Bayesian Panel Local Projections

Marco Schwarzbach*
Justus Liebig University Gießen

March 25, 2026

Abstract

Local projections are widely used for impulse response analysis, but in panel settings they face a familiar trade-off: imposing homogeneous responses across units can be too restrictive, while estimating each unit separately can be too noisy, especially at longer horizons. This paper proposes a Bayesian Panel Local Projections (BPLP) framework that combines horizon-by-horizon local projection estimation with hierarchical partial pooling and structured regularization across units and horizons. Estimation is based on a Gaussian working likelihood, while uncertainty is quantified separately using misspecification-robust long-run variance methods. The framework is designed to stabilize heterogeneous unit-level impulse responses, preserve comparability across units, and provide transparent panel and subgroup benchmarks. We apply the framework to monetary policy transmission in eight euro area economies using high-frequency identified ECB shocks. The estimates suggest a negative inflation response to contractionary monetary policy, alongside notable heterogeneity in the timing and magnitude of adjustment across countries. In particular, the Periphery is estimated to exhibit earlier and stronger inflation responses than the Core, while some countries show persistent deviations from the cross-country benchmark. The proposed framework provides a practical and interpretable tool for applied impulse response analysis in heterogeneous panels.

JEL codes: C11, C32, C33, E52, E58
Keywords: Bayesian panel local projections, cross-unit heterogeneity, hierarchical pooling, impulse responses, monetary policy transmission, panel data

*Justus Liebig University Gießen, Marco.Schwarzbach@uni-giessen.de. All errors are my own.

Contents

1 Introduction

Empirically, dynamic responses to shocks are often studied using vector autoregressions (VARs). VARs are efficient and provide a compact summary of aggregate dynamics, but they impose tightly parameterized structure and rely on the adequacy of a finite-order approximation to recover impulse responses. Furthermore, when a VAR is misspecified, its associated impulse responses can be biased, raising concerns about reliability in empirical applications.[1] Local projections (LPs), introduced by Jordà (2005), provide a flexible alternative by estimating each horizon separately. LPs are straightforward to implement and interpret, are often more robust to misspecified lag dynamics, and naturally extend to nonlinear environments (Montiel Olea & Plagborg-Møller, 2021; Plagborg-Møller & Wolf, 2021). This flexibility, however, comes with higher variance relative to VARs, so the choice between VARs and LPs reflects a bias–variance trade-off (Ferreira et al., 2025). The trade-off becomes especially sharp in panel settings, where the time dimension is often short while the parameter space expands quickly. Estimating LPs separately for each cross-sectional unit can therefore yield unstable impulse responses and imprecise inference, particularly at longer horizons, creating a natural role for regularization.

Several recent contributions introduce Bayesian local projection approaches that stabilize estimation through shrinkage and smoothing across horizons. Yet most of the existing Bayesian LP literature is focused on univariate settings and does not develop a general framework for sharing information across units in a panel without resorting to full pooling. What remains largely missing, to our knowledge, is a framework that combines Bayesian regularization across horizons with a principled mechanism for partial pooling across units in panel local projections. This paper develops such a framework: Bayesian Panel Local Projections (BPLP) integrate horizon-by-horizon local projections with hierarchical partial pooling across cross-sectional units, allowing unit-specific impulse responses to differ while letting the data determine the degree of commonality. The resulting structure is designed to deliver more stable impulse response estimates and a coherent separation between regularized estimation and robust inference, while preserving comparability across units.

Frequentist panel LP approaches with unit-specific effects typically treat units as unrelated, forcing a choice between precision (full pooling) and heterogeneity (fully separate estimation). Yet in many applications units are neither identical nor independent, making neither extreme particularly appealing. A hierarchical (multilevel) specification offers a natural intermediate position: unit-specific impulse responses are allowed to differ, but they are modeled as draws from a common distribution centred on a cross-sectional benchmark. The degree of similarity is inferred from the data, delivering more stable estimates while preserving economically meaningful heterogeneity.

[1] A related motivation is that macroeconomic time series may not be well approximated by low-order VARs, because early evidence suggests that many variables are more naturally represented by vector autoregressive moving-average processes (Wallis, 1977; Zellner & Palm, 1974).

The contributions of the paper are threefold. First, it develops the BPLP methodology, a general framework that combines local projections, Bayesian regularization, and hierarchical partial pooling for panels subject to common shocks. Second, it illustrates the framework by quantifying heterogeneity in monetary policy transmission within the European Monetary Union (EMU), showing how a common policy stance can generate a shared directional response alongside heterogeneous adjustment dynamics. In this sense, the primary contribution of the paper is methodological, while the euro-area application is intended both to illustrate the framework and to show that it yields economically interpretable results in a setting where cross-country heterogeneity is central. Third, it provides a transparent and replicable structure that can be adapted to other contexts involving state dependence, subgroup analysis, or cross-sectional interactions.

To demonstrate the usefulness of the framework in a setting where heterogeneity is central, we apply BPLP to monetary policy shocks in a subgroup of euro area member countries. The euro area provides a natural laboratory: countries share a common monetary policy stance, yet differ in economic structure and financial conditions, so the transmission of a common shock need not be uniform. Recent episodes of heightened macroeconomic uncertainty underscore the value of estimating impulse responses in a way that is both flexible across horizons and disciplined by cross-sectional information sharing.[2]

The remainder of the paper is organized as follows. Section 2 reviews the related literature on heterogeneous transmission and local projections. Section 3 introduces the BPLP framework and prior specification. Section 4 presents the empirical results, and Section 5 concludes.

[2]One prominent example is the Covid-19 pandemic and its aftermath, which amplified cross-country differences in inflation dynamics and real activity while monetary policy conditions were shared within the EMU.

2 Related Literature

This paper contributes to the empirical literature on heterogeneous dynamic responses to common shocks in panel settings and to the growing methodological literature that uses local projections with shrinkage and Bayesian structure. A long tradition studies monetary policy transmission in Europe with the objective of assessing whether a single monetary authority can operate effectively across economically diverse member states. Early contributions, using pre-EMU national data and VAR frameworks, document noticeable cross-country differences in responses to monetary policy shocks (Mojon & Peersman, 2001; Peersman, 2004; Peersman & Smets, 2001). Subsequent work connects these differences to structural developments and evolving propagation mechanisms. For example, Ciccarelli and Rebucci (2006) show that output sensitivity to monetary policy changed considerably in the run-up to the EMU, while Georgiadis (2015) illustrates how sectoral composition and labor market rigidities can generate asymmetric propagation.

More recent research moves from documenting heterogeneity toward identifying concrete channels through which it arises. Holm-Hadulla and Thürwächter (2021) highlight how differences in firm balance sheets translate into heterogeneous responses, Almgren et al. (2022) emphasize the role of household liquidity, and Pica (2021) discuss collateral and credit mechanisms in housing markets. Datsenko and Fleck (2024) show that economies with a larger manufacturing sector experience more pronounced contractions following a monetary tightening. Broader perspectives on fragmentation and institutional frictions include Lane (2012) on the sovereign debt crisis and Wächter et al. (2024) on Taylor-gap differentials. Taken together, this literature motivates empirical frameworks that can quantify heterogeneity in a disciplined way, especially when shocks are common but propagation differs across units.

The econometric relevance of such heterogeneity has motivated work that allows unit-specific dynamics while borrowing strength across the cross-section. Panel VAR approaches combine shared dynamics with unit-level behavior and use cross-sectional information to improve precision while permitting heterogeneity (Canova & Ciccarelli, 2013). Empirical applications confirm that heterogeneity is quantitatively important: Mandler et al. (2022), employing a multicountry BVAR, finds substantial variation in the magnitude and persistence of monetary policy effects across EMU member states. Related evidence is provided by Afonso et al. (2025), who links differences in fiscal regimes to heterogeneous transmission, and by Cima and Moreno (2025), who emphasize how wealth distribution shapes asymmetric macroeconomic outcomes.

At the same time, VAR-based approaches impose a tightly parameterized dynamic structure, which can be restrictive in short samples and vulnerable to misspecification when underlying dynamics are more complex than a finite-order VAR can capture. This has sustained interest in local projections (LPs), introduced by Jordà (2005), which estimate dynamic responses horizon by horizon. LPs are straightforward to implement and interpret and are comparatively robust to

misspecified lag dynamics. Methodological contributions include Jordà (2009) on simultaneous inference and Plagborg-Møller and Wolf (2021), who clarify conditions under which LPs and VARs recover the same structural impulse responses. Recent work such as Li et al. (2024) and Olea et al. (2025) emphasizes the bias–variance trade-off inherent in LP estimation and motivates shrinkage as a way to stabilize horizon-specific estimates, especially when the time dimension is short and long-horizon responses are noisy.

These considerations have led to a rapidly growing literature that introduces Bayesian structure into LP estimation. Ferreira et al. (2025) regularize coefficients across horizons using informative priors, improving finite-sample behavior and producing smoother impulse response profiles. Tanaka (2025) propose a quasi-Bayesian estimator based on a Laplace-type quasi-likelihood that yields well-calibrated inference and accommodates LP-IV settings. Brugnolini et al. (2023) develop Bayesian Flexible Local Projections capable of capturing nonlinear and time-varying impulse responses. However, most Bayesian LP methods remain univariate in scope and therefore do not provide a principled, data-driven mechanism to share information across units in a panel while preserving heterogeneity. This limitation is particularly salient in applications with genuinely common shocks, where separate estimation wastes information while full pooling can mechanically suppress heterogeneity.

A closely related methodological issue concerns inference under the error structure induced by LP construction. With overlapping horizons, LP forecast errors are typically serially correlated and heteroskedastic, and panel aggregates can inherit cross-sectional dependence. In such settings, interpreting a Gaussian likelihood as an exact data-generating model is often not compelling. A growing literature therefore views Gaussian specifications as convenient working approximations used to obtain stable regularized estimates, while uncertainty statements are calibrated in a misspecification-robust way. This perspective builds on classical results on quasi-maximum likelihood under misspecification (Huber, 1967; White, 1982) and on the misspecification-robust "artificial posterior" approach of Müller (2013). In practice, it corresponds to anchoring inference in HAC long-run variance estimators (Newey & McFadden, n.d.) and, for panel aggregates, in DK-type long-run variance estimators intended to remain robust to broad forms of cross-sectional dependence under fixed-N, large-T asymptotics (Driscoll & Kraay, 1998). The key implication is that regularization and inference can be separated: Bayesian structure can stabilize high-dimensional, high-variance estimates, while uncertainty can be reported in a way that is coherent with the dependence structure induced by LP estimation.

Overall, the existing literature delivers three relevant insights. First, heterogeneity in monetary policy transmission within the EMU is empirically important and can be traced to structural differences and identifiable channels. Second, LPs provide a flexible and transparent way to trace dynamic effects, but their horizon-by-horizon estimates can be noisy in short samples and benefit from stabilization. Third, Bayesian shrinkage and smoothing can improve the empirical performance

of LPs, yet existing approaches largely do not exploit cross-sectional information in a way that adapts endogenously to the degree of commonality across units. What remains largely missing, to our knowledge, is a framework that combines the flexibility of LPs with Bayesian regularization across horizons and a principled mechanism for partial pooling across units, paired with misspecification-robust inference appropriate for overlapping-horizon dependence and cross-sectional aggregation. The Bayesian Panel Local Projections framework developed in this paper is designed to address this gap.

3 Bayesian Panel Local Projections Framework

This section introduces the Bayesian Panel Local Projections (BPLP) framework. The objective is to estimate impulse responses in macroeconomic panels while (i) preserving the transparency of local projections, (ii) stabilizing estimation in short samples and at longer horizons through regularization, and (iii) allowing for cross-sectional heterogeneity without imposing either full pooling or fully separate estimation. A central feature of the framework is the separation between estimation and inference: coefficient updates are obtained under a Gaussian working likelihood that is convenient for regularization and computation, whereas uncertainty is quantified using pointwise misspecification-robust methods. Accordingly, estimation is based on a regularized quasi-posterior, while reported uncertainty is calibrated separately using robust long-run variance methods.

3.1 Panel Local Projections: Model Setup

The BPLP framework builds on the Local Projections approach of Jordà (2005). For a single unit, let $y_t \in \mathbb{R}^n$ denote an $n \times 1$ vector of outcomes and let z_t denote an identified scalar shock. Then, the horizon-h local projection is

$$y_{t+h} = c_h + \alpha_h z_t + \sum_{\ell=1}^{L} \Phi_{h,\ell}\, y_{t-\ell} + u_{t+h}, \tag{1}$$

where c_h is an intercept, $\alpha_h \in \mathbb{R}^n$ collects the horizon-h impulse responses to the shock, $\Phi_{h,\ell}{}_{\ell=1}^{L}$ are lag coefficients, and u_{t+h} is the h-step ahead forecast error. Since $y_t = (y_{t,1}, \ldots, y_{t,n})'$ stacks the n outcome variables, each horizon-h projection can be viewed as a system of n equations estimated separately horizon by horizon.

We now extend this setup to a panel with N units observed over T periods. Throughout, units are indexed by $i = 1, \ldots, N$, time by $t = 1, \ldots, T$, horizons by $h = 0, \ldots, H$, and lags by $\ell = 1, \ldots, L$. The horizon index has its usual interpretation: $h = 0$ is the impact response, and $h > 0$ is the response h periods after the shock. For unit i, let $y_{i,t} \in \mathbb{R}^n$ denote the outcome vector and let z_t denote the realization of the identified shock. The horizon-h local projection for unit i is

$$y_{i,t+h} = c_{i,h} + \alpha_{i,h} z_t + \sum_{\ell=1}^{L} \Phi_{i,h,\ell}\, y_{i,t-\ell} + u_{i,t+h}, \qquad t = 1, \ldots, T_h, \tag{2}$$

where, $T_h = T - L - h$ is the effective sample size after constructing lags and aligning the horizon-h dependent variable. The vector $\alpha_{i,h} \in \mathbb{R}^n$ contains the unit-specific impulse responses at horizon h, while $\Phi_{i,h,\ell}{}_{\ell=1}^{L}$ allows for heterogeneous dynamic adjustment across units. Deterministic terms such as time trends can be added without changing the basic structure of the model.

For compact notation, let $X_{i,h}$ denote the $T_h \times k$ regressor matrix collecting an intercept, the shock, and L lags of $y_{i,t}$, where k is the total number of regressors. Furthermore, let $Y_{i,h}$ be the

$T_h \times n$ matrix stacking the horizon-h outcomes, and collect coefficients in the $k \times n$ matrix

$$B_{i,h} = \begin{bmatrix} c_{i,h} \\ \alpha_{i,h} \\ \Phi_{i,h,1} \\ \vdots \\ \Phi_{i,h,L} \end{bmatrix}, \qquad B_{i,h} \in \mathbb{R}^{k \times n}. \tag{3}$$

Stacking over time gives the horizon-h system

$$Y_{i,h} = X_{i,h} B_{i,h} + U_{i,h}, \tag{4}$$

where $U_{i,h}$ is the $T_h \times n$ matrix of forecast errors.

Equation (4) highlights the two familiar polar cases. Estimating each unit separately corresponds to no pooling, whereas imposing common coefficients across units corresponds to full pooling. In macroeconomic panels, neither extreme is typically attractive. Units often share broad structural features, but may still differ meaningfully in the transmission of shocks. The BPLP framework targets the intermediate case by using a hierarchical structure that lets the data determine the degree of cross-unit similarity, thereby allowing for heterogeneous unit-level impulse responses while delivering a coherent panel response.

3.2 Gaussian Likelihood and Misspecification

Local projections generate horizon-h forecast errors whose dependence structure is induced by the way outcomes are constructed across horizons. With overlapping horizons, these forecast errors are typically serially correlated and may also be heteroskedastic. In multivariate settings, they may additionally be contemporaneously correlated across outcomes. A fully parametric model for this dependence structure is rarely compelling in applied work and would substantially increase complexity.

The BPLP framework therefore uses a Gaussian working likelihood to obtain stable and regularized updates for the horizon-specific coefficient matrices. The resulting object is interpreted as a quasi-posterior, because under misspecification, likelihood-based updates concentrate around pseudo-true parameter values, while first-order sampling uncertainty is characterized by a sandwich form (Huber, 1967; White, 1982).

Crucially, reported uncertainty is not taken from the dispersion implied by this working likelihood. Instead, inference follows the misspecification-robust "artificial posterior" logic of Müller (2013). Intuitively, the Gaussian likelihood is used as a convenient device for obtaining regularized coefficient updates, while posterior dispersion is replaced by a robust long-run variance calibration

that is better aligned with the sampling uncertainty of the object of interest. Concretely, dispersion is calibrated using heteroskedasticity- and autocorrelation-consistent long-run variance estimators that capture the serial dependence induced by overlapping horizons. For aggregate objects, such as subgroup impulse responses or panel benchmark responses, uncertainty is calibrated using robust HAC/DK-type long-run variance estimators.[3]

Conditional on the regressors, the horizon-specific coefficient matrix $B_{i,h}$ is updated using the matrix-normal working likelihood

$$Y_{i,h} \mid B_{i,h} \;\sim\; \mathcal{MN}_{T_h \times n}\Big(X_{i,h}B_{i,h},\, I_{T_h},\, \Sigma_{i,h}\Big). \tag{5}$$

The matrix $\Sigma_{i,h} \in \mathbb{R}^{n\times n}$ is not interpreted as the true covariance matrix of the LP forecast errors. Instead, it is a fixed working covariance matrix that scales the Gaussian update across outcomes within a given horizon-h regression.[4] To construct $\Sigma_{i,h}$, define plug-in horizon-h regression residuals

$$\widehat{U}_{i,h} = Y_{i,h} - X_{i,h}\widehat{B}_{i,h}, \tag{6}$$

where $\widehat{B}_{i,h}$ is obtained from the initialization step. The matrix $\Sigma_{i,h}$ is then formed as a plug-in residual covariance estimate based on $\widehat{U}_{i,h}$ and, in line with the quasi-likelihood specification, is held fixed when drawing $B_{i,h}$, though it may vary across units and horizons. Equivalently, (5) can be written in vectorized form,

$$\operatorname{vec}(Y_{i,h}) \mid B_{i,h} \sim \mathcal{N}_{T_h n}\Big(\operatorname{vec}\big(X_{i,h}B_{i,h}\big),\, I_{T_h} \otimes \Sigma_{i,h}\Big). \tag{7}$$

This representation is useful because it shows directly how the Gaussian working update combines conveniently with Bayesian regularization, while robust uncertainty quantification is handled separately.

3.3 Gaussian Prior Structure

The prior for the horizon-specific coefficient matrices $\{B_{i,h}\}$ introduces two complementary forms of regularization. First, a Minnesota-type component shrinks lag coefficients toward economically plausible values and thereby stabilizes the dynamic part of the local projection, especially when

[3] Appendix A.10, when included, provides an application-specific fixed-N, fixed-H, large-T_h justification for this pointwise robust long-run variance calibration. Throughout, the reported intervals should be interpreted as pointwise misspecification-robust frequentist intervals around a regularized center, rather than as exact Bayesian credible sets or simultaneous bands over the full horizon path.

[4] More precisely, $\Sigma_{i,h}$ determines the quadratic form used in the Gaussian working likelihood and hence the metric in which deviations of $Y_{i,h}$ from $X_{i,h}B_{i,h}$ are measured. If $\Sigma_{i,h}$ is diagonal, each outcome equation is scaled by its own residual variance. If $\Sigma_{i,h}$ is unrestricted, the same logic extends to a Mahalanobis-type metric that also accounts for contemporaneous covariance across outcomes. Thus, $\Sigma_{i,h}$ improves the geometry of the quasi-posterior update by providing scale normalization and, when allowed, covariance adjustment across equations. It is not intended to capture the serial dependence induced by overlapping horizons; that dependence is handled separately through the misspecification-robust long-run variance calibration used for inference.

effective sample sizes become small at longer horizons. Second, a hierarchical component introduces partial pooling across units for the coefficients on the identified shock, allowing unit-level impulse responses to remain heterogeneous while linking them through a common cross-sectional structure. Both components are Gaussian and therefore combine naturally with the Gaussian working likelihood, yielding conditionally Gaussian updates for $B_{i,h}$. Furthermore, regularization is targeted: the Minnesota-type prior acts on lag dynamics, whereas the hierarchical prior acts on the shock response coefficients. This separation stabilizes propagation dynamics without mechanically shrinking the impulse responses through the Minnesota-type component.

