Analyse

Uncertainty and the transmission of monetary policy in the euro area

September 2026

Maria Sole Pagliari & Minke van der Heijden

DeNederlandscheBank

EUROSYSTEEM

Uncertainty and the transmission of monetary policy in the euro area

Authors: Maria Sole Pagliari & Minke van der Heijden. With thanks to colleagues at DNB, and in particular Bas ter Weel, Niels Gilbert and Sebastian Rast for the useful exchange of views. All remaining errors are ours.

De Nederlandsche Bank NV
P.O. Box 98
1000 AB Amsterdam
Internet: www.dnb.nl
Email: info@dnb.nl

Executive Summary

This analysis examines whether monetary policy transmission in the euro area is uncertainty dependent. Do output and inflation respond differently to monetary policy when uncertainty is high rather than low? We distinguish between financial, macroeconomic and geopolitical uncertainty and construct euro-area-specific measures for each dimension from a broad set of indicators. These measures are used to identify low- and high-uncertainty regimes and estimate state-dependent local projections to compare the responses of output growth and inflation to a 25-basis-point monetary policy tightening across uncertainty regimes.

The results indicate that monetary policy transmission tends to be weaker when uncertainty is high. In low-uncertainty environments, a 25-basis-point policy tightening reduces output growth by around 0.39 percentage points at its peak and leads to a 0.07 percentage-point decline in inflation. By contrast, in high-uncertainty periods, the effects on both output and inflation are considerably weaker and generally statistically insignificant. While this pattern emerges across all three uncertainty dimensions, the attenuation of monetary policy transmission appears particularly pronounced during periods of elevated financial and geopolitical uncertainty.

The finding that monetary policy transmission weakens when uncertainty is high has implications for the interpretation of policy outcomes. The results do not imply that monetary policy becomes ineffective in periods of elevated uncertainty, but rather that the speed, strength and predictability of transmission may vary with the state of the economy and with the type of uncertainty that prevails. Therefore, to achieve comparable effects on inflation in a high-uncertainty environment, the central bank may need to act more forcefully. Furthermore, policymakers should pay close attention not only to the level of uncertainty, but also to its composition, as financial and geopolitical uncertainty seem to be particularly relevant for the effectiveness of monetary policy. This implies that it is important to treat uncertainty as a multidimensional phenomenon, and it suggests that policymakers should rely on a broad set of uncertainty indicators when assessing the monetary policy stance.

1 Introduction

The euro area has faced a sequence of unusually large shocks in recent years. The pandemic, Russia's war against Ukraine, energy-market disruptions, trade tensions and renewed geopolitical conflict in the Middle East have all contributed to a more uncertain economic environment. At the same time, monetary policy has aggressively tightened and subsequently eased its stance in response to such shocks. This combination raises an important question for central banks: does uncertainty merely complicate the outlook, or can it also change the way monetary policy is transmitted to the economy?

This study analyses whether monetary policy transmission depends on the type and level of uncertainty prevailing in the euro area. A policy rate change may have different effects depending on whether households, firms and financial markets face a relatively calm environment or one marked by elevated uncertainty. If high uncertainty makes economic agents postpone decisions, increases risk premia or weakens balance sheets, the response of demand and inflation to monetary policy may become slower and its effect may be smaller. Understanding this state dependence is relevant to central banks to assess both the monetary policy stance and the speed at which policy change affects the economy.

We approach this question by treating uncertainty as a multi-dimensional phenomenon and distinguishing between financial, macroeconomic and geopolitical uncertainty. Rather than relying on a single indicator, we construct euro-area uncertainty factors from a broad set of measures. This allows us to distinguish between episodes of financial stress, uncertainty about the macroeconomic outlook and geopolitical tensions, as these dimensions may operate through different channels. We also separate euro-area-specific uncertainty from global uncertainty, because global shocks may affect the euro area differently from uncertainty in the euro-area environment. The empirical strategy has two steps. First, we construct euro-area-specific uncertainty factors from a broad set of indicators and use them to classify months into low- and high-uncertainty regimes. Second, we estimate state-dependent local projections to compare the responses of output growth and inflation to the same monetary policy shock across regimes. Monetary policy shocks are measured using high-frequency monetary policy surprises and scaled to a 25-basis-point tightening.

