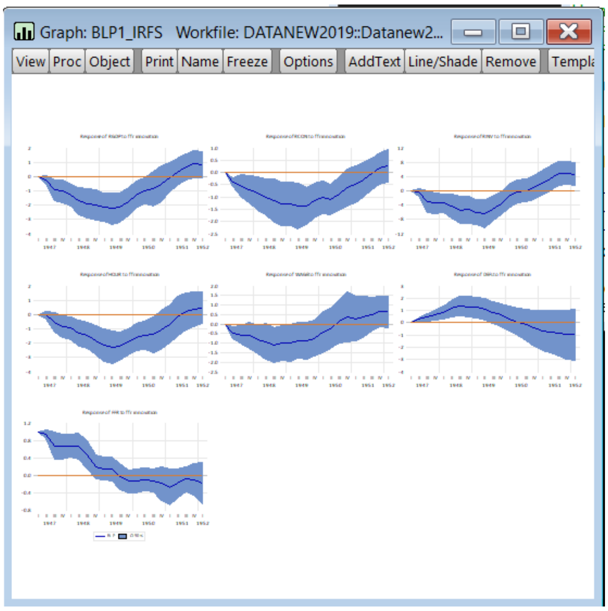

Figure 2: FRED replication (`gr_jptq_all`): responses to a Cholesky-identified FFR innovation

6 Scope and Limitations

This add-in implements only the Random-Walk-prior, Cholesky-identification path of the original MATLAB toolbox. The following toolbox features are *not* included here:

- VAR-prior and DSGE-prior variants for the local-projection coefficient mean (only the Random-Walk prior mean is implemented).
- Proxy-SVAR / external-instrument identification (only Cholesky).
- The plain (non-Bayesian) local-projection and unrestricted BVAR comparison output that the toolbox's main replication script also produces.

Users needing those variants should refer to the original MATLAB toolbox.

7 References

- Ferreira, L. N., Miranda-Agrippino, S., and Ricco, G. (2023). Bayesian Local Projections. Working paper / MATLAB toolbox.
- Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics*, 97(2):436–451.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.
- Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions: Five years of experience. *Journal of Business & Economic Statistics*, 4(1):25–38.
- Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.