3.3.1 Minnesota-type Prior for Lag Dynamics.

The first regularization component is a Minnesota-type prior. Its role is to stabilize the dynamic part of the local projection, especially at longer horizons where effective sample sizes are smaller and unrestricted lag dynamics can become noisy. As usual, the prior mean reflects the distinction between persistent variables in levels and stationary or detrended variables: the first own lag is centered at unity for variables in levels and at zero for stationary variables, while all remaining coefficients are centered at zero,

$$b_{0,h}^{(\ell=1)} = \begin{cases} I_n, & \text{for variables expressed in levels,} \\ 0_n, & \text{for stationary or detrended variables.} \end{cases} \tag{8}$$

This shrinkage mechanism is targeted to the coefficients of lagged outcomes, which are shrunk toward the prior means set in Equation 8 using a horizon-specific tightness parameter λ_h together with the usual lag-decay structure. By contrast, deterministic components are given diffuse prior dispersion, and the coefficients on the identified shock are excluded from Minnesota-type shrinkage. Their regularization comes instead through the cross-unit pooling structure introduced below. This separation stabilizes lag dynamics without mechanically attenuating impulse responses.

Conditional on $\lambda_h > 0$, the Minnesota component can be written as a Gaussian prior centered at $b_{0,h}$, with coefficient-specific shrinkage governed by the diagonal precision matrix $\Omega_{i,h}(\lambda_h)$. Specifically,

$$B_{i,h} \mid \lambda_h, \Sigma_{i,h} \sim \mathcal{MN}_{k\times n}\Big(b_{0,h},\, \Omega_{i,h}(\lambda_h)^{-1},\, \Sigma_{i,h}\Big), \tag{9}$$

Here, $\Omega_{i,h}(\lambda_h) \in \mathbb{R}^{k\times k}$ determines how strongly individual coefficients are shrunk, while $\Sigma_{i,h} \in \mathbb{R}^{n\times n}$ is the same working covariance matrix introduced in the Gaussian working likelihood. Its role here is again one of scaling across outcomes, so that the amount of shrinkage is comparable across variables measured in different units or exhibiting different residual variability. This is distinct from the Minnesota-specific scale normalization: $\Sigma_{i,h}$ rescales the prior across outcome equations, whereas $s_{i,h,j}$ governs how strongly individual lag coefficients are shrunk within a given equation. Let $s_{i,h,j} > 0$ denote a unit-, horizon-, and variable-specific scale-normalization term, which is

treated as fixed.[5] In the full specification, this Minnesota block is combined with the hierarchical pooling prior, so the complete prior for $B_{i,h}$ reflects both within-unit shrinkage of lag dynamics and cross-unit pooling of shock responses.

$$\mathrm{Var}\left(B_{i,h,j}^{(\ell)}\right) \;\propto\; \frac{\lambda_h^2}{\ell^{\,2}\, s_{i,h,j}^2}, \qquad \ell = 1,\dots,L, \qquad j = 1,\dots,n, \tag{10}$$

or, equivalently, the corresponding prior precision satisfies

$$\kappa_{h,\ell,j} \;\propto\; \frac{\ell^{\,2}\, s_{i,h,j}^2}{\lambda_h^2}. \tag{11}$$

This structure imposes stronger shrinkage for higher-order lags, uses $s_{i,h,j}$ to normalize Minnesota shrinkage across variables and horizons, and lets λ_h determine the overall strength of lag shrinkage at horizon h.

Horizon-Specific Shrinkage.
The strength of Minnesota-type shrinkage is allowed to vary with the horizon. This is useful because longer-horizon local projections are typically estimated less precisely, because forecast errors become noisier and the effective sample size declines with h. Allowing λ_h to vary across horizons therefore makes it possible to apply stronger shrinkage where the signal is weaker, while keeping shrinkage comparatively mild at short horizons. Each λ_h is treated as unknown and is assigned a Gamma prior,

$$\lambda_h \sim \Gamma(k_\lambda, \theta_\lambda) \qquad h = 0,\dots,H, \tag{12}$$

with shape $k_\lambda > 0$ and scale $\theta_\lambda > 0$. To impose positivity conveniently, we work with the log-tightness parameter $z_h = \log \lambda_h$. The sequence $\{z_h\}_{h=0}^H$ is then regularized through a Gaussian random walk prior,

$$z_h - z_{h-1} \sim \mathcal{N}(0, \sigma_z^2), \qquad h = 1,\dots,H, \tag{13}$$

together with a weakly informative prior for the initial level, $z_0 \sim \mathcal{N}(m_z, V_z)$. The innovation variance σ_z^2 governs how smoothly shrinkage strength is allowed to evolve across horizons and is assigned an inverse-Gamma hyperprior,

$$\sigma_z^2 \sim \text{Inv-Gamma}(a_z, b_z). \tag{14}$$

This specification regularizes the path of shrinkage parameters $\{\lambda_h\}_{h=0}^H$ rather than the impulse responses themselves. As a result, the data determine both the overall amount of lag shrinkage and how smoothly that shrinkage changes across horizons, without imposing a fixed functional form. Without the random-walk prior in (13), the shrinkage parameters λ_h would be treated

[5] Details are provided in Appendix A.2.3.

independently across horizons under the Gamma prior in (12). The random-walk specification links adjacent horizons and thereby smooths the evolution of shrinkage strength, while preserving the interpretation of λ_h as the horizon-specific tightness parameter governing lag shrinkage.

3.3.2 Cross-Unit Pooling Prior.

In addition to the Minnesota-type shrinkage on lag dynamics, the BPLP framework links units through a cross-unit pooling prior on the shock response coefficients. This is the key device that allows unit-level impulse responses to remain heterogeneous while still being connected through a common cross-sectional structure. The prior is applied only to the coefficients on the identified shock, because these are the impulse responses of interest. By contrast, the remaining local projection coefficients remain unit-specific, so dynamic propagation can differ flexibly across units.

For each horizon h, the unit-specific shock response coefficients satisfy

$$\alpha_{i,h} \mid \mu_h, \tau_i^2, \Sigma_{i,h} \sim \mathcal{N}_n\Big(\mu_h,\ \tau_i^2 \Sigma_{i,h}\Big), \qquad i = 1, \ldots, N,\ h = 0, \ldots, H, \tag{15}$$

where $\mu_h \in \mathbb{R}^n$ denotes the horizon-specific latent benchmark response. It is an unobserved benchmark estimated within the hierarchical model and serves as the reference point for cross-unit pooling.[6][7]

The matrix $\Sigma_{i,h}$ is again the same working covariance matrix introduced in the Gaussian working likelihood. In the pooling prior, it determines how deviations from the latent benchmark are scaled across outcomes. If $\Sigma_{i,h}$ is diagonal, pooling is applied outcome by outcome. If it is unrestricted, pooling additionally reflects contemporaneous comovement across outcomes. In either case, $\Sigma_{i,h}$ serves only as a scaling device, so that the degree of pooling is not mechanically driven by measurement units or by outcomes with larger residual variability. The heterogeneity parameters are assigned inverse-Gamma hyperpriors,

$$\tau_i^2 \sim \text{Inv-Gamma}(a_\tau, b_\tau), \qquad i = 1, \ldots, N. \tag{16}$$

Conditional on $\{\alpha_{i,h}\}_{i=1}^N$ and $\{\tau_i^2\}_{i=1}^N$, the cross-unit pooling prior implies a Gaussian updating equation for the latent benchmark at horizon h. Define

$$\Lambda_h = \sum_{i=1}^N \tau_i^{-2} \Sigma_{i,h}^{-1}, \qquad R_h = \Lambda_h^{-1}, \qquad y_h^\star = R_h \left(\sum_{i=1}^N \tau_i^{-2} \Sigma_{i,h}^{-1} \alpha_{i,h} \right). \tag{17}$$

[6] Here, latent means unobserved and estimated within the hierarchical model. It does not refer to a latent factor in the sense of factor models or principal components. When horizon smoothing is imposed, the sequence $\{\mu_h\}_{h=0}^H$ is additionally treated as a latent state path across horizons.

[7] This latent benchmark should be distinguished from the reported panel benchmark discussed in subsection 3.4, which is constructed separately as a post hoc aggregation of regularized unit-level responses.

Then the pooling block can be written as

$$y_h^\star \mid \mu_h \sim \mathcal{N}_n(\mu_h, R_h). \tag{18}$$

This shows that the latent benchmark is informed by a precision-weighted combination of the unit-level responses, where units with smaller τ_i^2 receive more weight. In the scalar case, or when the $\Sigma_{i,h}$ are proportional across units, these weights reduce to being proportional to τ_i^{-2}.

If no smoothing is imposed on the sequence $\{\mu_h\}_{h=0}^H$, the conditional mean is determined horizon by horizon and simplifies to

$$\mu_h = y_h^\star = \left(\sum_{i=1}^N \tau_i^{-2}\Sigma_{i,h}^{-1}\right)^{-1}\left(\sum_{i=1}^N \tau_i^{-2}\Sigma_{i,h}^{-1}\alpha_{i,h}\right). \tag{19}$$

If horizon smoothing is imposed, however, the latent benchmark path $\{\mu_h\}_{h=0}^H$ is updated jointly across horizons by the state-space smoother introduced below, with (17) providing the horizon-specific Gaussian measurement equation. Without smoothing, the update reduces to horizon-by-horizon Gaussian conditioning based on (17)[8]

Horizon Smoothness.
Impulse responses are often expected to vary smoothly with the horizon, reflecting gradual adjustment rather than erratic horizon-by-horizon movements. The BPLP framework allows such smoothness at two distinct levels. First, the latent benchmark path $\{\mu_h\}_{h=0}^H$ can be smoothed across horizons. Second, unit-specific shock-response paths $\{\alpha_{i,h}\}_{h=0}^H$ can, optionally, also be smoothed directly. These two layers serve different purposes: benchmark smoothing regularizes the common cross-sectional component, whereas unit-level smoothing regularizes the time profile of individual responses.

Latent Benchmark Path.
When smoothing is imposed, the latent benchmark path $\{\mu_h\}_{h=0}^H$ follows a first-order Gaussian random walk,

$$\mu_h \mid \mu_{h-1}, \sigma_\mu^2, \Sigma_\mu \sim \mathcal{N}_n\big(\mu_{h-1},\, \sigma_\mu^2\,\Sigma_\mu\big), \qquad h = 1, \ldots, H, \tag{20}$$

with a diffuse prior for μ_0. This prior penalizes abrupt changes in the latent benchmark across adjacent horizons, while still allowing the overall response profile to be learned from the data. The matrix Σ_μ scales horizon-to-horizon innovations across outcomes. With Σ_μ being diagonal, smoothness is imposed outcome by outcome, whereas a full Σ_μ also allows innovations to co-move across outcomes. The overall degree of smoothness is governed by

$$\sigma_\mu^2 \sim \text{Inv-Gamma}(a_\mu, b_\mu). \tag{21}$$

[8] Forward-filtering backward-sampling is only needed in the smoothed case.

Unit-level Shock Response Paths.
In addition to smoothing the latent benchmark path, the BPLP framework can also smooth unit-specific shock responses directly. When smoothing for the unit-level shock response paths is used, each path $\{\alpha_{i,h}\}_{h=0}^{H}$ follows

$$\alpha_{i,h} \mid \alpha_{i,h-1}, \xi^2, \Sigma_\alpha \sim \mathcal{N}_n\big(\alpha_{i,h-1},\, \xi^2\, \Sigma_\alpha\big), \qquad i = 1, \ldots, N, \;\; h = 1, \ldots, H, \tag{22}$$

where Σ_α plays the same scaling role as before. When Σ_α is diagonal smoothing is applied outcome by outcome, while an unrestricted Σ_α allows for co-movement across outcomes. The corresponding smoothness parameter is assigned

$$\xi^2 \sim \text{Inv-Gamma}(a_\xi, b_\xi). \tag{23}$$

Smaller values of ξ^2 imply smoother unit-level paths, whereas larger values allow more local variation across horizons. When this option is not used, unit-level shock responses remain unrestricted across horizons apart from the cross-unit pooling prior.

3.3.3 Joint prior structure.

The prior for the coefficient matrices $\{B_{i,h}\}$ combines the two regularization components introduced above: Minnesota-type shrinkage for lag dynamics and a cross-unit pooling prior for the coefficients on the identified shock. Because these components act on different parts of the coefficient vector, they combine additively in precision form. This yields a joint Gaussian prior that separates within-unit stabilization of lag dynamics from cross-unit pooling of impulse responses.

Let $\Omega_{i,h}(\lambda_h)$ denote the Minnesota-type precision matrix at horizon h. It is diagonal, implements lag decay and scale normalization, and is governed overall by the horizon-specific tightness parameter λ_h. By construction, this block shrinks lag coefficients, while deterministic terms are treated diffusely and the coefficients on the identified shock are excluded from Minnesota-type shrinkage. Let $T_i(\tau_i^2)$ denote the precision contribution from the cross-unit pooling prior for unit i. It acts only on the pooled coefficients, contributing precision proportional to τ_i^{-2} around the latent benchmark response μ_h.

The resulting joint prior is

$$B_{i,h} \mid \mu_h, \tau_i^2, \lambda_h, \Sigma_{i,h} \sim \mathcal{MN}_{k\times n}\Big(m_{i,h},\, P_{i,h}^{-1},\, \Sigma_{i,h}\Big), \tag{24}$$

with joint precision

$$P_{i,h} = \Omega_{i,h}(\lambda_h) + T_i(\tau_i^2), \tag{25}$$

and corresponding prior mean

$$m_{i,h} = P_{i,h}^{-1}\big(\Omega_{i,h}(\lambda_h)b_{0,h} + T_i(\tau_i^2)\mu_h\big). \tag{26}$$

This representation makes the roles of the hyperparameters transparent: λ_h governs the strength of shrinkage on lag dynamics, while τ_i^2 governs the strength of cross-unit pooling for the shock responses.

Quasi-Posterior and Posterior Simulation.
Combining the Gaussian working likelihood with the joint Gaussian prior yields a conditionally Gaussian quasi-posterior for $B_{i,h}$ and the hierarchical parameters. Conditional on $(\mu_h, \tau_i^2, \lambda_h)$ and the fixed working covariance matrix $\Sigma_{i,h}$, the coefficient matrix $B_{i,h}$ is updated from a matrix-normal full conditional with precision matrix

$$X_{i,h}'X_{i,h} + \Omega_{i,h}(\lambda_h) + T_i(\tau_i^2). \tag{27}$$

The remaining quantities are then updated from their corresponding conditional distributions. The pooling parameters τ_i^2 are updated from the cross-sectional dispersion of unit-level shock responses around the latent benchmark, the horizon-specific shrinkage parameters λ_h are updated through the log-tightness representation under the random-walk prior in (13), and, when benchmark smoothing is imposed, the latent benchmark path $\{\mu_h\}_{h=0}^H$ is updated jointly across horizons using forward-filtering backward-sampling under equations (20)–(21). When unit-level horizon smoothness is imposed, the paths $\{\alpha_{i,h}\}_{h=0}^H$ are updated analogously under equations (22)–(23).[9]

Summary.
Taken together, the algorithm produces MCMC draws for the unit-level impulse responses and hierarchical parameters,

$$\{\alpha_{i,h}^{(d)}, \tau_i^{2,(d)}, \lambda_h^{(d)}, \mu_h^{(d)}, \sigma_\mu^{2,(d)}\}, \tag{28}$$

and, when unit-level horizon smoothness is imposed, also $\{\xi^{2,(d)}\}$. These draws deliver the regularized quasi-posterior center and the associated hierarchical latent states. Reported uncertainty for the resulting impulse responses is then calibrated separately using the robust procedures described next.

3.4 Impulse Responses and Robust Inference.

This subsection defines the impulse responses. Throughout, point estimates are taken from the regularized center implied by the Gaussian working likelihood together with the Minnesota-type and cross-unit pooling priors. Uncertainty, however, is reported using pointwise misspecification-robust artificial-posterior methods, so that interval estimation is not tied to the misspecified Gaussian error

[9] Detailed derivations of these conditional updates are given in the technical appendix A.

structure.[10] Accordingly, the reported intervals should be interpreted as pointwise misspecification-robust frequentist intervals around a regularized center, rather than as exact Bayesian credible sets.

(1) Unit-level Impulse Response.
For unit i, horizon h, and outcome variable j, the impulse response is the coefficient on the identified shock in the corresponding horizon-h local projection. Let $\alpha_{i,h,j}^{(d)}$ denote the d-th retained draw of this coefficient, obtained from the corresponding draw of $B_{i,h}$. The reported unit-level point estimate is the regularized quasi-posterior center

$$\widehat{\alpha}_{i,h,j} = \mathbb{E}[\alpha_{i,h,j} \mid \cdot]\,. \tag{29}$$

The draws $\{\alpha_{i,h,j}^{(d)}\}$ are useful for summarizing the quasi-posterior fit and for constructing the reported regularized center, but the reported uncertainty bands are not taken directly from their dispersion. Instead, unit-level inference is based on a heteroskedasticity- and autocorrelation-consistent long-run variance estimator that captures the serial dependence induced by overlapping horizons.

(2) Reported Panel Benchmark Response.
It is useful to distinguish between the latent benchmark μ_h, which enters the cross-unit pooling prior inside the hierarchical model, and the *reported* panel benchmark response, which is constructed post hoc as a fixed-weight aggregation of regularized unit-level impulse responses, where the weights are given by posterior means of draw-wise normalized precision weights. Let $G = \{1, \ldots, N\}$ denote the full set of units. For each draw d, horizon h, and outcome variable j, define the normalized precision weights

$$w_i^{(d)} = \frac{\omega_i^{(d)}}{\sum_{m=1}^{N} \omega_m^{(d)}}, \qquad \omega_i^{(d)} = \frac{1}{\tau_i^{2,(d)}}.$$

Using the posterior mean of these draw-wise normalized weights,

$$\bar{w}_i := \mathbb{E}\Big[w_i^{(d)} \mid \text{data}\Big]\,,$$

the reported panel benchmark point estimate is

$$\widehat{\mu}_{h,j}^{\text{panel}} = \sum_{i=1}^{N} \bar{w}_i\, \widehat{\alpha}_{i,h,j}. \tag{30}$$

Units with smaller pooling variances receive larger weight, so the reported panel benchmark places more mass on responses that are estimated to be more coherent with the pooled component of

[10] For unit-level impulse responses, the robust calibration follows the HAC-based artificial-posterior logic used in Ferreira et al. (2025). For aggregated objects such as subgroup and panel benchmarks, Appendix A.10, when included, provides an application-specific theorem establishing the asymptotic validity of pointwise DK-style long-run variance calibration under the maintained asymptotic assumptions.

the model. Uncertainty for this reported panel benchmark is calibrated post hoc by combining the unit-level influence contributions with the same weights used in the benchmark itself and then applying a robust HAC/DK-type long-run variance estimator to the resulting aggregate series. The resulting intervals are pointwise misspecification-robust intervals around this reported regularized benchmark.

(3) Reported Leave-One-Out Response.
To compare unit i with the rest of the panel, we construct a leave-one-out benchmark that excludes that unit. Let $G_{-i} = \{1, \dots, N\} \setminus \{i\}$ denote the set of remaining units. For each draw d, define the leave-one-out normalized precision weights

$$w_{m,(-i)}^{(d)} = \frac{\omega_m^{(d)}}{\sum_{\ell \neq i} \omega_\ell^{(d)}}, \qquad m \neq i, \qquad \omega_m^{(d)} = \frac{1}{\tau_m^{2,(d)}}.$$

Using the posterior mean of these draw-wise normalized weights,

$$\bar{w}_{m,(-i)} := \mathbb{E}\Big[w_{m,(-i)}^{(d)} \mid \text{data}\Big], \qquad m \neq i,$$

the reported leave-one-out benchmark is defined as

$$\widehat{\mu}_{h,j}^{(-i)} = \sum_{m \neq i} \bar{w}_{m,(-i)}\, \widehat{\alpha}_{m,h,j}. \tag{31}$$

This benchmark summarizes the response profile implied by the remaining units only. Comparing the reported unit-level response to the corresponding leave-one-out benchmark therefore isolates whether unit i is broadly aligned with the pooled cross-sectional pattern or exhibits systematic deviations from it. Uncertainty for the leave-one-out benchmark is calibrated post hoc by combining the unit-level influence contributions of the remaining units with the same leave-one-out weights used in the benchmark itself and then applying a robust HAC/DK-type long-run variance estimator to the resulting aggregate series.