The rest of the analysis is organised as follows. Section 2 explains how the uncertainty factors are constructed and shows that the three dimensions capture different episodes of stress. Section 3 describes the regime classification and presents the state-dependent monetary policy responses. Section 4 concludes by discussing the implications for policy assessment and communication.

2 Uncertainty factors

We summarise euro-area uncertainty in three dimensions: financial, macroeconomic and geopolitical. This distinction is important because not all uncertainty shocks are alike. Financial stress, uncertainty about the macroeconomic outlook and geopolitical tensions may affect the economy through different channels and may therefore matter differently for monetary policy transmission. To capture these dimensions, we start from six euro-area indicators: (i) the Composite Index of Financial Systemic Stress by Holló et al. (2012), (ii) the VSTOXX, (iii) the European Economic Policy Uncertainty index by Baker et al. (2016), (iv) the dispersion of inflation forecasts from the ECB Survey of Professional Forecasters, (v) the Euro Area Geopolitical Risk index by Bondarenko et al. (2024), and (vi) the Geoeconomic Tension index by Ioannou et al. (2026).

As euro-area indicators also move with global uncertainty, we remove common global movements from each series. Financial indicators are purged using the CBOE Volatility Index (VIX) and the one-year-ahead financial uncertainty measure of Jurado et al. (2015); macroeconomic indicators using the Global Economic Policy Uncertainty index by Baker et al. (2016) and the one-year-ahead macroeconomic uncertainty measure of Jurado et al. (2015); and geopolitical indicators using the Global Geopolitical Risk index by Caldara and Iacoviello (2022).[1] The remaining variation can be interpreted as the euro-area-specific component of each uncertainty indicator.

We then extract three factors from the euro-area-specific indicators to create our financial, macroeconomic and geopolitical uncertainty measures. Each factor is constructed so that it loads positively on the indicators associated with its own uncertainty dimension, as illustrated in Table 1. This gives three interpretable measures: financial uncertainty, macroeconomic uncertainty and geopolitical uncertainty. Technical details on the residualisation and factor-extraction procedure are reported in the Appendix.

Table 1. Positive factor loadings

Factor	Positively loading on:
Financial uncertainty	CISS, VSTOXX
Macroeconomic uncertainty	EPUEU, SPF disagreement
Geopolitical uncertainty	GPREA, GEOECON

Figure 1 shows that the three uncertainty factors capture different episodes of stress. Financial uncertainty rises most strongly during periods of market turmoil, such as the global financial crisis, the euro-area sovereign debt crisis and the first phase of the Covid-19 pandemic. Macroeconomic uncertainty is more persistent and increases around broad economic disruptions, including the pandemic recession, the energy-price shock following Russia's invasion of Ukraine and the subsequent inflation surge. Geopolitical uncertainty is more episodic, with spikes around major geopolitical and trade-related events which mainly took place in the past 10 years. These differences show that uncertainty is not a single phenomenon. They also motivate analysing monetary policy transmission separately across the three uncertainty dimensions.

[1] Refer to Appendix A for a description of all underlying measures.