(4) Reported Group-specific Response.
The same aggregation logic can be applied to any subset of units. For a group $G \subset \{1, \dots, N\}$, define the draw-wise normalized subgroup precision weights

$$w_{i,G}^{(d)} = \frac{\omega_i^{(d)}}{\sum_{m \in G} \omega_m^{(d)}}, \qquad i \in G, \qquad \omega_i^{(d)} = \frac{1}{\tau_i^{2,(d)}}.$$

Using the posterior mean of these draw-wise normalized subgroup weights,

$$\bar{w}_{i,G} := \mathbb{E}\Big[w_{i,G}^{(d)} \mid \text{data}\Big], \qquad i \in G,$$

the reported group-specific response is defined as

$$\widehat{\mu}_{G,h,j} = \sum_{i \in G} \bar{w}_{i,G} \, \widehat{\alpha}_{i,h,j}. \tag{32}$$

This construction yields group-specific responses that remain fully consistent with the cross-unit pooling logic of the model. As for the panel benchmark, uncertainty for these aggregated responses is calibrated post hoc by combining the unit-level influence contributions with the same subgroup weights used in the benchmark itself and then applying a robust HAC/DK-type long-run variance estimator to the resulting aggregate series.

4 Monetary Policy Transmission in the Euro Area

This section illustrates the BPLP framework introduced in Section 3 by examining monetary policy transmission across eight euro area economies. The analysis combines a balanced monthly panel for Austria, France, Germany, the Netherlands, Italy, Spain, Portugal, and Greece with high-frequency identified monetary policy surprises. The empirical objective is twofold: (i) to characterize the common component of monetary transmission and (ii) to quantify systematic cross-country heterogeneity. Reflecting the ECB's primary mandate, the discussion focuses on inflation, while results for industrial production and unemployment are used mainly to frame real-side adjustment costs.[11]

The data consist of monthly industrial production, the harmonised index of consumer prices (HICP), and unemployment rates from January 2001 to April 2025. Industrial production and HICP are expressed as year-on-year percentage changes, which smooth high-frequency volatility and enhance comparability across countries. The unemployment rate enters in levels, consistent with standard macroeconomic specifications. Monetary policy shocks are taken from the official euro-area high-frequency shock series associated with Jarociński and Karadi (2020), which is available for the sample used here from January 2001 to April 2025 and uses surprises around ECB policy announcements together with sign restrictions to isolate monetary policy shocks from central bank information shocks.[12] Because euro area monetary policy is set centrally, the same identified shock series is assigned to all countries and normalized to a one-standard-deviation innovation.

Estimation proceeds by running Bayesian Panel Local Projections up to 24 horizons. For each country and horizon, the local projection includes an intercept, the common monetary policy shock, and 12 lags of the outcome variables. Using 12 lags is standard for monthly macroeconomic data and allows the specification to capture annual dynamics and residual persistence in a transparent way. Given these transformations, the Minnesota prior mean for lag dynamics is set to zero in the baseline specification. Cross-country heterogeneity is modeled through partial pooling of the shock coefficients so that country-specific impulse responses are shrunk toward a horizon-specific panel benchmark path, while intercepts and lag dynamics remain country-specific. The strength of pooling is governed by the unit-specific heterogeneity parameters τ_i^2, so countries with more coherent shock responses are pooled more tightly. Computation is based on a misspecified Gaussian working likelihood with a diagonal scale matrix $\Sigma_{i,h}$, constructed prior to MCMC from preliminary OLS residuals and treated as fixed thereafter. Smoothness across horizons is introduced via random-walk priors for the benchmark path and for the horizon-specific tightness parameter λ_h. In addition, in the baseline specification each country's shock response path is smoothed across horizons. Inference

[11] Results for industrial production and unemployment are reported in Appendix B.

[12] High-frequency identification for monetary policy shocks in the euro area builds on asset-price movements in narrow windows around ECB announcements; see Altavilla et al. (2019). Jarociński and Karadi (2020) discuss the role of information effects in ECB announcement windows, and Ramey (2016) surveys identification strategies. In the empirical application, these shocks are treated as externally identified monetary policy innovations for the local-projection specification.

is reported using pointwise misspecification-robust intervals: unit-level uncertainty is calibrated by Newey–West long-run variance estimates that account for the serial dependence induced by overlapping horizons, and aggregate (panel or group) responses use a DK-style long-run variance estimator applied to the corresponding precision-weighted aggregated influence process. Markov chain Monte Carlo sampling is conducted with 20,000 draws, discarding the first 10,000 as burn-in. This choice is guided by the goal of obtaining stable posterior summaries and satisfactory convergence diagnostics for the main hierarchical parameters.

The remainder of this section is organized as follows. We first describe the precision-weighted panel response for inflation, which summarizes the common component of monetary transmission across the eight euro area countries. We then turn to country-level inflation dynamics to identify how atypical each country's response is relative to its peers and evaluate regional Core–Periphery patterns. The final subsection discusses the implications of these findings for monetary policy and communication in a heterogeneous currency union.

4.1 The Common Component of Euro Area Inflation

The precision-weighted panel impulse response provides a summary of the common cross-country component of monetary transmission by aggregating the country-level responses into a single benchmark path. The panel response at each horizon is computed as a precision-weighted average, where weights are proportional to $1/\tau_i^2$, the inverse of the unit-specific pooling variances.[13] This scheme places greater weight on countries whose shock responses are estimated to align more tightly with the evolving cross-sectional pattern, and correspondingly downweights economies whose responses are more idiosyncratic.

Figure 1 summarizes in Panel A the Bayesian Panel Local Projection inflation response, whereas Panel B shows the corresponding panel Local Projection benchmark estimated without Bayesian regularization or horizon smoothing. The panel LP benchmark indicates some evidence that inflation reacts to a contractionary monetary policy shock in the short- and medium-term horizons. Its point estimates and Driscoll–Kraay bands are visibly more irregular across horizons, which is consistent with the absence of the shrinkage and smoothing mechanisms built into the BPLP specification.

[13] The reported panel benchmark is not intended as a GDP-weighted euro-area aggregate. Rather, it is a model-based precision-weighted summary that places greater weight on responses estimated to be more coherent with the pooled component of the hierarchical specification.

Figure 1: Panel Response of HICP inflation.

Panel A: Bayesian Panel Local Projections.

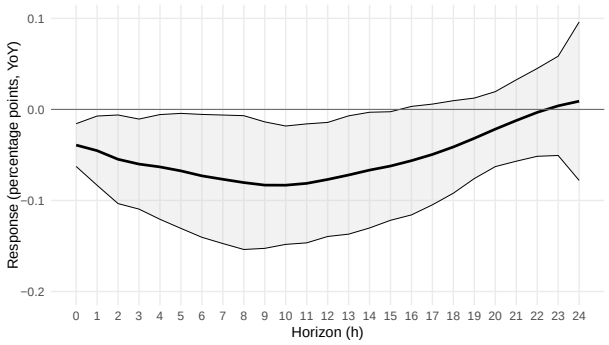

Precision-weighted panel response of HICP inflation after a one-standard-deviation monetary policy shock. Shown are the regularized panel point estimates and 68% pointwise misspecification-robust intervals based on Driscoll–Kraay-style long-run variance calibration.

Panel B: Panel Local Projections.

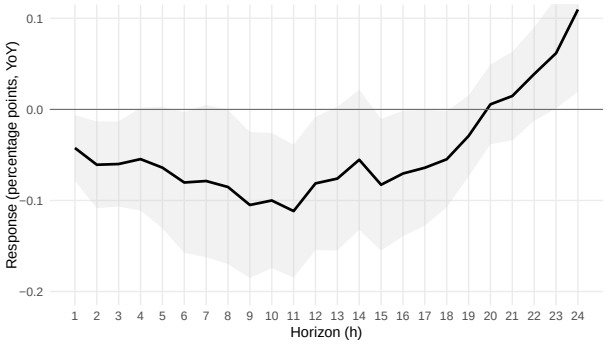

Panel Local Projection response of HICP inflation after a one-standard-deviation monetary policy shock. Shown are unsmoothed OLS point estimates with 68% pointwise Driscoll–Kraay confidence intervals.

In contrast, the Bayesian Panel Local Projections deliver a more consistently negative regularized response over short-to-medium horizons together with narrower pointwise misspecification-robust intervals. A natural interpretation is that the Bayesian specification stabilizes estimation through prior regularization and cross-sectional pooling, while uncertainty is calibrated separately through misspecification-robust long-run variance estimators. By comparison, Panel B deliberately leaves these regularizing mechanisms absent, so its more jagged profile provides a useful visual benchmark for the stabilization delivered by the BPLP framework.

In what follows, the focus is on the BPLP benchmark response. The panel mean is negative already at short horizons and becomes more pronounced over the first year. The response declines from a small initial drop to a trough of roughly -0.1 percentage points around months nine to eleven. The pointwise misspecification-robust intervals (shaded area) lie entirely below zero over a broad short- to mid-horizon window, indicating evidence of a decline in inflation over that range.

After the trough, the response steadily attenuates, approaching zero around months $21 - 22$. At the longest horizons the point estimate turns slightly positive, but the uncertainty bands widen and overlap zero, so the data do not support a precise statement about the sign of the response beyond roughly 18 months. The smooth profile is consistent with the model's regularization: Minnesota-type shrinkage disciplines lag dynamics within each horizon, while random-walk smoothing across horizons dampens erratic month-to-month movements without imposing a parametric shape on the impulse response.

Figure 2 provides a direct view of the pooling mechanism. Panel A summarizes posterior dispersion in $\log(\tau_i^2)$ across countries, and Panel B reports the corresponding posterior mean precision weights (normalized to sum to one) used to construct the panel benchmark.

Figure 2: Heterogeneity Measures.

Panel A: Posterior Pooling Variances

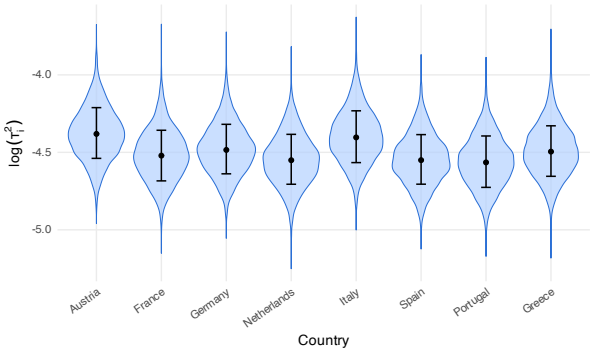

Posterior densities of the country-specific pooling variances τ_i^2 (log scale). Smaller values indicate stronger shrinkage toward the common component, meaning that the corresponding country's inflation response is estimated to be more coherent with the cross-sectional pattern captured by the panel benchmark.

Panel B: Posterior Precision Weights

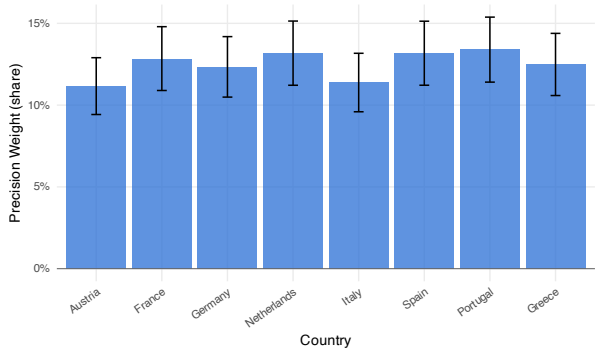

Posterior mean precision weights used in constructing the common inflation response. Weights are normalized to sum to one and are proportional to the inverse pooling variances $1/\tau_i^2$. Countries with responses estimated to be more coherent with the pooled component receive larger weights.

Consequently, Panel B yields precision weights that are fairly even, with each economy contributing on the order of one eighth of the benchmark and no single country dominating the aggregate response. This matters for interpretation: the panel benchmark is not dominated by any single member state in the precision-weighted aggregation, while remaining slightly more anchored in those units whose responses are estimated to be more coherent within the model's pooling structure.

Taken together, the evidence indicates that inflation declines following a contractionary monetary policy shock across the euro area sample, although the strength and timing of this adjustment may vary meaningfully across economies. These differences motivate the more granular country-level and regional analysis in the next subsections.

4.2 Country-Specific Inflation Responses

To assess cross-country heterogeneity, we compare each country's inflation response to a leave-one-out (LOO) benchmark constructed from the remaining seven countries in the full sample. This ensures that the reference path is not mechanically influenced by the country under consideration. Panel A plots the regularized country response together with the corresponding LOO benchmark, each accompanied by 68% pointwise misspecification-robust intervals. Panel B plots the difference (country minus LOO benchmark): values above zero indicate that the country's response lies above the benchmark at that horizon, whereas values below zero indicate a response below the benchmark. Deviations are interpreted as pointwise evidence of differential responses when the 68% pointwise interval of the difference excludes zero.

Among the core economies (Austria, France, Germany, and the Netherlands), France and Germany track their respective LOO benchmarks closely. In both cases, the country-specific regularized point estimate follows the same broad shape as the benchmark across horizons, and the corresponding difference plots show that the 68% pointwise intervals overlap zero throughout, indicating no systematic departure from the cross-country reference path. Austria, in contrast, exhibits a pronounced deviation from its LOO benchmark: while the benchmark implies a negative inflation response over the medium horizons, Austria's regularized point estimate lies above that reference, and the difference plot is positive for essentially all horizons. Moreover, the gap for Austria is sizeable over the medium-to-long horizons, and the 68% pointwise interval of the difference excludes zero over a wide range, indicating that Austria's inflation response lies pointwise above what is implied by the other seven countries over that range.

Figure 3: Core Countries' HICP Inflation Response.

Panel A: Country's Response relative to the Leave-One-Out Benchmark.

Panel B: Difference between Country and Leave-One-Out Benchmark.

Austria

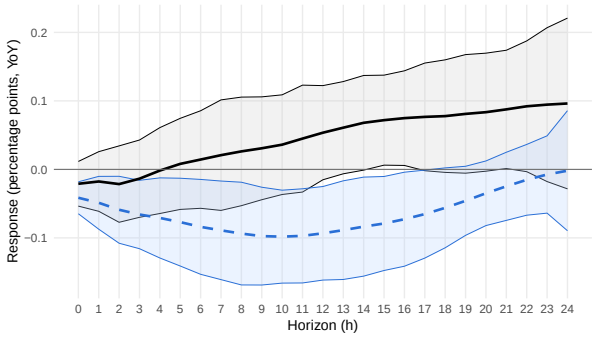

(a) *Austria's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

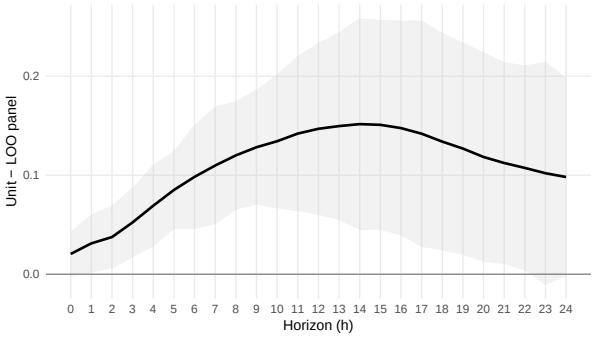

(b) *Difference between Austria's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

France

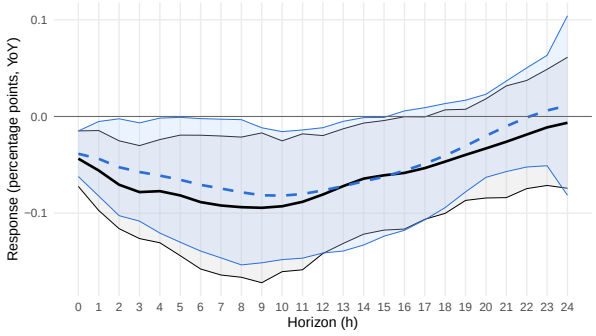

(c) *France's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

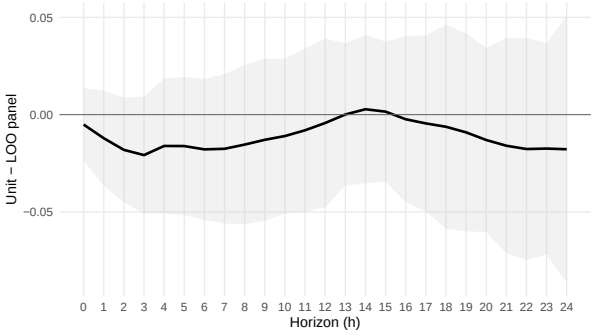

(d) *Difference between France's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

Figure 3: Core Countries' HICP Inflation Response. (continued).

Panel A: Country's Response relative to the Panel.

Panel B: Difference Plot between Country and Panel.

Germany

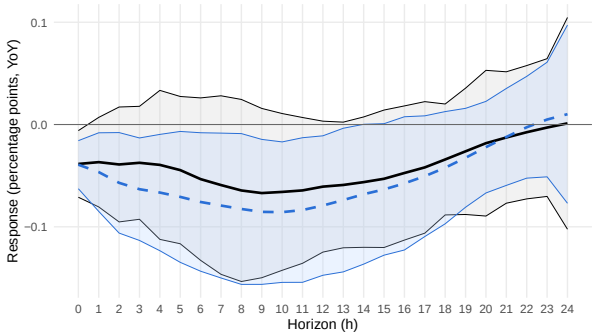

(e) *Germany's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

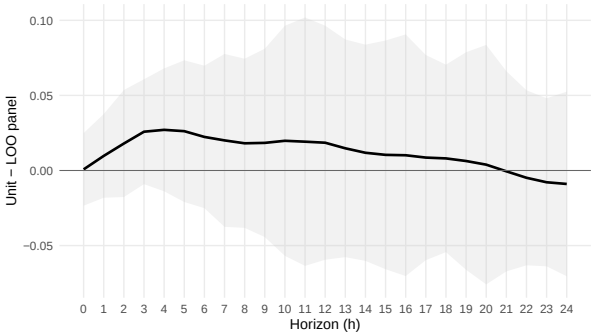

(f) *Difference between Germany's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

Netherlands

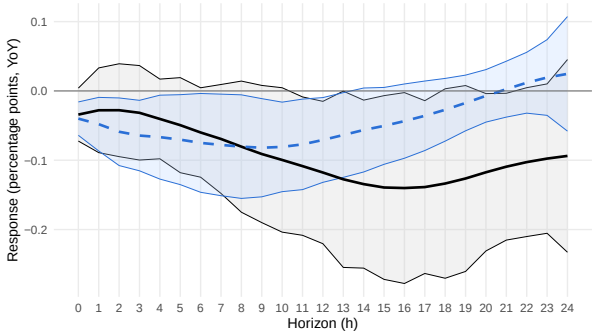

(g) *The Netherlands' inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

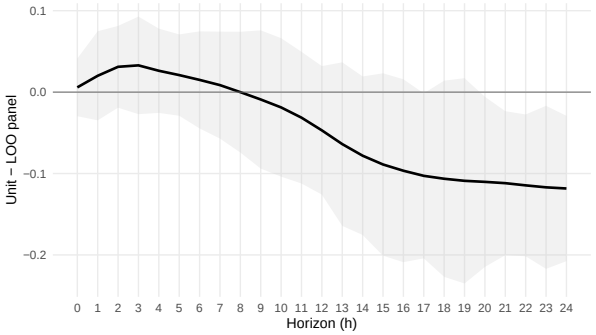

(h) *Difference between the Netherlands' inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

The Netherlands display a different form of heterogeneity. The difference plot is mildly positive at short horizons but turns negative after roughly the first year and becomes strongly negative at longer horizons. This implies that the Dutch response starts out slightly above the LOO benchmark but then falls below it later on, with a persistent medium-to-long horizon gap.

Taken together, the core group contains two countries that are close to the benchmark (France and Germany) and two that deviate in economically distinct ways: Austria is consistently above its peers' benchmark, whereas the Netherlands transitions from a small short-run positive gap to a pronounced negative gap at longer horizons.

Country-Specific Inflation Responses: Periphery

Figure 4 repeats the leave-one-out (LOO) comparison for Italy, Spain, Portugal, and Greece. In Panel A, the solid line is the country-specific regularized response and the dashed line is the corresponding LOO benchmark (the precision-weighted panel response recomputed without that country). Panel B reports the gap between the two (country minus LOO benchmark). Negative values indicate a more negative inflation response than implied by peers, while positive values indicate a less negative response (or a relative increase). Statistical relevance is assessed by whether the 68% pointwise interval in Panel B excludes zero.