Figure 1. Uncertainty differs by factor

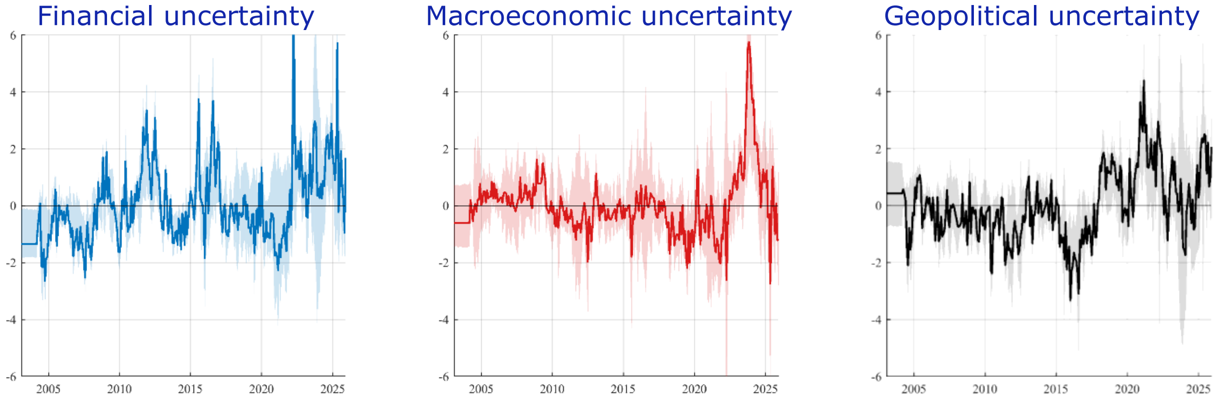

Notes: Solid lines report median extracted factors. Shaded areas depict 68% confidence bands based on moving-block bootstrap replications. Time period: 2003M1-2025M11. Once extracted, the factors are demeaned, while the standard deviation is close to 1.

3 Monetary policy transmission and uncertainty

Next, we estimate how our uncertainty measures impact the transmission of monetary policy. We use the state-contingent local projections framework of Jordà (2005) and Cloyne et al. (2023) to estimate the response of output and inflation to monetary policy shocks separately across low- and high-uncertainty regimes.[2] The uncertainty threshold is selected endogenously in the estimation. This allows us to quantify whether uncertainty affects the response of macro variables to a monetary policy shock. As a measure of output we use euro-area industrial production growth and for inflation the HICP. Monetary policy shocks are measured using the target factor of Altavilla et al. (2019), that captures movement in the short-end of yield curve as a proxy for conventional monetary policy.[3] The factor is aggregated at the monthly level using the weighted averaging of Jarociński and Karadi (2020). Responses are scaled to a 25-basis-point tightening. The local projections include lagged macroeconomic controls and are estimated with Bayesian methods (see appendix).

3.1 Regime identification

The regimes are identified separately for each of the three uncertainty factors. A month is classified as high uncertainty when the relevant factor is above an estimated threshold; all other months are classified as low uncertainty. The threshold is selected from the data, rather than imposed in advance, so that the regime split reflects the behaviour of each uncertainty dimension. This gives three regime classifications: one based on financial uncertainty, one on macroeconomic uncertainty and one on geopolitical uncertainty (Figure 2).

Figure 2. Uncertainty Regimes by Nature of Uncertainty

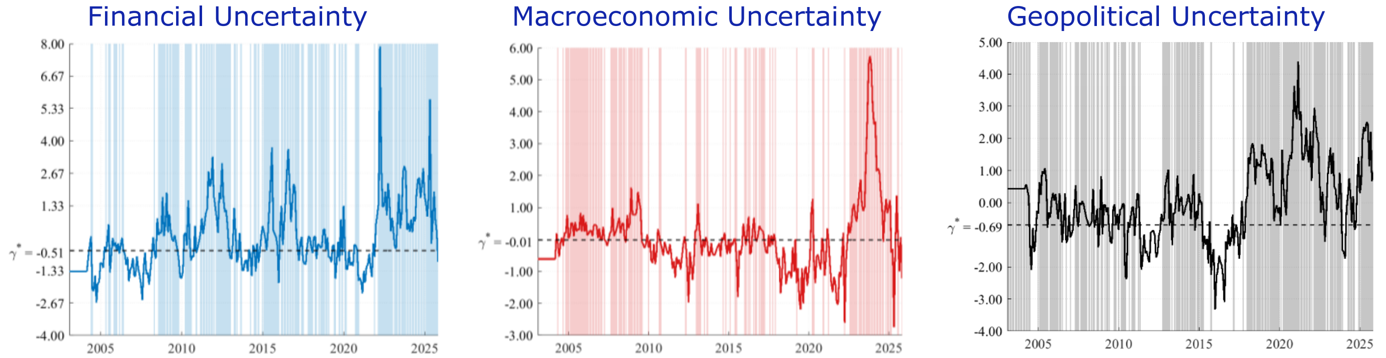

Notes: Solid lines report median extracted factors. Shaded areas depict periods of elevated uncertainty. γ^* indicates the endogenous threshold value. Time period: 2003M1-2025M11.