Figure 4: Periphery Countries' HICP Inflation Response.

Panel A: Country's Response relative to the Panel. | *Panel B: Difference Plot between Country and Panel.*

Italy

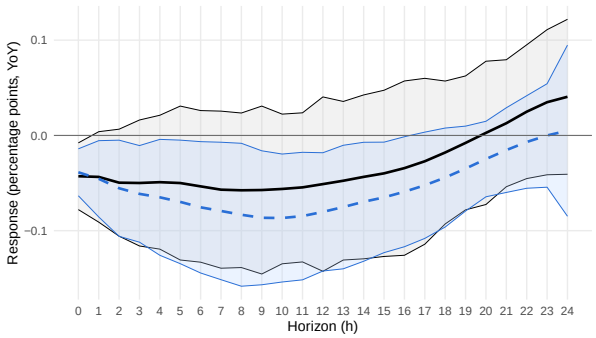

(a) Italy's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.

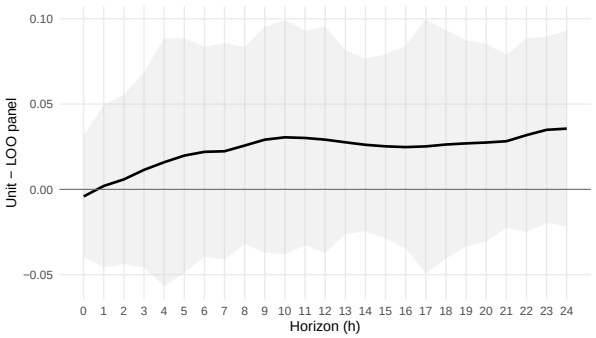

(b) Difference between Italy's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.

Spain

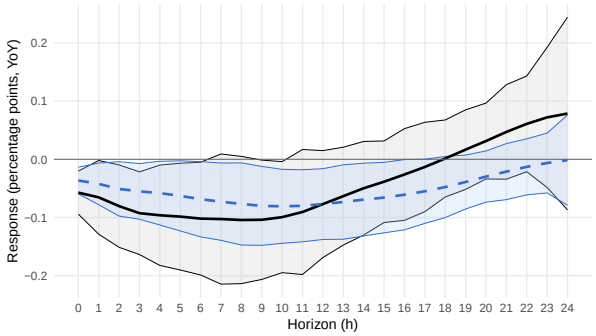

(c) *Spain's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

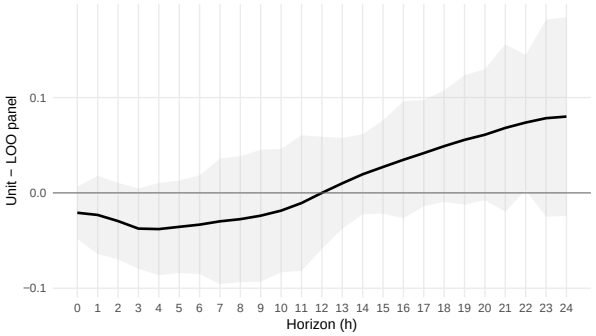

(d) *Difference between Spain's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

Figure 4: Periphery Countries' HICP Inflation Response. (continued).

Panel A: Country's Response relative to the Panel.

Panel B: Difference Plot between Country and Panel.

Portugal

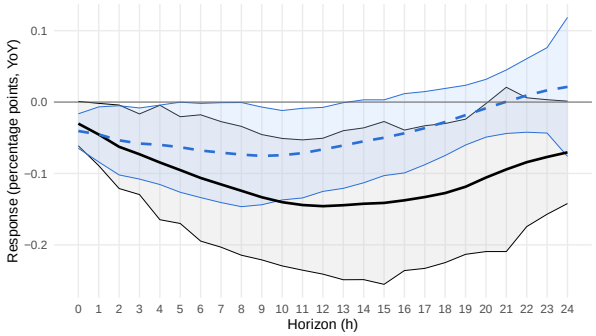

(e) *Portugal's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

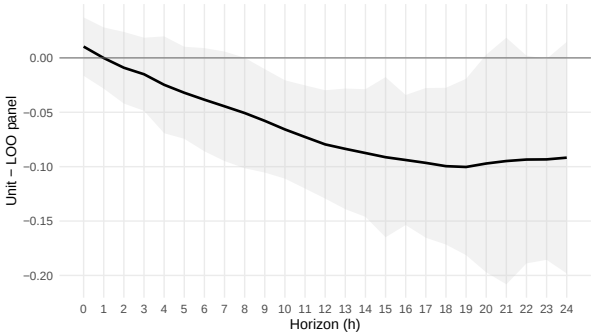

(f) *Difference between Portugal's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

Greece

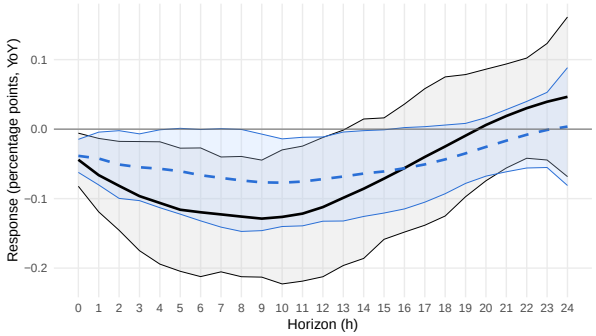

(g) *Greece's inflation response (black) relative to its leave-one-out benchmark (blue). Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals.*

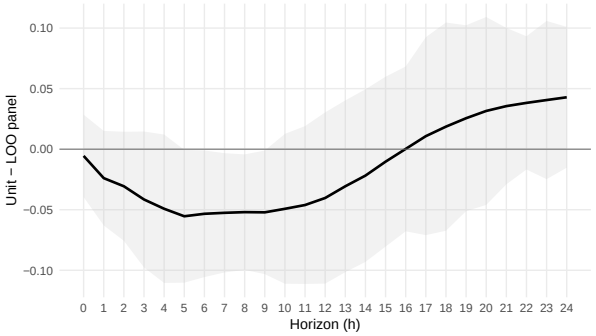

(h) *Difference between Greece's inflation response and its leave-one-out benchmark. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

Two patterns stand out. First, Portugal appears as the most pronounced outlier within the periphery. Its response falls sharply and remains substantially below the LOO benchmark throughout the horizons. The difference plot declines steadily and stays negative at medium and long horizons. Relative to the rest of the sample, Portugal is estimated to exhibit a larger and more persistent decline in inflation rather than a short-lived deviation concentrated at a particular horizon.

Second, Spain and Greece share a common shape relative to their LOO benchmarks. Both start below the benchmark at short horizons (negative differences), but the gap narrows over time and turns positive at longer horizons. Concretely, Spain's inflation drops more than the LOO benchmark in the first year, after which the country-specific path catches up and eventually moves above the benchmark toward the end of the horizon. Greece shows the same qualitative pattern, with a

clearer early shortfall relative to the benchmark and a gradual turn toward positive differences later on. These profiles suggest that these countries are not uniformly associated with a larger inflation decline than the panel, instead, they load more on the early part of the adjustment and less on the late part.

Italy is the closest to the cross-sectional pattern in the sense that its country path is broadly aligned with the LOO benchmark in Panel A. The difference plot is modest and mostly positive, indicating that Italy's response tends to be slightly less negative than what is implied by the rest of the sample, especially around medium horizons, but without the pronounced and persistent gap observed for Portugal.

Taken together, the periphery does not move as a single block relative to the panel benchmark. Portugal contributes a distinctly larger and longer-lasting decline in inflation, whereas Spain and Greece deviate mainly through a re-timing of the response, i.e., stronger early movements paired with comparatively weaker (and eventually higher) outcomes at longer horizons. Italy, in contrast, remains comparatively close to its peers. These country-level patterns motivate the subsequent aggregation into Core and Periphery responses, which clarifies how much of the regional difference is driven by a persistent level shift (Portugal) versus differences in timing (Spain and Greece).

Regional Patterns: Core versus Periphery

Figure 5 aggregates the country-level inflation responses into two groups, the Core (Austria, France, Germany, the Netherlands) and the Periphery (Italy, Spain, Portugal, Greece), using the same precision-weighting logic as for the panel benchmark. Panel A plots the two regional impulse responses with 68% pointwise misspecification-robust intervals. Both regions display the expected sign following a contractionary monetary policy shock, i.e., inflation moves below zero for an extended stretch of horizons. The Core response is comparatively modest: after a small negative movement at short horizons, it reaches its most negative values around the medium horizons and then steadily drifts back toward zero. The pointwise intervals remain relatively tight throughout, consistent with Core-country responses being more mutually coherent.

The Periphery response is qualitatively similar but quantitatively stronger. Inflation falls more sharply and reaches a noticeably deeper trough around the same medium horizons, with a magnitude roughly about twice as large as the Core at the low point. Uncertainty is also visibly higher in the Periphery, reflected in wider pointwise intervals, indicating greater within-group heterogeneity even after pooling. Toward the end of the horizon, both regional responses approach small values and the uncertainty bands widen, so differences become harder to pin down precisely.

Figure 5: Core vs Periphery Response of HICP inflation.

Panel A: Core vs Periphery

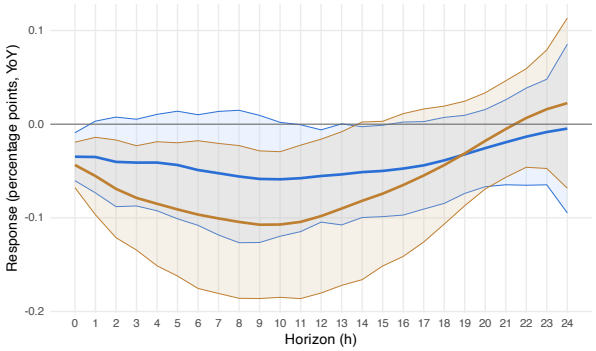

(a) *Responses of the Core (Austria, France, Germany, the Netherlands) and the Periphery (Italy, Spain, Portugal, Greece) for HICP inflation after a one-standard-deviation monetary policy shock. The Core response is shown in blue and the Periphery response in brown. Shown are the regularized group point estimates together with 68% pointwise misspecification-robust intervals.*

Panel B: Difference Plot

(b) *Difference between the Core and Periphery responses of HICP inflation after a one-standard-deviation monetary policy shock. Shown is the regularized difference together with the 68% pointwise misspecification-robust interval.*

Panel B reports the difference (Core minus Periphery). The difference curve lies predominantly above zero over the horizons where the inflation decline is largest, suggesting that the Periphery experiences a more negative inflation response than the Core in the medium run. The largest gaps occur around the trough region, aligning with the visual separation between the two paths in Panel A. Overall, the regional comparison shows that cross-country heterogeneity documented in the leave-one-out figures is not only idiosyncratic at the country level, but also broadly consistent with a noticeable Core–Periphery pattern.

Implications for Monetary Policy and Communication

The preceding subsection summarizes how the panel benchmark, country-specific deviations, and regional patterns jointly characterize monetary transmission across the eight euro area economies. These results have several implications for the interpretation and communication of monetary policy in a heterogeneous currency union.

First, the inflation response of the panel benchmark provides a useful summary measure of the euro-area-wide effect of a monetary policy tightening. The estimated path is persistent, suggesting that the common component of monetary transmission is well behaved and that the short- to medium-term inflation trajectory is reasonably stable in the data. This benchmark can serve as a reference point for communicating the expected effects of policy decisions, while acknowledging that the underlying country-level responses are not identical.

Second, the cross-country and regional differences imply that a policy-rate change generates unequal timing and, in some cases, unequal intensity in the decline of inflation across member states. Economies with faster transmission experience an earlier fall in inflation and may perceive the stance as comparatively tighter, while those with slower transmission face a more gradual adjustment. The posterior pooling variances and precision weights provide diagnostics for these asymmetries: countries with larger pooling variances contribute less to the panel benchmark because their responses are more idiosyncratic or more uncertain, while countries with tightly estimated profiles anchor the benchmark. These diagnostics offer a disciplined way to relate national perspectives to the aggregate inflation outlook.

Third, the results bear on the interpretation of forward-looking communication. Forward guidance aims to coordinate expectations about the future path of monetary policy and inflation, but heterogeneous transmission suggests that the same guidance may be interpreted differently across member states. While the analysis does not estimate expectations or guidance effects directly, it clarifies the empirical environment in which they operate: some countries exhibit relatively predictable and tightly estimated inflation responses, whereas others face more uncertain and more dispersed adjustment paths.

Finally, the results suggest that the common component and the group-level responses are directionally aligned in this sample, even though the timing and magnitude vary across countries. The hierarchical structure makes this transparent: a single monetary policy shock can generate a shared negative inflation response while allowing for meaningful variation in timing and persistence. Robustness checks, when included in Appendix B, vary the degree of cross-country pooling, the tightness of lag shrinkage, and related hyperparameters, and are intended to assess whether the main heterogeneity patterns are driven by a particular calibration.

In sum, the BPLP framework provides a transparent decomposition between common and country-specific components of monetary policy transmission in the eight euro area member countries. The precision-weighted panel benchmark and the leave-one-out diagnostics offer interpretable summaries of the cross-sectional pattern and its deviations, helping to characterize where transmission is cohesive and where it differs in timing or magnitude.

5 Conclusion

This paper proposes a Bayesian Panel Local Projections framework for estimating the transmission of monetary policy in a heterogeneous currency union. The econometric contribution is twofold: the model combines country-specific local projections with hierarchical pooling that adapts endogenously to the cross-sectional structure of the data, and it embeds Minnesota-type lag shrinkage with horizon-smoothing priors that stabilize medium-term IRFs without imposing excessive structure. The resulting hierarchical posterior objects, in particular the pooling variances and precision weights, provide a transparent decomposition of heterogeneity and common dynamics, while uncertainty for reported impulse responses is calibrated separately using misspecification-robust methods.

The empirical application to an eight-country euro area sample illustrates that the framework can deliver coherent and economically plausible impulse responses at both the aggregate and disaggregate levels. The estimates suggest a common component of monetary transmission: inflation responds negatively to a contractionary shock in a persistent manner, with pointwise misspecification-robust intervals that remain interpretable across horizons. At the same time, the model uncovers meaningful cross-country variation in both the timing and the magnitude of inflation. While the panel benchmark is not dominated by a single economy in the precision-weighted aggregation, the Periphery is estimated to exhibit earlier and more pronounced adjustment than the Core. The hierarchical prior structure is designed to help distinguish systematic heterogeneity from idiosyncratic noise, making the estimated differences more economically interpretable and less sensitive to sampling variability.

These findings illustrate that a single monetary policy can generate a unified directional response while still producing heterogeneous trajectories across member states. This has direct relevance for policy communication, interpretation of national perspectives within the Governing Council, and assessment of the distribution of adjustment costs. Robustness exercises, when included in the full version, are intended to assess whether the qualitative patterns remain stable across reasonable prior choices.

More broadly, the BPLP framework offers a flexible and computationally efficient tool for studying policy transmission in panels where heterogeneity is expected but should not be imposed mechanically. Future work may extend the framework to allow for time-varying pooling or to incorporate alternative shock identification schemes. Nonetheless, the results presented here highlight that carefully designed Bayesian structure can reconcile common monetary dynamics with meaningful cross-country variation, providing a richer and more transparent depiction of policy transmission in a heterogeneous monetary union.

References

Afonso, A., Alves, J., & Ionta, S. (2025). Monetary policy surprise shocks under different fiscal regimes: A panel analysis of the euro area. *Journal of International Money and Finance*, *156*, 103341. https://doi.org/10.1016/j.jimonfin.2025.103341

Almgren, M., Gallegos, J.-E., Kramer, J., & Lima, R. (2022). Monetary policy and liquidity constraints: Evidence from the euro area. *American Economic Journal: Macroeconomics*, *14*(4), 309–340. https://doi.org/10.1257/mac.20200096

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., & Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, *108*, 162–179. https://doi.org/10.1016/j.jmoneco.2019.08.016

Bobasu, A., Ciccarelli, M., Berg, T. O., Brázdik, F., Buss, G., Conti, A. M., Covarrubias, M., Delis, P., Giammaria, A., Gonçalves, N. V., Grimaud, A., Kornprobst, A., Kortelainen, M., Lozej, M., Mandler, M., Müller, G., Nilavongse, R., Notarpietro, A., Odendahl, F., … Valderrama, M. T. (2025). Monetary policy transmission: A reference guide through escb models and empirical benchmarks. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5715582

Brugnolini, L., Catania, L., & Santucci de Magistris, P. (2023). Bayesian flexible local projections. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4669495

Calza, A., Monacelli, T., & Stracca, L. (2013). Housing finance and monetary policy. *Journal of the European Economic Association*, *11*, 101–122. https://doi.org/10.1111/j.1542-4774.2012.01095.x

Canova, F., & Ciccarelli, M. (2013). Panel vector autoregressive models: A survey. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2201610

Ciccarelli, M., & Rebucci, A. (2006). Has the transmission mechanism of european monetary policy changed in the run-up to emu? *European Economic Review*, *50*(3), 737–776. https://doi.org/10.1016/j.euroecorev.2005.01.001

Cima, S., & Moreno, M. (2025). Monetary policy shocks and their effects across the wealth distribution: Evidence from new european data. *The Journal of Economic Inequality*. https://doi.org/10.1007/s10888-025-09687-4

Datsenko, R., & Fleck, J. (2024). Country-specific effects of euro-area monetary policy: The role of sectoral differences. *FEDS Notes*, (2024-11-12-2), None–None. https://doi.org/10.17016/2380-7172.3649

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, *80*(4), 549–560. https://doi.org/10.1162/003465398557825

Ferreira, L. N., Miranda-Agrippino, S., & Ricco, G. (2025). Bayesian local projections. *Review of Economics and Statistics*, *107*(5), 1424–1438. https://doi.org/10.1162/rest{\textunderscore}a{\textunderscore}01334

Georgiadis, G. (2015). Examining asymmetries in the transmission of monetary policy in the euro area: Evidence from a mixed cross-section global var model. *European Economic Review*, *75*, 195–215. https://doi.org/10.1016/j.euroecorev.2014.12.007

Holm-Hadulla, F., & Thürwächter, C. (2021). Heterogeneity in corporate debt structures and the transmission of monetary policy. *European Economic Review*, *136*, 103743. https://doi.org/10.1016/j.euroecorev.2021.103743

Huber, F., Matthes, C., & Pfarrhofer, M. (2024). General seemingly unrelated local projections. https://doi.org/10.48550/ARXIV.2410.17105

Huber, P. J. (1967). The behavior of maximum likelihood estimates under nonstandard conditions. *Proceedings of the fifth Berkeley Symposium on Mathematical Statistics and Probability*, (1), 221–2333. Retrieved November 30, 2025, from https://projecteuclid.org/ebook/Download?urlid=bsmsp%2F1200512988&isFullBook=False

Jarociński, M., & Karadi, P. (2020). Deconstructing monetary policy surprises— the role of information shocks. *American Economic Journal: Macroeconomics*, *12*(2), 1–43. https://doi.org/10.1257/mac.20180090

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182. https://doi.org/10.1257/0002828053828518

Jordà, Ò. (2009). Simultaneous confidence regions for impulse responses. *Review of Economics and Statistics*, *91*(3), 629–647. https://doi.org/10.1162/rest.91.3.629

Lane, P. R. (2012). The european sovereign debt crisis. *Journal of Economic Perspectives*, *26*(3), 49–68. https://doi.org/10.1257/jep.26.3.49

Li, D., Plagborg-Møller, M., & Wolf, C. K. (2024). Local projections vs. vars: Lessons from thousands of dgps. *Journal of Econometrics*, *244*(2), 105722. https://doi.org/10.1016/j.jeconom.2024.105722

Mandler, M., Scharnagl, M., & Volz, U. (2022). Heterogeneity in euro area monetary policy transmission: Results from a large multicountry bvar model. *Journal of Money, Credit and Banking*, *54*(2-3), 627–649. https://doi.org/10.1111/jmcb.12859

Mojon, B., & Peersman, G. (2001). A var description of the effects of monetary policy in the individual countries of the euro area. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.303801

Montiel Olea, J. L., & Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, *89*(4), 1789–1823. https://doi.org/10.3982/ECTA18756

Müller, U. K. (2013). Risk of bayesian inference in misspecified models, and the sandwich covariance matrix. *Econometrica*, *81*(5), 1805–1849. https://doi.org/10.3982/ECTA9097

Newey, W. K., & McFadden, D. (n.d.). Chapter 36 large sample estimation and hypothesis testing. https://doi.org/10.1016/S1573-4412(05)80005-4

Olea, J. L. M., Plagborg-Møller, M., Qian, E., & Wolf, C. K. (2025). Local projections or vars? a primer for macroeconomists. https://doi.org/10.48550/ARXIV.2503.17144

Peersman, G. (2004). The transmission of monetary policy in the euro area: Are the effects different across countries?*. *Oxford Bulletin of Economics and Statistics*, *66*(3), 285–308. https://doi.org/10.1111/j.1468-0084.2004.00080.x

Peersman, G., & Smets, F. (2001). The monetary transmission mechanism in the euro area: More evidence from var analysis (mtn conference paper). *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.356269

Pica, S. (2021). Housing markets and the heterogeneous effects of monetary policy across the euro area. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4060424

Plagborg-Møller, M., & Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, *89*(2), 955–980. https://doi.org/10.3982/ECTA17813

Ramey, V. (2016). Macroeconomic shocks and their propagation. https://doi.org/10.3386/w21978

Slacalek, J., Tristani, O., & Violante, G. L. (2020). Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area. *Journal of Economic Dynamics and Control*, *115*, 103879. https://doi.org/10.1016/j.jedc.2020.103879

Tanaka, M. (2025). Quasi-bayesian local projections: Simultaneous inference and extension to the instrumental variable method. https://doi.org/10.48550/ARXIV.2503.20249

Wächter, M., Proano, C., & Peña, J. C. (2024). *How fitting is "one-size-fits-all"? revisiting the dynamic effects of ecb's interest policy on euro area countries* (Vol. no. 199 (November 2024)). Bamberg Economic Research Group, Bamberg University.

Wallis, K. F. (1977). Multiple time series analysis and the final form of econometric models. *Econometrica*, *45*(6), 1481. https://doi.org/10.2307/1912313

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, *50*(1), 1. https://doi.org/10.2307/1912526

Zellner, A., & Palm, F. (1974). Time series analysis and simultaneous equation econometric models. *Journal of Econometrics*, *2*(1), 17–54. https://doi.org/10.1016/0304-4076(74)90028-1

A Technical Appendix: Bayesian Panel Local Projections (BPLP)

This appendix provides the technical details for the Bayesian Panel Local Projections (BPLP) framework. It derives the main posterior updating blocks and documents the implemented regularized quasi-posterior machinery. The presentation is aligned with the empirical implementation, including the use of a fixed working covariance matrix and a post hoc robust inference step.

We consider a panel of N units and horizons $h \in \{0, 1, \dots, H\}$. Let L denote the number of lags in the local projection design. For each horizon h, define the effective sample size as

$$T_h = T - L - h.$$

For each unit i and horizon h, let

- $Y_{i,h} \in \mathbb{R}^{T_h \times n}$ denote the stacked outcomes,
- $X_{i,h} \in \mathbb{R}^{T_h \times k}$ denote the stacked regressors, and
- $B_{i,h} \in \mathbb{R}^{k \times n}$ denote the corresponding coefficient matrix.

The impulse responses of interest are given by the coefficients on the identified shock and may optionally be smoothed across horizons depending on the specification.

A.1 Local Projection Regression System

For each unit i and horizon h, the local projection system is

$$Y_{i,h} = X_{i,h} B_{i,h} + U_{i,h}, \qquad U_{i,h} \in \mathbb{R}^{T_h \times n}. \tag{A.1}$$

Gaussian Working Likelihood.
Conditional on $B_{i,h}$, the forecast errors are modeled using a Gaussian *working* likelihood with independence over time and a horizon-specific covariance matrix $\Sigma_{i,h}$ across equations:

$$U_{i,h} \mid \Sigma_{i,h} \sim \mathcal{MN}_{T_h \times n}(0,\, I_{T_h},\, \Sigma_{i,h}), \tag{A.2}$$

or, equivalently,

$$Y_{i,h} \mid B_{i,h}, \Sigma_{i,h} \sim \mathcal{MN}_{T_h \times n}\big(X_{i,h} B_{i,h},\, I_{T_h},\, \Sigma_{i,h}\big). \tag{A.3}$$

The matrix $\Sigma_{i,h}$ is treated as fixed during posterior simulation and is computed outside the Markov chain; see Section A.8. It serves as a scaling device in the quasi-posterior update and is not intended to represent the true covariance structure of the local-projection forecast errors. In particular, serial dependence induced by overlapping horizons is not modeled here but is handled separately through the misspecification-robust long-run variance calibration used for inference.

Under misspecification, the resulting quasi-posterior concentrates around pseudo-true parameter values defined as minimizers of the corresponding population least-squares criterion. The asymptotic justification for inference around these pseudo-true objects is provided in Appendix A.10.

Sufficient Statistics.
Define the cross-product matrices

$$S_{i,h}^{XX} = X_{i,h}' X_{i,h} \in \mathbb{R}^{k\times k}, \qquad S_{i,h}^{XY} = X_{i,h}' Y_{i,h} \in \mathbb{R}^{k\times n}. \tag{A.4}$$

Conditional on $\Sigma_{i,h}$ and the hyperparameters, the full conditional update for $B_{i,h}$ depends on the data only through these sufficient statistics.

A.2 Prior for $B_{i,h}$: Minnesota Shrinkage and Cross-Unit Pooling

A.2.1 Pooled Coefficient Block

Let $\mathcal{R} \subseteq \{1, \dots, k\}$ denote the set of regressor indices whose coefficients are pooled across units. In the empirical application, this set corresponds to the identified shock, although the framework allows additional pooling (e.g., for intercepts or trends). Let $R = |\mathcal{R}|$.

Write $B_{i,h}^{(\mathcal{R})} \in \mathbb{R}^{R\times n}$ for the submatrix of $B_{i,h}$ formed by the pooled rows, and let $\mu_h^{(\mathcal{R})} \in \mathbb{R}^{R\times n}$ denote the corresponding latent benchmark block.

A.2.2 Minnesota-type Prior Mean $b_{0,h}$

Let $b_{0,h} \in \mathbb{R}^{k\times n}$ denote the Minnesota-type prior mean at horizon h. In the current implementation, this object is horizon-invariant, so $b_{0,h} \equiv b_0$. The own first lag is centered according to the persistence properties of the data, while all remaining coefficients are centered at zero.

A.2.3 Minnesota Precision $\Omega_{i,h}(\lambda_h)$

For each unit i and horizon h, let $\Omega_{i,h}(\lambda_h) \in \mathbb{R}^{k\times k}$ denote a diagonal prior precision matrix,

$$\Omega_{i,h}(\lambda_h) = \operatorname{diag}(\omega_{i,h,1}, \dots, \omega_{i,h,k}).$$

For lag coefficients at lag $\ell \in \{1, \dots, L\}$ and variable $s \in \{1, \dots, n\}$, the precision takes the Minnesota form

$$\omega_{i,h,\ell,s} = \frac{\ell^2 \operatorname{scale}_{i,h,s}^2}{\lambda_h^2}, \tag{A.5}$$

where $\operatorname{scale}_{i,h,s} > 0$ is a fixed scale-normalization term. These scale factors are treated as fixed within the posterior simulation and can be constructed from preliminary residual-based estimates to normalize coefficients across variables and horizons.

Intercept, shock, and trend regressors are assigned diffuse Minnesota dispersion. In particular, for every pooled regressor index $r \in \mathcal{R}$,

$$\omega_{i,h,r} = 0, \qquad r \in \mathcal{R}. \tag{A.6}$$

Thus, the Minnesota block is flat on the pooled coefficients and, taken on its own, is generally improper when $\mathcal{R} \neq \varnothing$. Propriety is ensured by combining it with the cross-unit pooling block introduced below.

A.2.4 Pooling Precision

Each unit i has a heterogeneity parameter $\tau_i > 0$ with variance τ_i^2.

Case A (independent pooled coefficients).
If pooled coefficients are independent across the R pooled regressors, the pooling block uses an identity structure.

Case B (dependent pooled coefficients).
More generally, dependence across pooled coefficients can be introduced through a positive definite matrix $\Sigma_B \in \mathbb{R}^{R\times R}$ with inverse Σ_B^{-1}.

Both cases are covered by defining

$$A = \begin{cases} I_R, & \text{independent case,} \\ \Sigma_B^{-1}, & \text{dependent case.} \end{cases}$$

The pooling precision matrix $T_i(\tau_i^2) \in \mathbb{R}^{k\times k}$ is then defined by

$$T_i(\tau_i^2)\big|_{\mathcal{R},\mathcal{R}} = \frac{1}{\tau_i^2} A, \qquad T_i(\tau_i^2)\big|_{\mathcal{R}^c,\cdot} = 0. \tag{A.7}$$

A.2.5 Joint Prior Distribution of $B_{i,h}$

Define the total prior precision matrix as

$$P_{i,h}(\lambda_h, \tau_i^2) = \Omega_{i,h}(\lambda_h) + T_i(\tau_i^2). \tag{A.8}$$

The corresponding prior mean is the precision-weighted combination of the Minnesota prior mean and the pooled benchmark,

$$m_{i,h} = P_{i,h}^{-1}\Big(\Omega_{i,h}(\lambda_h) b_{0,h} + T_i(\tau_i^2)\mu_h\Big). \tag{A.9}$$

The resulting joint prior is matrix-normal,

$$B_{i,h} \mid \lambda_h, \tau_i^2, \mu_h, \Sigma_{i,h} \sim \mathcal{MN}_{k\times n}\Big(m_{i,h},\, P_{i,h}^{-1},\, \Sigma_{i,h}\Big). \tag{A.10}$$

Thus, Minnesota shrinkage and cross-unit pooling combine additively in precision form, while $\Sigma_{i,h}$ provides a common scaling across outcome equations.

A.3 Posterior of $B_{i,h}$

Because both the Gaussian working likelihood and the joint prior for $B_{i,h}$ are matrix-normal, the conditional posterior of $B_{i,h}$ is available in closed form.

A.3.1 Derivation by Completing the Square

To simplify notation, suppress the (i,h) subscripts and write $\Sigma = \Sigma_{i,h}$, $X = X_{i,h}$, $Y = Y_{i,h}$, $P = P_{i,h}$, and $m = m_{i,h}$. The derivation follows by collecting the terms in B from the working likelihood and the prior, and then completing the square.

The likelihood kernel in B is

$$\log p(Y \mid B, \Sigma) \propto -\frac{1}{2}\mathrm{tr}\Big(\Sigma^{-1}(Y - XB)'(Y - XB)\Big). \tag{A.11}$$

Expanding the quadratic form gives

$$\begin{aligned} (Y - XB)'(Y - XB) &= Y'Y - Y'XB - B'X'Y + B'X'XB, \\ \Rightarrow \quad \mathrm{tr}\Big(\Sigma^{-1}(Y - XB)'(Y - XB)\Big) &= \mathrm{tr}(\Sigma^{-1}Y'Y) - 2\,\mathrm{tr}\Big(\Sigma^{-1}B'X'Y\Big) + \mathrm{tr}\Big(\Sigma^{-1}B'X'XB\Big). \end{aligned} \tag{A.12}$$

The prior kernel in B is

$$\log p(B \mid \lambda_h, \tau_i^2, \mu_h, \Sigma) \propto -\frac{1}{2}\mathrm{tr}\Big(\Sigma^{-1}(B - m)'P(B - m)\Big). \tag{A.13}$$

Expanding this term gives

$$\begin{aligned} (B - m)'P(B - m) &= B'PB - B'Pm - m'PB + m'Pm, \\ \Rightarrow \quad \mathrm{tr}\Big(\Sigma^{-1}(B - m)'P(B - m)\Big) &= \mathrm{tr}\Big(\Sigma^{-1}B'PB\Big) - 2\,\mathrm{tr}\Big(\Sigma^{-1}B'Pm\Big) + \mathrm{tr}\Big(\Sigma^{-1}m'Pm\Big). \end{aligned} \tag{A.14}$$

Combining (A.12) and (A.14), and dropping terms that do not depend on B, the quadratic term in B is

$$\operatorname{tr}\Big(\Sigma^{-1}B'(X'X+P)B\Big),$$

while the linear term is

$$-2\operatorname{tr}\Big(\Sigma^{-1}B'(X'Y+Pm)\Big).$$

It follows that the posterior is matrix-normal with posterior precision

$$Q^B_{i,h} = X'_{i,h}X_{i,h} + P_{i,h}(\lambda_h, \tau_i^2), \tag{A.15}$$

and posterior mean

$$\bar{B}_{i,h} = (Q^B_{i,h})^{-1}\Big(X'_{i,h}Y_{i,h} + P_{i,h}m_{i,h}\Big). \tag{A.16}$$

Using (A.9), note that

$$P_{i,h}m_{i,h} = \Omega_{i,h}b_{0,h} + T_i(\tau_i^2)\mu_h,$$

so the posterior mean can equivalently be written as

$$\boxed{\bar{B}_{i,h} = (Q^B_{i,h})^{-1}\Big(X'_{i,h}Y_{i,h} + \Omega_{i,h}(\lambda_h)b_{0,h} + T_i(\tau_i^2)\mu_h\Big).} \tag{A.17}$$

A.3.2 Full conditional

Therefore, the full conditional distribution of $B_{i,h}$ is

$$\boxed{B_{i,h} \mid Y_{i,h}, X_{i,h}, \lambda_h, \tau_i^2, \mu_h, \Sigma_{i,h} \sim \mathcal{MN}_{k\times n}\Big(\bar{B}_{i,h},\, (Q^B_{i,h})^{-1},\, \Sigma_{i,h}\Big).} \tag{A.18}$$

A.4 Hierarchical Pooling: Priors and Posteriors for τ_i^2 and Σ_B

A.4.1 Prior for τ_i^2

The unit-specific pooling variance is assigned an inverse-gamma prior,

$$\tau_i^2 \sim \text{Inv-Gamma}(a_\tau, b_\tau). \tag{A.19}$$

A.4.2 Pooling Block as a Function of τ_i^2

Let

$$E_{i,h} = B^{(\mathcal{R})}_{i,h} - \mu^{(\mathcal{R})}_h \in \mathbb{R}^{R\times n}$$

denote the deviation of the pooled coefficient block from the latent benchmark at horizon h. The pooling prior for the pooled submatrix can be written as

$$B^{(\mathcal{R})}_{i,h} \mid \mu^{(\mathcal{R})}_h, \tau_i^2, \Sigma_B, \Sigma_{i,h} \sim \mathcal{MN}_{R\times n}\Big(\mu^{(\mathcal{R})}_h,\, \tau_i^2\Sigma_B,\, \Sigma_{i,h}\Big), \tag{A.20}$$

with $\Sigma_B = I_R$ in the case where pooled coefficients are independent across rows.

Conditional on $\{\mu_h^{(\mathcal{R})}\}_{h=0}^H$, the contribution of this pooling block to the kernel in τ_i^2 is, up to constants that do not depend on τ_i^2,

$$p(\{B_{i,h}^{(\mathcal{R})}\}_h \mid \{\mu_h^{(\mathcal{R})}\}_h, \tau_i^2, \cdot) \propto (\tau_i^2)^{-\frac{Rn(H+1)}{2}} \exp\left\{ -\frac{1}{2\tau_i^2} \sum_{h=0}^{H} \mathrm{tr}\Big(\Sigma_{i,h}^{-1} E_{i,h}' \Sigma_B^{-1} E_{i,h} \Big) \right\}. \tag{A.21}$$

If $\Sigma_B = I_R$, the factor Σ_B^{-1} drops out. In the empirical application, the pooled block is one-dimensional ($R = 1$), so this expression reduces to the scalar pooling case.

A.4.3 Full Conditional for τ_i^2

Combining (A.19) and (A.21) yields the conjugate full conditional

$$\boxed{\tau_i^2 \mid \{B_{i,h}\}_h, \{\mu_h\}_h, \{\Sigma_{i,h}\}_h, \Sigma_B \sim \text{Inv-Gamma}\Big(a_\tau + \frac{Rn(H+1)}{2},\ b_\tau + \frac{1}{2} \sum_{h=0}^{H} \mathrm{tr}\big(\Sigma_{i,h}^{-1} E_{i,h}' \Sigma_B^{-1} E_{i,h} \big) \Big).} \tag{A.22}$$

A.4.4 Optional Prior and Full Conditional for Σ_B

If the optional covariance matrix Σ_B is used, it is assigned an inverse-Wishart prior,

$$\Sigma_B \sim \text{Inv-Wishart}(\nu_{0,B}, S_{0,B}). \tag{A.23}$$

Given $\{B_{i,h}\}$ and $\{\mu_h\}$, define the pooled deviations

$$D_{i,h} = B_{i,h}^{(\mathcal{R})} - \mu_h^{(\mathcal{R})}.$$

Under (A.20), the corresponding conjugate full conditional is

$$\boxed{\Sigma_B \mid \{B_{i,h}\}_{i,h}, \{\mu_h\}_h, \{\tau_i^2\}_i, \{\Sigma_{i,h}\}_{i,h} \sim \text{Inv-Wishart}\Big(\nu_{0,B} + nN(H+1),\ S_{0,B} + \sum_{h=0}^{H} \sum_{i=1}^{N} \frac{1}{\tau_i^2} D_{i,h} \Sigma_{i,h}^{-1} D_{i,h}' \Big).} \tag{A.24}$$

This block is included for completeness. It is not used in the empirical application, which works with the simpler specification that does not introduce cross-row dependence within the pooled coefficient block.

A.5 Latent Benchmark Path $\{\mu_h\}$: Smoothing and State-Space Update

This subsection describes the update of the latent benchmark path $\{\mu_h\}_{h=0}^{H}$ that enters the cross-unit pooling prior. This latent benchmark should be distinguished from the reported panel benchmark in the main text, which is constructed separately as a post hoc aggregation of regularized unit-level responses. When no smoothing across horizons is imposed, the latent benchmark is updated horizon by horizon. The present subsection gives the corresponding state-space representation for the smoothed case. Without benchmark smoothing, the update reduces to horizon-by-horizon Gaussian conditioning based on the measurement equation (A.30), so FFBS is only needed in the smoothed case with state equation (A.25).

A.5.1 State-Space Representation

Stack the pooled coefficients into a vector

$$\theta_h = \mathrm{vec}\big(\mu_h^{(\mathcal{R})}\big) \in \mathbb{R}^d, \qquad d = Rn.$$

Let $m_0 \in \mathbb{R}^d$ and $C_0 \in \mathbb{R}^{d\times d}$ denote the prior mean and covariance of θ_0. In the default prior, $C_0 = (1/V_0^{-1})I_d$ if $V_0^{-1} > 0$ and is diffuse otherwise.

RW1 State Equation.
The smoothed latent benchmark follows a first-order Gaussian random walk,

$$\theta_0 \sim \mathcal{N}(m_0, C_0), \qquad \theta_h = \theta_{h-1} + u_h, \quad u_h \sim \mathcal{N}(0, W_\mu), \quad h = 1, \dots, H, \tag{A.25}$$

with innovation covariance

$$\boxed{W_\mu = \sigma_\mu^2 \left(\Sigma_\mu \otimes I_R\right).} \tag{A.26}$$

Here $\Sigma_\mu \in \mathbb{R}^{n\times n}$ is a fixed correlation-normalized covariance matrix constructed from the working covariance matrices; see Section A.8. It satisfies $\mathrm{diag}(\Sigma_\mu) = \mathbf{1}$ and determines how innovations in the latent benchmark co-move across outcomes. The scalar σ_μ^2 governs the overall degree of smoothness.

Prior for σ_μ^2.
The smoothing parameter σ_μ^2 is assigned an inverse-gamma prior

$$\sigma_\mu^2 \sim \text{Inv-Gamma}(a_\mu, b_\mu). \tag{A.27}$$

A.5.2 Measurement Equation Implied by the Pooling Prior

At each horizon h, let

$$b_{i,h} = \mathrm{vec}\big(B_{i,h}^{(\mathcal{R})}\big) \in \mathbb{R}^d$$

denote the pooled coefficient block for unit i in vectorized form. The pooling prior implies that, conditional on θ_h and τ_i^2, this block is Gaussian with precision

$$\Lambda_{i,h} = \frac{1}{\tau_i^2}\big(\Sigma_{i,h}^{-1} \otimes A\big), \tag{A.28}$$

where $A = I_R$ in the independent case and $A = \Sigma_B^{-1}$ in the dependent case.