The three classifications point to partly different high-uncertainty environments. Financial uncertainty is elevated in about 60% of the sample, macroeconomic uncertainty in about 45%, and geopolitical uncertainty in about 70%. High-uncertainty episodes also differ in persistence: they last on average 6 months for financial uncertainty, 4 months for macroeconomic uncertainty and 7 months for geopolitical uncertainty.[4] The correlation coefficients reported in Table 2 confirm this pattern. Financial and macroeconomic uncertainty are negatively and significantly correlated, macroeconomic and geopolitical uncertainty are positively but only moderately correlated in terms of size, and the correlation between financial and geopolitical uncertainty is small and statistically insignificant. This supports treating the three dimensions separately when analysing monetary policy transmission.

[2] Cloyne et al. (2023) show that fiscal multipliers vary considerably with the monetary policy regime. Our focus, however, is on whether the transmission of exogenous monetary policy shocks differs across uncertainty regimes. As a result, the estimated responses should be interpreted as reduced-form effects that may also reflect other features of high-uncertainty episodes, including differences in fiscal policy support.

[3] Refer to Figure E1 in the appendix for a depiction of the shocks.

[4] In Table E1 in the appendix we also report the mean and standard deviation of the relevant outcome variables (output growth and inflation) across the different regimes.

Table 2. Linear Correlation across Factors

	Financial	Macroeconomic	Geopolitical
Financial	1.000	-0.377***	-0.063
Macroeconomic	-0.377***	1.000	0.196***
Geopolitical	-0.063	0.196***	1.000

*Notes: correlations are computed across first-differenced series; *** p<0.01, ** p<0.05, * p<0.10.*

3.2 Effects on monetary policy transmission

Across models, the results point to stronger transmission when uncertainty is low, but the pattern differs by source of uncertainty (Figure 3). In the financial-uncertainty model, output reacts clearly in the low-uncertainty regime: the peak response to a 25-basis-point tightening is -0.50 pp and the cumulated response is also negative, while high-uncertainty responses are not statistically significant. In the macroeconomic uncertainty model, the low-uncertainty peak response is smaller, at -0.30 pp, with also a smaller negative cumulated effect; high-uncertainty responses are again not significant. The geopolitical uncertainty model shows the largest low-uncertainty output response, with a peak of -0.67 pp and a sizeable cumulated effect, whereas high-uncertainty responses are not significant. Inflation responses are more muted: significant peak and cumulated effects appear mainly in the financial and geopolitical models under low uncertainty, while the macroeconomic model shows no robust inflation response on impact in either regime.[5] Overall, elevated uncertainty – especially of financial or geopolitical nature – appears to weaken monetary policy transmission, in line with the nonlinear uncertainty literature, including Caggiano et al. (2022), Alessandri and Mumtaz (2019), and Pellegrino et al. (2023), as well as recent ECB policy work.[6]