Because the cross-sectional pooling contributions are Gaussian and conditionally independent across units, they aggregate into a single Gaussian measurement equation at horizon h:

$$\Lambda_h = \sum_{i=1}^{N} \Lambda_{i,h}, \qquad R_h = \Lambda_h^{-1}, \qquad y_h = R_h \left(\sum_{i=1}^{N} \Lambda_{i,h} b_{i,h} \right), \tag{A.29}$$

so that

$$\boxed{y_h \mid \theta_h \sim \mathcal{N}(\theta_h, R_h).} \tag{A.30}$$

Thus, at each horizon, the cross section reduces to a single Gaussian measurement pair (y_h, R_h) for the latent benchmark state.

A.5.3 FFBS Update for $\{\theta_h\}_{h=0}^{H}$

Given (A.25) and (A.30), the smoothed latent benchmark path is updated in a linear Gaussian state-space system. Sampling proceeds by forward-filtering backward-sampling (FFBS).

Forward filter.
Initialize

$$a_0 = m_0, \qquad P_0 = C_0.$$

At horizon $h = 0$, update using (y_0, R_0):

$$S_0 = P_0 + R_0, \tag{A.31}$$
$$K_0 = P_0 S_0^{-1}, \tag{A.32}$$
$$m_0^\star = a_0 + K_0(y_0 - a_0), \tag{A.33}$$
$$C_0^\star = (I - K_0)P_0. \tag{A.34}$$

For $h = 1, \ldots, H$, first predict:

$$a_h = m^\star_{h-1}, \qquad P_h = C^\star_{h-1} + W_\mu, \tag{A.35}$$

and then update:

$$S_h = P_h + R_h, \tag{A.36}$$
$$K_h = P_h S_h^{-1}, \tag{A.37}$$
$$m^\star_h = a_h + K_h(y_h - a_h), \tag{A.38}$$
$$C^\star_h = (I - K_h)P_h. \tag{A.39}$$

Backward sampling. Draw $\theta_H \sim \mathcal{N}(m^\star_H, C^\star_H)$. For $h = H-1, \dots, 0$ define

$$J_h = C^\star_h P^{-1}_{h+1}. \tag{A.40}$$

Then

$$\boxed{\theta_h \mid \theta_{h+1}, y_{0:H} \sim \mathcal{N}\Big(m^\star_h + J_h(\theta_{h+1} - a_{h+1}),\ C^\star_h - J_h P_{h+1} J_h'\Big).} \tag{A.41}$$

Finally, reshape θ_h back into $\mu_h^{(\mathcal{R})} \in \mathbb{R}^{R\times n}$ and embed it in $\mu_h \in \mathbb{R}^{k\times n}$.

A.5.4 Full Conditional for σ^2_μ

Under (A.26), the increments $\Delta_h = \theta_h - \theta_{h-1}$ satisfy

$$\Delta_h \sim \mathcal{N}\Big(0,\ \sigma^2_\mu(\Sigma_\mu \otimes I_R)\Big).$$

The conjugate update therefore depends on the quadratic form

$$\mathrm{SS}_\mu = \sum_{h=1}^{H} \Delta_h' \big(\Sigma_\mu^{-1} \otimes I_R\big)\Delta_h. \tag{A.42}$$

Let

$$\mathrm{df}_\mu = HRn.$$

Then

$$\boxed{\sigma_\mu^2 \mid \{\mu_h\} \sim \text{Inv-Gamma}\Big(a_\mu + \frac{\mathrm{df}_\mu}{2},\ b_\mu + \frac{\mathrm{SS}_\mu}{2}\Big).} \tag{A.43}$$

A.6 Horizon-specific Tightness λ_h

The Minnesota tightness parameter is allowed to vary across horizons. To impose positivity, write

$$z_h = \log \lambda_h.$$

The prior for λ_h is Gamma with shape k_λ and scale θ_λ:

$$\lambda_h \sim \text{Gamma}(k_\lambda, \theta_\lambda). \tag{A.44}$$

Optionally, smoothness across horizons is introduced through an RW1 prior on z_h:

$$z_0 \sim \mathcal{N}(m_z, V_z), \qquad z_h \mid z_{h-1}, \sigma_z^2 \sim \mathcal{N}(z_{h-1}, \sigma_z^2), \quad h = 1, \dots, H, \tag{A.45}$$

with

$$\sigma_z^2 \sim \text{Inv-Gamma}(a_z, b_z). \tag{A.46}$$

A.6.1 Conditional Target for z_h

The conditional update of z_h combines four ingredients: the prior density of $B_{i,h}$ through its dependence on λ_h via $\Omega_{i,h}(\lambda_h)$, the Gamma prior for λ_h, the RW1 prior for z_h, and the Jacobian term implied by the transformation $\lambda_h = e^{z_h}$.

Fix h and write $\lambda = e^{z_h}$. For each unit i, define

$$P_{i,h}(\lambda) = \Omega_{i,h}(\lambda) + T_i(\tau_i^2), \qquad m_{i,h}(\lambda) = P_{i,h}(\lambda)^{-1}\Big(\Omega_{i,h}(\lambda) b_{0,h} + T_i(\tau_i^2)\mu_h\Big),$$

and

$$D_{i,h}(\lambda) = B_{i,h} - m_{i,h}(\lambda).$$

From the matrix-normal prior (A.10), the part of the log-kernel that depends on λ is

$$\log p(B_{i,h} \mid \lambda, \tau_i^2, \mu_h, \Sigma_{i,h}) \doteq \frac{n}{2} \log |P_{i,h}(\lambda)| - \frac{1}{2}\mathrm{tr}\Big(\Sigma_{i,h}^{-1} D_{i,h}(\lambda)' P_{i,h}(\lambda) D_{i,h}(\lambda)\Big). \tag{A.47}$$

In the z_h parametrization, the Gamma prior contributes

$$\log p(\lambda) = (k_\lambda - 1)z_h - \frac{e^{z_h}}{\theta_\lambda} + \text{const},$$

and the transformation from λ_h to z_h contributes the Jacobian term $+z_h$. The RW1 prior adds the corresponding quadratic term in neighboring values of z_h, with the boundary case at $h = 0$ determined by $(z_0 - m_z)^2/V_z$.

Collecting these components, the conditional target for z_h is

$$\begin{aligned}\log p(z_h \mid \cdot) \doteq \sum_{i=1}^{N} \left[\frac{n}{2}\log|P_{i,h}(e^{z_h})| - \frac{1}{2}\mathrm{tr}\Big(\Sigma_{i,h}^{-1} D_{i,h}(e^{z_h})' P_{i,h}(e^{z_h}) D_{i,h}(e^{z_h})\Big)\right] \\ + (k_\lambda - 1)z_h - \frac{e^{z_h}}{\theta_\lambda} + z_h \;+\; \log p(z_h \mid z_{-h}, m_z, V_z, \sigma_z^2), \end{aligned} \tag{A.48}$$

where $\log p(z_h \mid z_{-h}, \cdot)$ denotes the RW1 conditional prior term implied by (A.45). Because (A.48) is non-conjugate, z_h is updated by a one-dimensional MCMC step, such as slice sampling or random-walk Metropolis.

A.6.2 Full Conditional for σ_z^2

Let $\Delta z_h = z_h - z_{h-1}$. Under (A.45) and (A.46), the conditional distribution of σ_z^2 is

$$\boxed{\sigma_z^2 \mid \{z_h\} \sim \text{Inv-Gamma}\Big(a_z + \frac{H}{2},\; b_z + \frac{1}{2}\sum_{h=1}^{H}(\Delta z_h)^2\Big).} \tag{A.49}$$

A.7 Optional Smoothing of Unit-Level Shock Responses across Horizons

An optional specification imposes smoothness directly on the unit-level shock response path across horizons. This block is distinct from the smoothing of the latent benchmark path discussed above: here the object being smoothed is the unit-specific shock coefficient itself.

Let $\beta_{i,h} \in \mathbb{R}^n$ denote the shock coefficient vector for unit i at horizon h, i.e. the unit-level impulse response vector. The smoothing prior is

$$\beta_{i,h} = \beta_{i,h-1} + \eta_{i,h}, \qquad \eta_{i,h} \sim \mathcal{N}\big(0,\; \xi^2 \Sigma_\alpha\big), \qquad h = 1, \ldots, H, \tag{A.50}$$

where Σ_α is a fixed positive definite matrix constructed from the working covariance matrices $\Sigma_{i,h}$ and $\xi^2 > 0$ is the smoothing parameter. In the implementation, the initial state is treated diffusely. In the univariate version, Σ_α is replaced by $\mathrm{diag}(\Sigma_\alpha)$, so smoothing is applied outcome by outcome.

A.7.1 Conditional State-Space Update

The smoothed shock path is updated conditional on the remaining coefficient rows. Let $r_\star$ denote the row index of the shock regressor in $B_{i,h}$ and define

$$\beta_{i,h} = B_{i,h}(r_\star, \cdot)' \in \mathbb{R}^n.$$

Partition the coefficient matrix as

$$B_{i,h} = \begin{bmatrix} \beta_{i,h}' \\ B_{i,h}^{(-\star)} \end{bmatrix},$$

where $B_{i,h}^{(-\star)} \in \mathbb{R}^{(k-1)\times n}$ collects all non-shock rows.

The state-space update is obtained by isolating the working-likelihood contribution of the shock row and, when shock pooling is active, combining it with the hierarchical pooling prior. To do so, partition the sample cross-product matrix $S_{i,h}^{XX} = X_{i,h}'X_{i,h}$ and the cross moment $S_{i,h}^{XY} = X_{i,h}'Y_{i,h}$ conformably as

$$S_{i,h}^{XX} = \begin{bmatrix} s_{\star\star,i,h} & s_{\star o,i,h} \\ s_{o\star,i,h} & S_{oo,i,h}^{XX} \end{bmatrix}, \qquad S_{i,h}^{XY} = \begin{bmatrix} s_{\star y,i,h} \\ S_{oy,i,h}^{XY} \end{bmatrix}.$$

Given the current non-shock rows $B_{i,h}^{(-\star)}$, define the implied working-likelihood pseudo-observation

$$\widehat{y}_{i,h} = \frac{1}{s_{\star\star,i,h}} \Big(s_{\star y,i,h} - s_{\star o,i,h} B_{i,h}^{(-\star)} \Big) \in \mathbb{R}^{1\times n}, \tag{A.51}$$

whenever $s_{\star\star,i,h} > 0$. This is the conditional Gaussian measurement for the shock row implied by the Gaussian working likelihood.

In the multivariate version, the corresponding working measurement equation is

$$\widehat{y}_{i,h}' \mid \beta_{i,h} \sim \mathcal{N}\Big(\beta_{i,h},\ V_{i,h}\Big), \qquad V_{i,h} = \frac{1}{s_{\star\star,i,h}} \Sigma_{i,h}. \tag{A.52}$$

In the univariate version, only the diagonal elements are retained:

$$\widehat{y}_{i,h,j} \mid \beta_{i,h,j} \sim \mathcal{N}\Big(\beta_{i,h,j},\ \Sigma_{i,h}(j,j)/s_{\star\star,i,h}\Big), \qquad j = 1, \ldots, n. \tag{A.53}$$

If shock pooling is active, the cross-unit pooling prior for the shock

$$\beta_{i,h} \mid \mu_h^{(\star)}, \tau_i^2 \sim \mathcal{N}\big(\mu_h^{(\star)},\ \tau_i^2 \Sigma_{i,h}\big)$$

is combined with (A.52) or (A.53) to produce a single Gaussian measurement equation for the state-space update. In multivariate form,

$$R_{i,h}^{-1} = V_{i,h}^{-1} + (\tau_i^2 \Sigma_{i,h})^{-1}, \tag{A.54}$$

$$y_{i,h}^{\star} = R_{i,h}\Big(V_{i,h}^{-1}\widehat{y}_{i,h}' + (\tau_i^2 \Sigma_{i,h})^{-1}\mu_h^{(\star)}\Big), \tag{A.55}$$

so that

$$y_{i,h}^{\star} \mid \beta_{i,h} \sim \mathcal{N}(\beta_{i,h}, R_{i,h}). \tag{A.56}$$

If shock pooling is not active, then $y_{i,h}^{\star} = \widehat{y}_{i,h}'$ and $R_{i,h} = V_{i,h}$.

Hence, for each unit i, the smoothed shock path is sampled from the linear Gaussian state-space system

$$\beta_{i,0} \sim \mathcal{N}(m_{i,0}^{\beta}, C_{i,0}^{\beta}), \tag{A.57}$$

$$\beta_{i,h} = \beta_{i,h-1} + \eta_{i,h}, \qquad \eta_{i,h} \sim \mathcal{N}(0, \xi^2 \Sigma_\alpha), \qquad h = 1, \dots, H, \tag{A.58}$$

$$y_{i,h}^{\star} \mid \beta_{i,h} \sim \mathcal{N}(\beta_{i,h}, R_{i,h}), \qquad h = 0, \dots, H. \tag{A.59}$$

In the univariate version, Σ_α is replaced by $\mathrm{diag}(\Sigma_\alpha)$. Conditional on the remaining coefficient rows, the path $\{\beta_{i,h}\}_{h=0}^{H}$ is sampled by FFBS.

Full Conditional for ξ^2.
If $\xi^2 \sim \text{Inv-Gamma}(a_\xi, b_\xi)$ and $\Delta\beta_{i,h} = \beta_{i,h} - \beta_{i,h-1}$, then

$$\xi^2 \mid \{\beta_{i,h}\} \sim \text{Inv-Gamma}\Big(a_\xi + \frac{NHn}{2},\ b_\xi + \frac{1}{2}\sum_{i=1}^{N}\sum_{h=1}^{H} \Delta\beta_{i,h}' \Sigma_\alpha^{-1} \Delta\beta_{i,h}\Big), \tag{A.60}$$

with Σ_α^{-1} replaced by $\mathrm{diag}(\Sigma_\alpha)^{-1}$ in the univariate version.

A.8 Fixed Working Covariance Matrices

This subsection describes the fixed working covariance matrices used in the quasi-posterior updates. These matrices are introduced for scaling and weighting within the Gaussian working specification; they are not intended to represent the true serial dependence of local-projection forecast errors and are not used as long-run variance estimators for inference.

For each unit i and horizon h, the working covariance matrix $\Sigma_{i,h}$ is constructed from plug-in regression residuals using the lag-0 covariance estimator

$$\Sigma_{i,h} = \frac{1}{T_h}\,\widehat{U}_{i,h}'\widehat{U}_{i,h}, \qquad \widehat{U}_{i,h} = Y_{i,h} - X_{i,h}\widehat{B}_{i,h}, \tag{A.61}$$

where $\widehat{B}_{i,h}$ is obtained from the initialization step. Once constructed, $\Sigma_{i,h}$ is held fixed during posterior simulation. Thus $\Sigma_{i,h}$ is a fixed lag-0 working covariance matrix that governs how residual variation is scaled across outcomes in the Gaussian update, but it is not interpreted as the true covariance matrix of the local-projection forecast errors.

In implementation, $\Sigma_{i,h}$ is symmetrized and regularized by a small ridge term to ensure numerical stability. Two working specifications are allowed:

- *Diagonal working covariance:* $\Sigma_{i,h} = \mathrm{diag}(\Sigma_{i,h})$, so scaling is applied outcome by outcome.
- *Full working covariance:* $\Sigma_{i,h}$ is unrestricted, so scaling can also reflect contemporaneous comovement across outcomes.

These working covariance matrices are used only inside the quasi-posterior updating steps. Misspecification-robust inference is handled separately through the post hoc long-run variance estimators described in the main text and in Appendix A.10.

The matrix Σ_α used in the optional smoothing of unit-level shock responses is constructed as the average of $\Sigma_{i,h}$ across all (i, h) and is then symmetrized and regularized to be positive definite.

The matrix Σ_μ used in the random-walk innovations for the latent benchmark path is a correlation-normalized version of Σ_α:

$$\Sigma_\mu = D^{-1}\Sigma_\alpha D^{-1}, \qquad D = \mathrm{diag}\big(\sqrt{\Sigma_\alpha(1,1)}, \ldots, \sqrt{\Sigma_\alpha(n,n)}\big). \tag{A.62}$$

This normalization removes scale differences across outcomes, so that benchmark smoothing is governed by correlation structure rather than by the measurement units of the variables. The resulting matrix is again regularized to be positive definite.

A.9 Posterior Simulation Algorithm

The matrices $\Sigma_{i,h}$ are treated as fixed inputs throughout the MCMC. One iteration proceeds through the following blocks, with some steps included only when the corresponding specification is active:

1. **Optional smoothing of unit-level shock paths.** If unit-level shock smoothing is active, first update $\{B_{i,h}\}$ conditional on the current state, then update the smoothed shock paths $\{\beta_{i,h}\}_{h=0}^{H}$ under the RW1 prior in (A.50), and finally redraw $\{B_{i,h}\}$ conditional on the updated shock paths. If ξ^2 is treated as random, update it from (A.60).
2. **Update of $\{B_{i,h}\}$.** If shock-path smoothing is not active, sample each $B_{i,h}$ directly from the full conditional in (A.18).
3. **Update of $\{\tau_i^2\}$.** For each unit i, sample the unit-specific pooling variance τ_i^2 from (A.22).

4. **Update of the latent benchmark path $\{\mu_h\}$.** If the latent benchmark is learned, update $\{\mu_h\}_{h=0}^{H}$ by FFBS using (A.25)–(A.41). When benchmark smoothing is not imposed, this step reduces to horizon-by-horizon Gaussian updating.

5. **Update of σ_μ^2.** If benchmark smoothing is active and σ_μ^2 is treated as unknown, sample it from (A.43).

6. **Update of $\{\lambda_h\}$.** Update $z_h = \log \lambda_h$ using the conditional target in (A.48). If the RW1 prior on z_h is active, update σ_z^2 from (A.49).

7. **Optional update of Σ_B.** If the optional cross-row covariance block is used, sample Σ_B from (A.24). This step is not used in the empirical application.

Bridge to robust inference.
The MCMC blocks described above deliver the regularized quasi-posterior center and the associated hierarchical latent states. Reported uncertainty for impulse responses is not taken from the dispersion of this quasi-posterior. Instead, pointwise uncertainty is calibrated separately using the misspecification-robust long-run variance procedures described in the main text. The resulting reported intervals should be interpreted as pointwise misspecification-robust intervals around a regularized center.

A.10 Asymptotic Justification for Pointwise Misspecification-Robust Inference on the Regularized Panel Estimator

The reported estimators are obtained from a regularized Gaussian working model combining Minnesota-type shrinkage and cross-sectional pooling.

In large samples with fixed cross-sectional dimension and horizon set, the regularization components act as higher-order terms in the corresponding first-order conditions. As a result, the reported regularized center behaves, to first order, like the unpenalized local-projection estimator targeting the pseudo-true parameter under misspecification (cf. Müller, 2013).

This motivates the use of misspecification-robust long-run variance estimators applied to the corresponding aggregate influence process for pointwise inference.

A formal asymptotic justification of the proposed regularized DK-style inference is currently being developed and will be provided in a subsequent version.

B Appendix Empirical Results and Robustness

B.1 Additional Empirical Results

This section complements the main empirical analysis and reports precision-weighted panel responses as well as the heterogeneity measures. Furthermore, country-specific and regional, i.e. Core versus Periphery, responses for industrial production, HICP inflation and unemployment are shown including the difference plots, respectively.

B.1.1 Precision-weighted Panel Responses

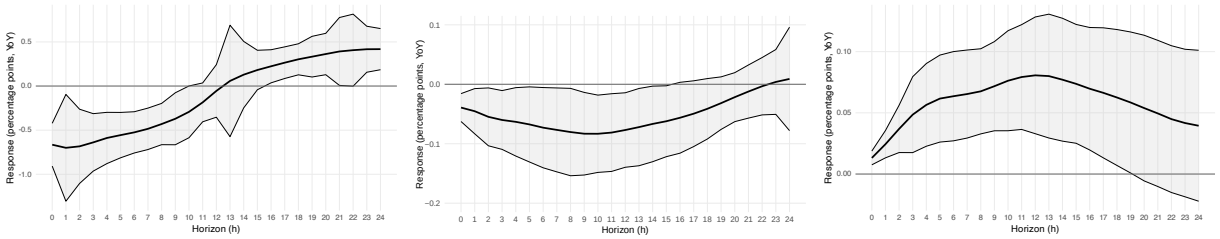

(a) Industrial production **(b)** HICP inflation **(c)** Unemployment

Notes: Each panel shows the regularized panel point estimates and 68% pointwise misspecification-robust intervals based on Driscoll–Kraay-style long-run variance calibration to a contractionary monetary policy shock.

B.1.2 Heterogeneity Measures

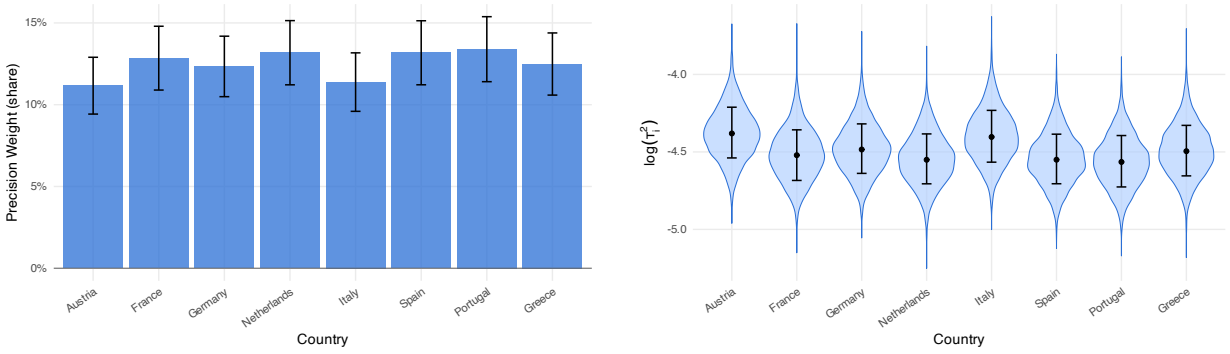

(a) *Posterior densities of the country-specific pooling variances τ_i^2 (log scale). Smaller values indicate stronger shrinkage toward the common component, meaning that the corresponding country's inflation response is estimated to be more coherent with the cross-sectional pattern captured by the panel benchmark.*

(b) *Posterior mean precision weights used in constructing the common inflation response. Weights are normalized to sum to one and are proportional to the inverse pooling variances $1/\tau_i^2$. Countries with responses estimated to be more coherent with the pooled component receive larger weights.*

B.1.3 Country-specific Responses: Industrial Production

Austria

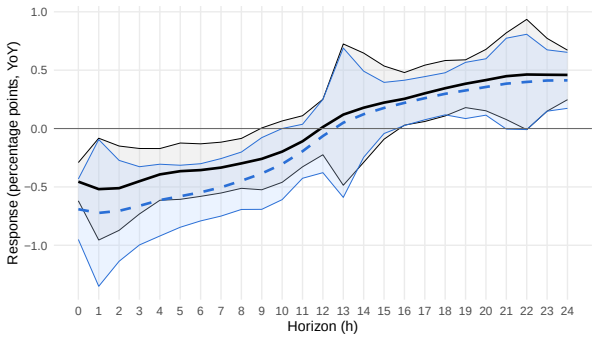

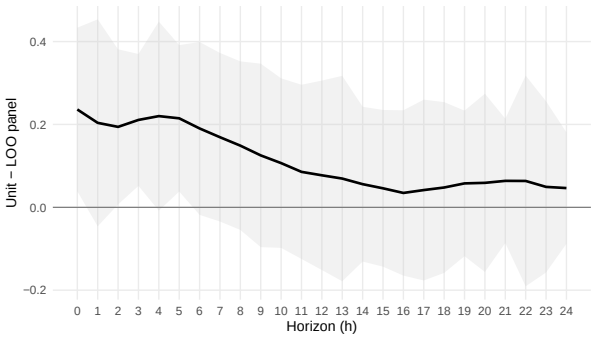

France

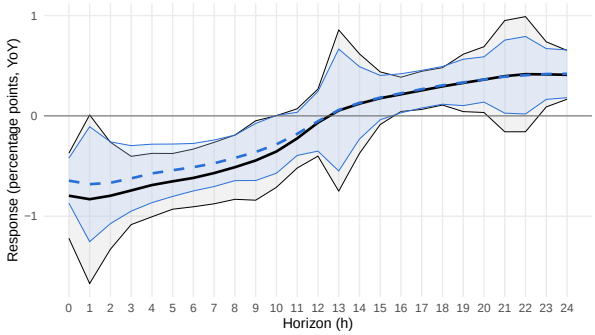

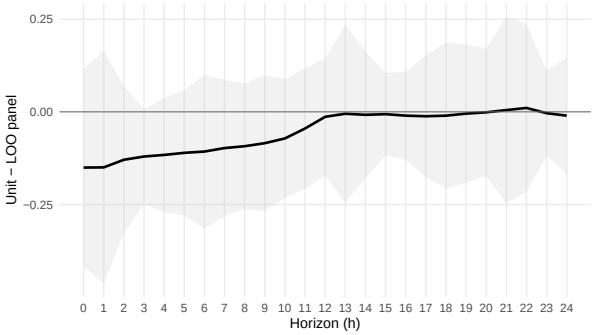

Germany

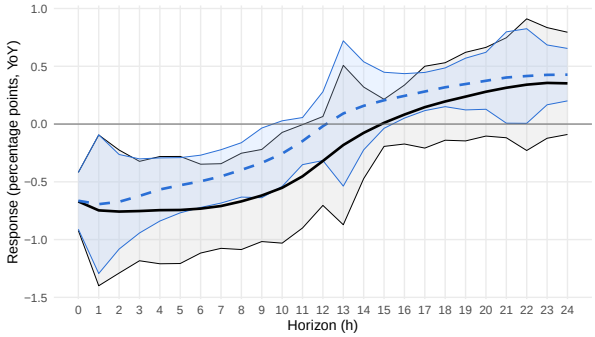

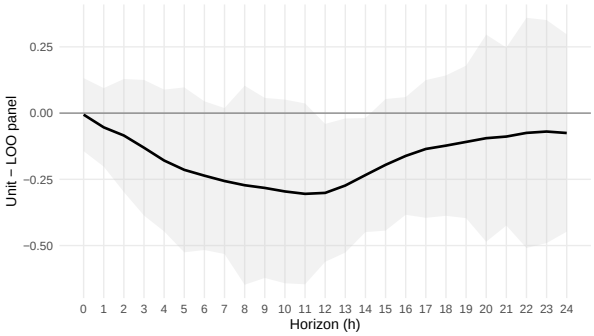

Netherlands

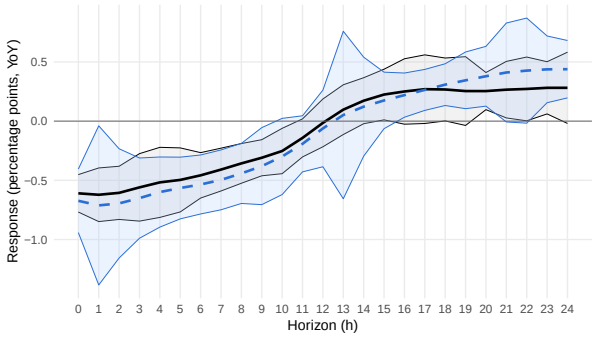

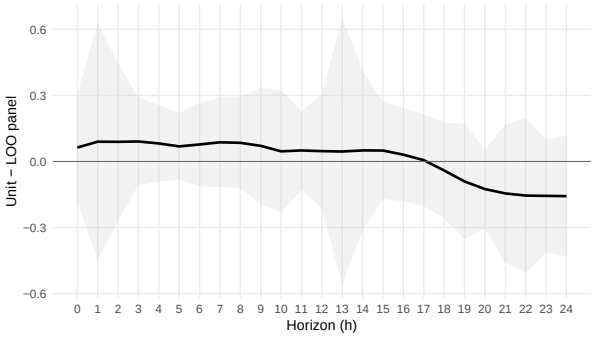

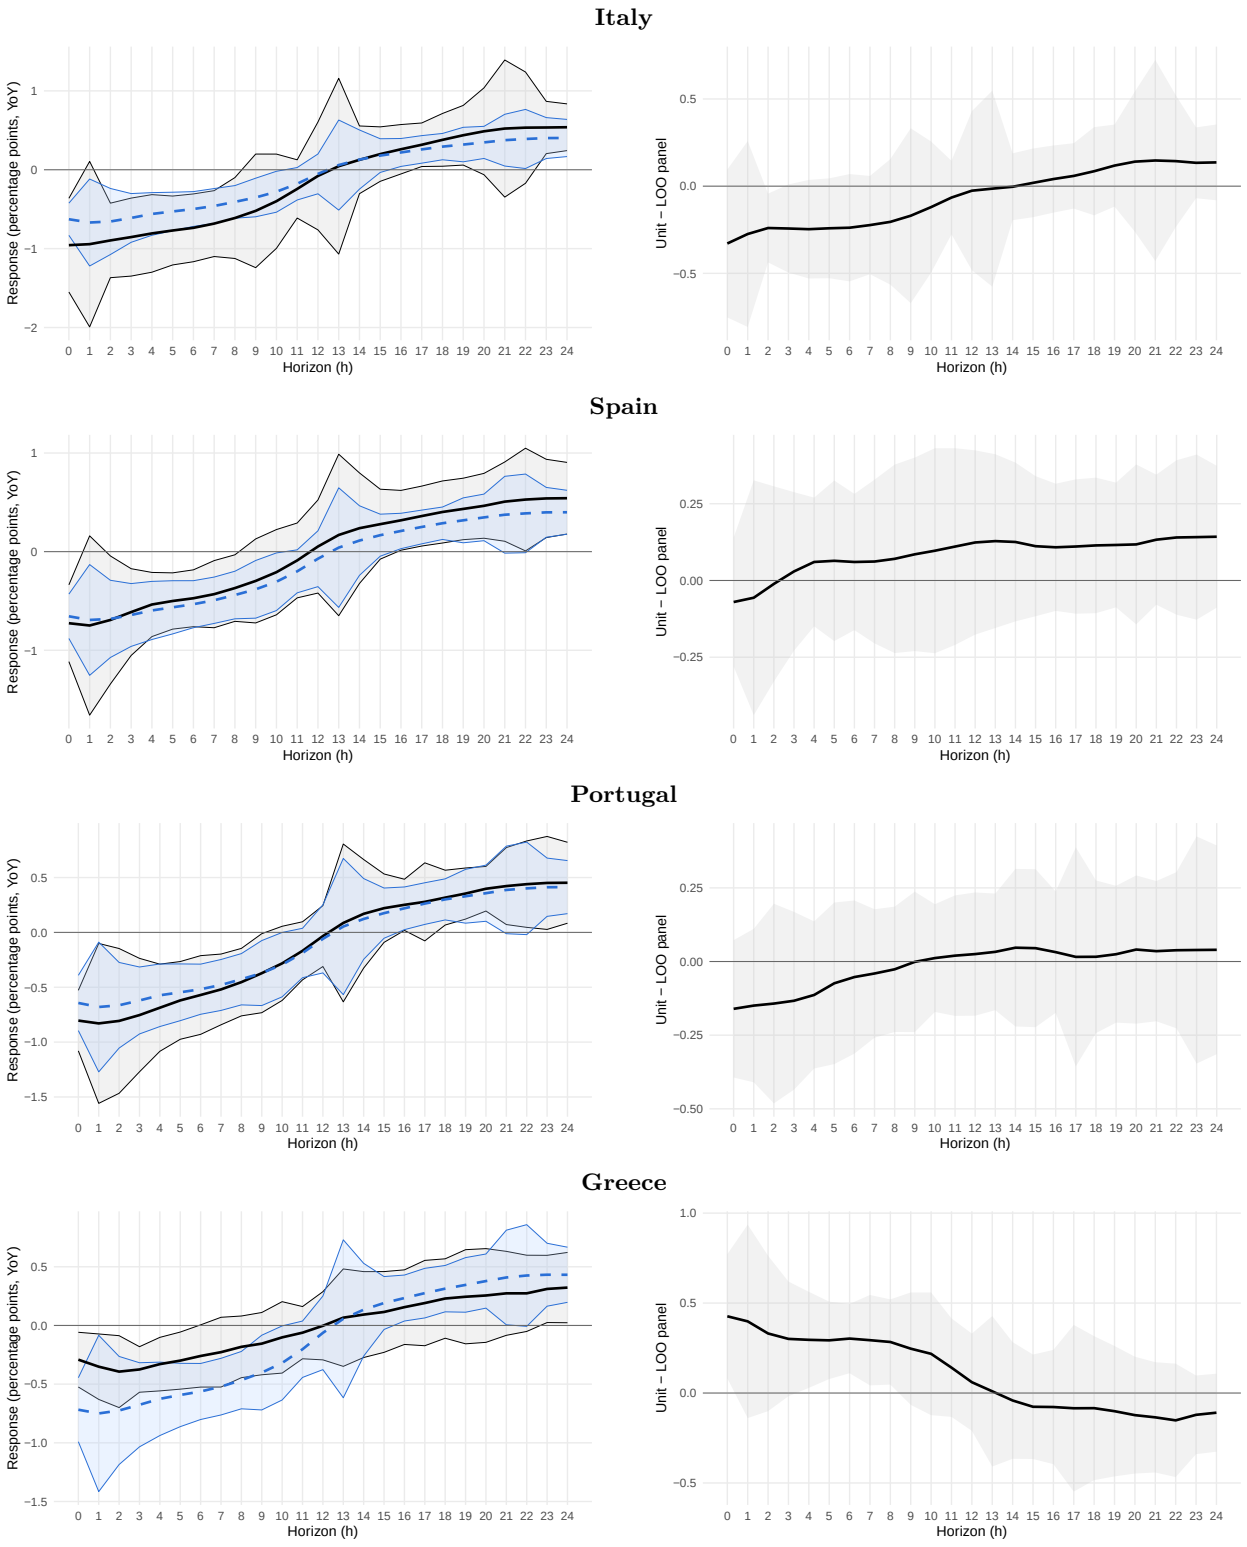

Notes: The left panel reports the country's (black line) response of industrial production relative to its leave-one-out benchmark (blue) after a contractionary monetary policy shock. Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals. The right panel shows the difference between the country's and the LOO's response, respectively. Again, with the 68% pointwise misspecification-robust intervals.