Figure 3. IRFs to a 25bps monetary policy shock

Peak response

Output growth

Inflation

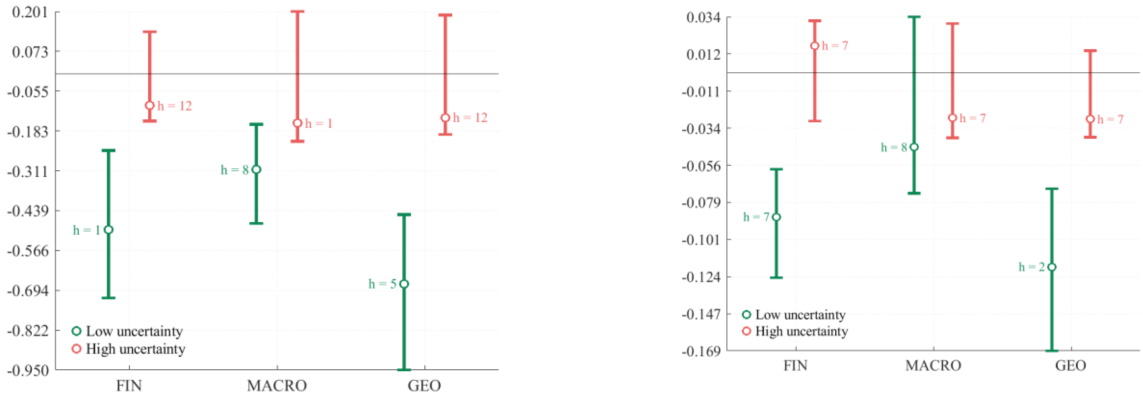

---
[5] Results are also summarized in Table E2 in the appendix, where we also report the responses over the whole sample, without distinguishing across the different uncertainty regimes.
[6] See the ECB blog by Falconio and Schumacher (September 2025) and the ECB blog by Allayioti et al. (October 2025).

Cumulated response

Output growth Inflation

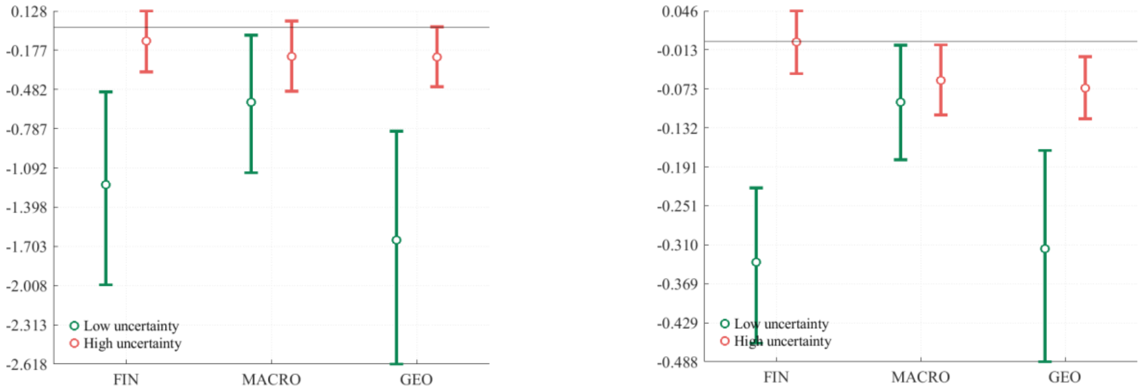

Notes: Variables are in y-o-y growth rates. Responses are in percentage points deviation from the baseline. The cumulated response is computed as the sum of the impulse responses over the relevant horizon. Dots report the median response, while whiskers indicate 68% confidence intervals.

We combine the three uncertainty factors and the aggregate IRFs show that monetary policy transmission is stronger and more persistent when uncertainty is low (Figure 4). To account for modelling uncertainty, we combine the information from the three specifications by means of Bayesian techniques.[7] In this aggregate specification, a 25-basis-point monetary policy shock lowers output growth by 0.39 pp at the peak in low-uncertainty periods, while the peak response in high-uncertainty periods is not statistically significant. Inflation also reacts mainly under low uncertainty, with a peak decline of 0.07 pp, compared to an insignificant response when uncertainty is high. The same pattern holds over the response horizon: the effects are bigger and more persistent in the low-uncertainty regime, whereas high-uncertainty responses are weaker, though more precisely estimated.[8]