Country-specific Responses: HICP inflation

Austria

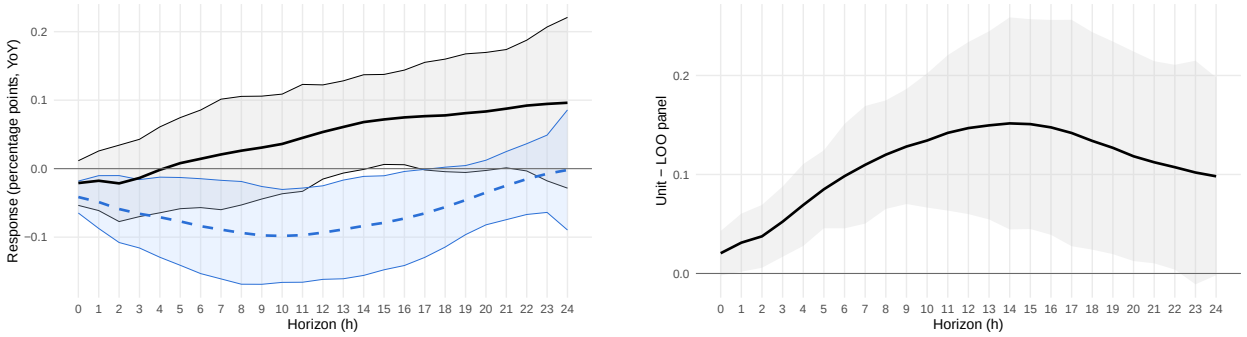

France

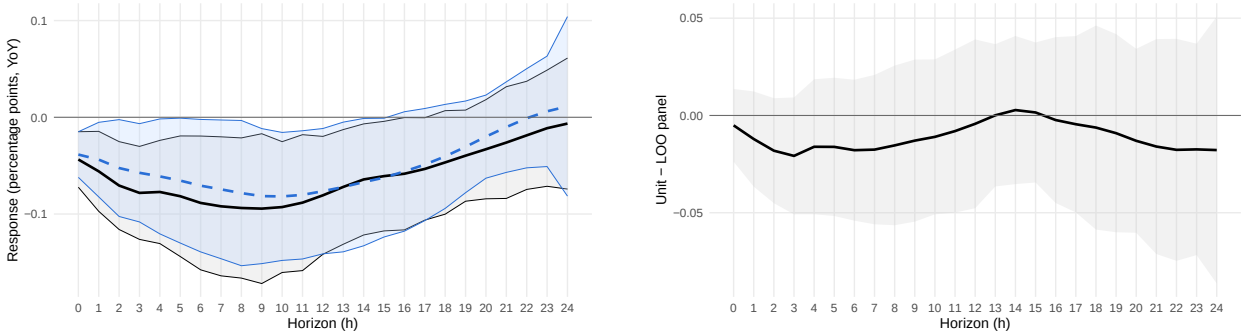

Germany

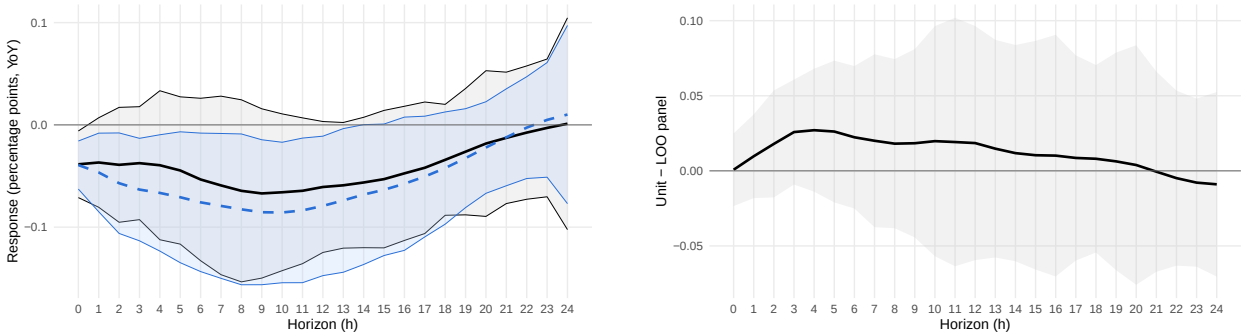

Netherlands

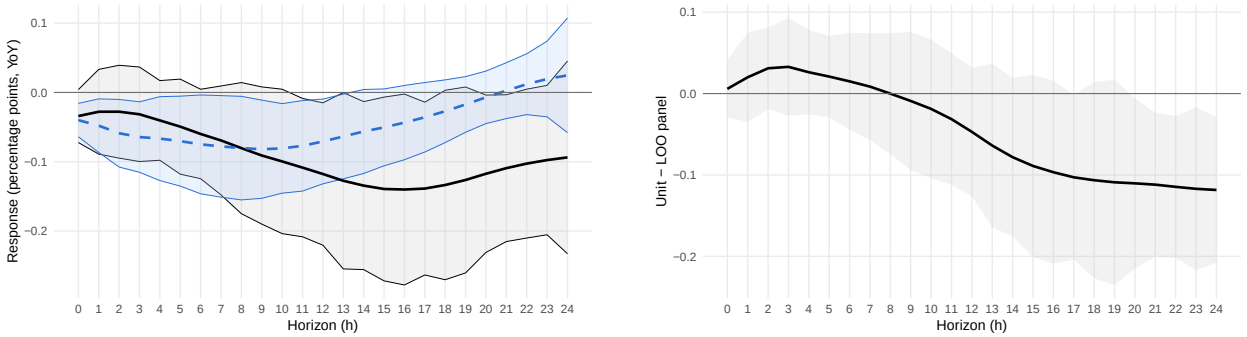

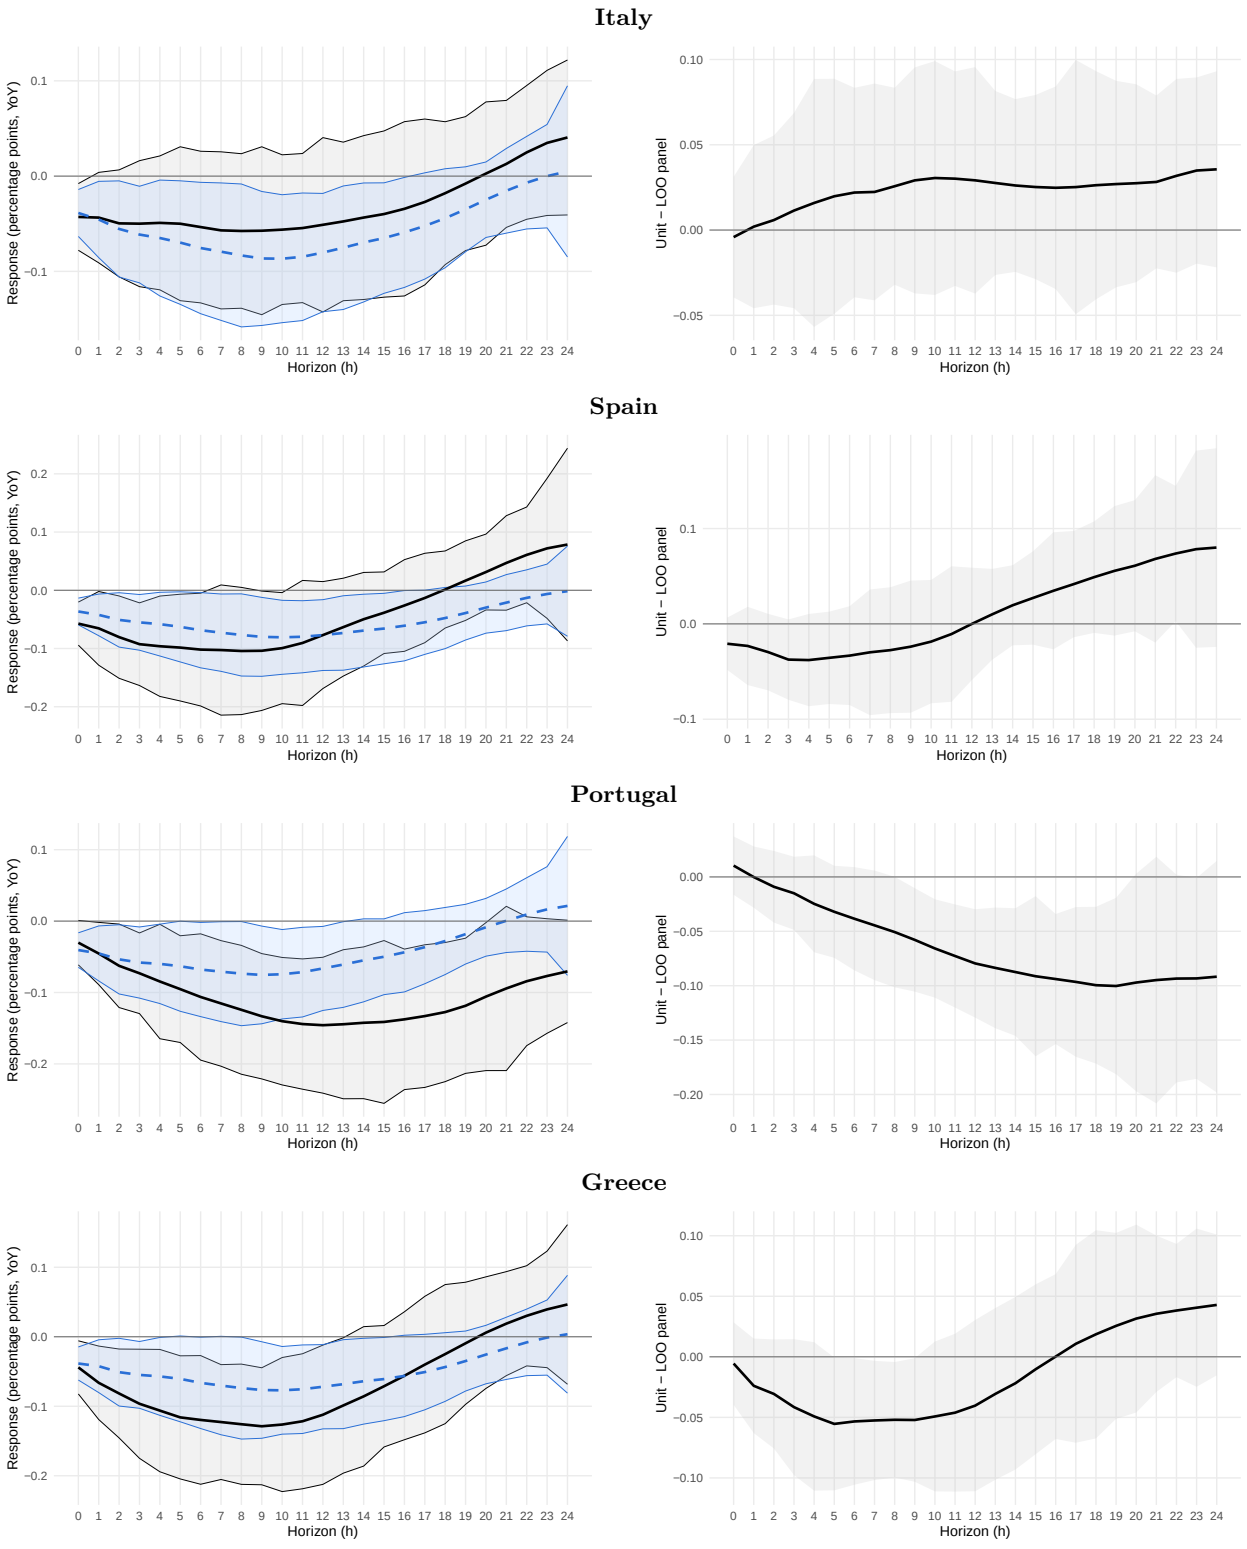

Notes: The left panel reports the country's inflation response (black line) relative to its leave-one-out benchmark (blue) after a contractionary monetary policy shock. Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals. The right panel shows the difference between the country's and the LOO's response, respectively. Again, with the 68% pointwise misspecification-robust intervals.

Country-specific Responses: Unemployment

Austria

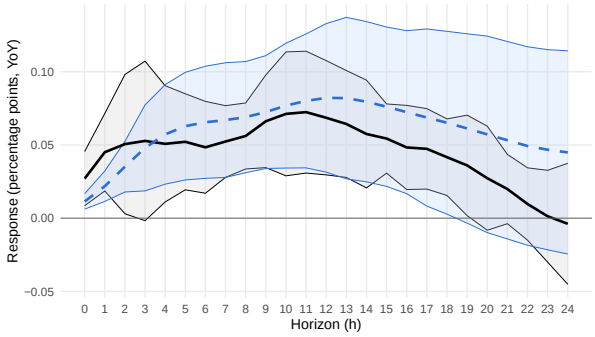

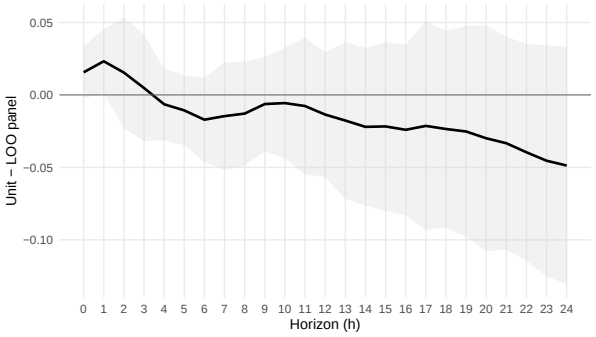

France

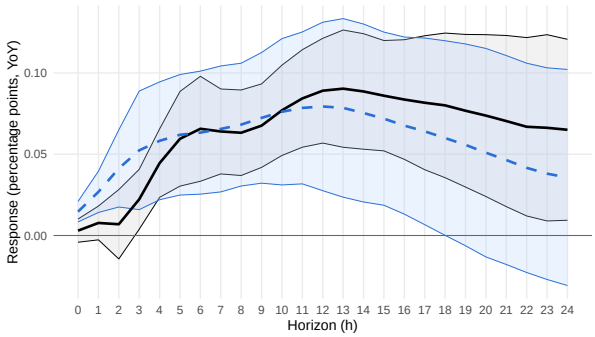

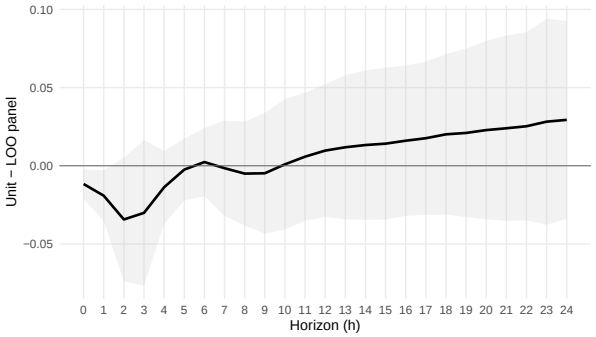

Germany

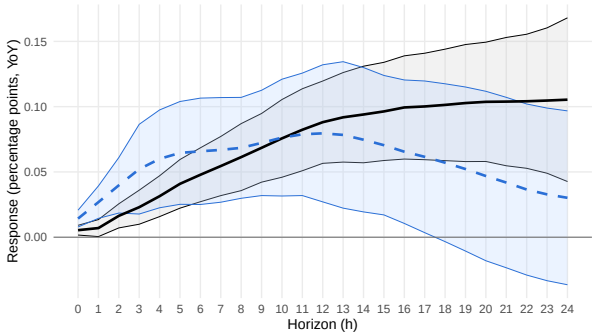

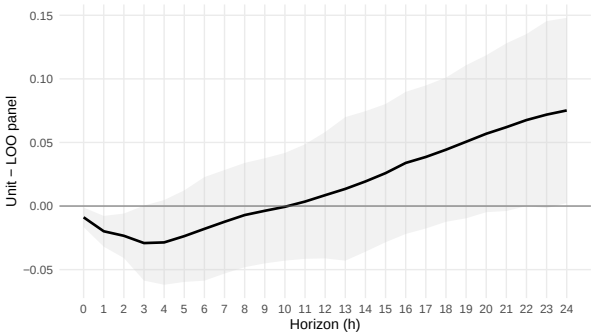

Netherlands

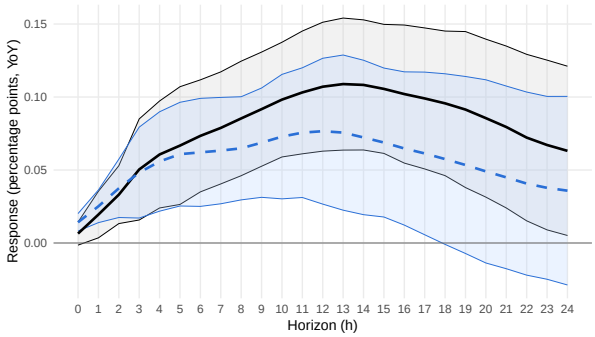

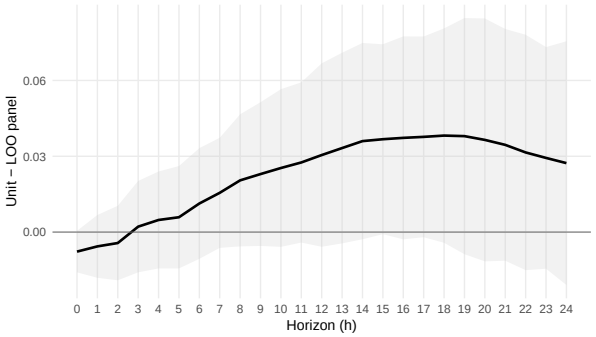

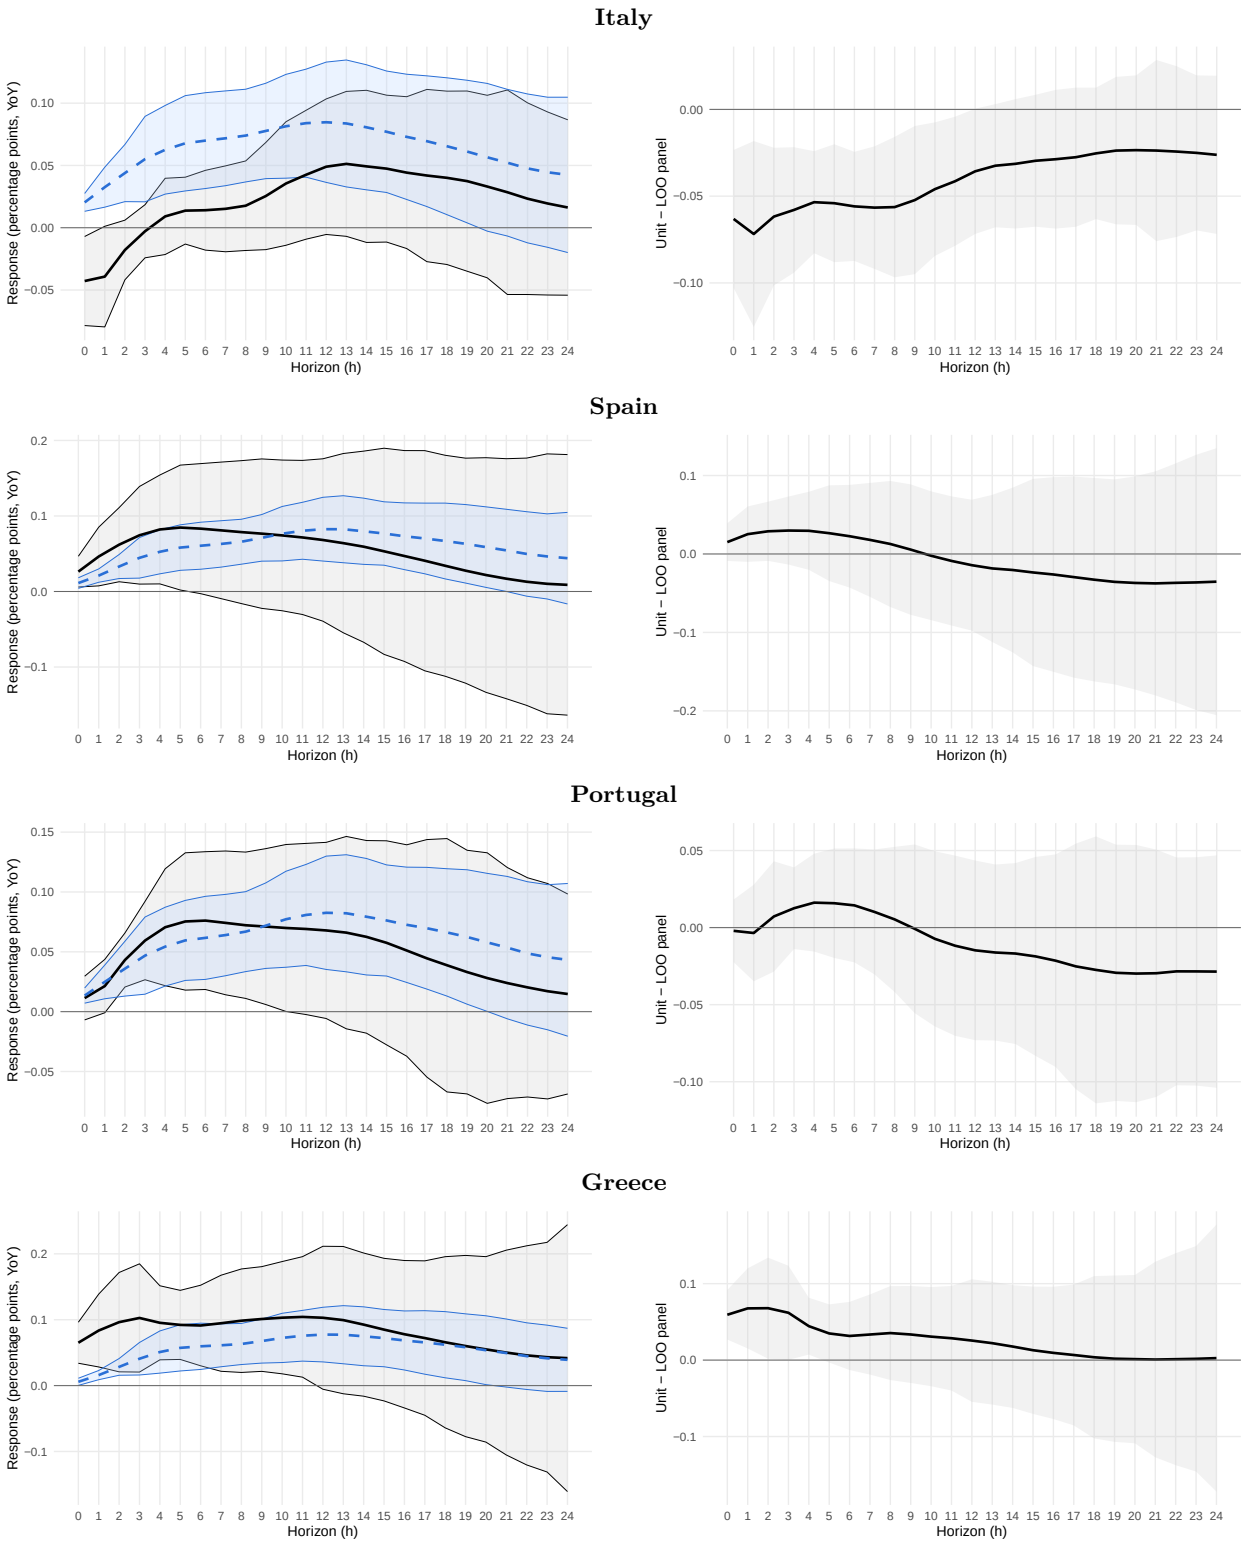

Notes: The left panel reports the country's (black line) response of unemployment relative to its leave-one-out benchmark (blue) after a contractionary monetary policy shock. Shown are the regularized point estimates together with 68% pointwise misspecification-robust intervals. The right panel shows the difference between the country's and the LOO's response, respectively. Again, with the 68% pointwise misspecification-robust intervals.

B.1.4 Regional Responses: Core versus Periphery

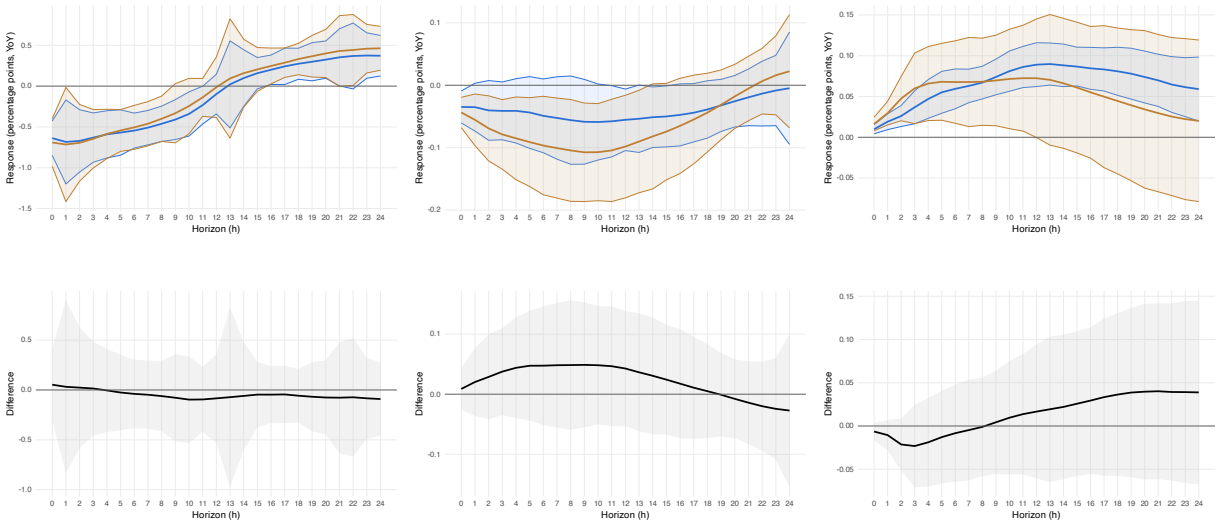

(a) Industrial production **(b)** HICP inflation **(c)** Unemployment

Notes: The upper panels display the regularized subgroup point estimates (Core: Austria, France, Germany, the Netherlands (blue); Periphery: Italy, Spain, Portugal, Greece (brown)) and 68% pointwise misspecification-robust intervals based on Driscoll–Kraay-style long-run variance calibration after a contractionary monetary policy shock. The lower panels show the difference between the Core and Periphery responses of the variables, respectively, after a contractionary monetary policy shock together with the 68% pointwise misspecification-robust interval.