Figure 4. IRFs to a 25bps monetary policy shock

Output growth Inflation

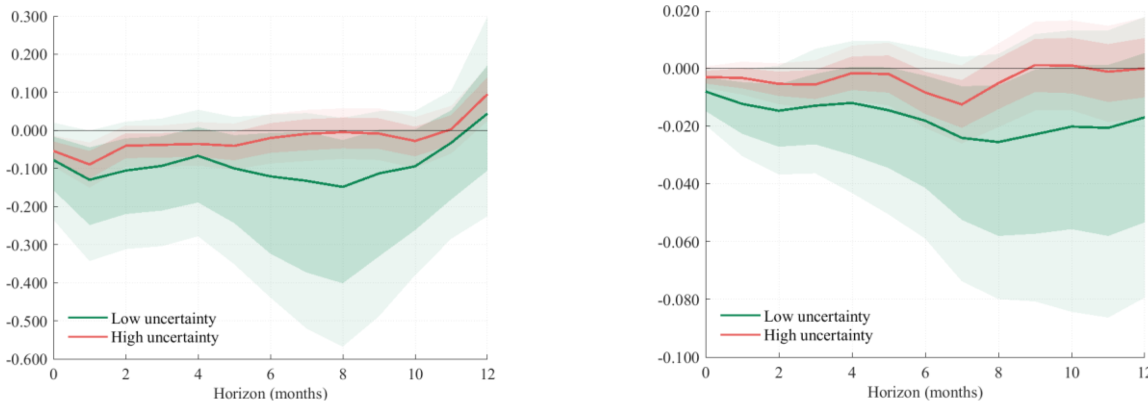

Notes: Solid lines indicate median IRFs averaged across models using BMA. Darker (lighter) shaded areas indicate 68% (90%) high density posterior sets.

---

[7] Notably, we use Bayesian Model Averaging (BMA), which is a weighting approach where weights are based on the relative marginal likelihoods of the different models, so that specifications with stronger empirical support receive more weight in the aggregate impulse responses. See Appendix D for details.

[8] This is because most months in our sample are classified as high uncertainty periods. With more observations assigned to the low-uncertainty regime, the low-uncertainty IRFs would likely be estimated more precisely. If anything, this would probably reinforce the difference between regimes, so the reported low-uncertainty effects can be interpreted as conservative estimates.

4 Conclusion

This analysis shows that the transmission of monetary policy depends on the uncertainty environment in which policy decisions are taken. Across the three uncertainty dimensions considered, a 25-basis-point monetary policy tightening has stronger and more persistent effects on output growth when uncertainty is low. By contrast, when uncertainty is high, output responses are generally weaker and often not statistically significant. Inflation responses are more muted overall, but the same broad pattern emerges: the disinflationary effect of monetary policy is clearer in low-uncertainty regimes, while high-uncertainty responses are less robust. The results of the aggregate uncertainty measures confirm this message. They show a significant peak decline in output growth of around 0.39 percentage points and a smaller but significant decline in inflation in low-uncertainty periods, whereas the corresponding high-uncertainty responses are not statistically significant.

A practical implication is that uncertainty should be monitored as a multi-dimensional phenomenon. Financial, macroeconomic and geopolitical uncertainty capture different episodes, have different origins and different persistence, and are only imperfectly correlated. They also deliver different estimates of the strength of monetary policy transmission, even though the qualitative pattern is consistent across models. This argues against drawing strong policy conclusions from a single uncertainty indicator or model specification. Instead, policy assessment should rely on a broad uncertainty dashboard, complemented by model averaging or other robustness checks, to evaluate not only the expected macroeconomic effects of monetary policy, but also the uncertainty surrounding those effects. This is particularly relevant because high-uncertainty environments appear to be associated not only with weaker average responses, but also with less precise estimates of the effects of policy.

In times of high-uncertainty central banks may need to act more forcefully to achieve comparable effects on inflation, while at the same time remaining agile to the uncertain environment. The results do not imply that monetary policy becomes ineffective in periods of elevated uncertainty. Rather, they suggest that the speed, strength and predictability of transmission may vary with the state of the economy and with the type of uncertainty that dominates. This has direct implications for policy assessment. If high uncertainty weakens the response of demand and inflation to policy tightening, then a given policy move may generate a smaller near-term impact compared to a period of low uncertainty. This may warrant a stronger policy intervention. At the same time, because the effects of policy are also surrounded by greater uncertainty in a high-uncertainty environment, forcefulness should be combined with agility: policy needs to remain data-dependent, continuously reassessed and capable of correcting course as new information arrives. This should also be reflected in the choice of instruments.

The appropriate policy response should consider all available information, not just the uncertainty level. The finding that uncertainty can dampen monetary transmission does not mechanically imply that policy should always do more in high-uncertainty periods. Other features of the macroeconomic environment may work in the opposite direction. For example, a recent DNB analysis for the euro area suggests that Phillips curves may steepen in high-inflation states, making disinflation more effective and the inflation-output trade-off state-dependent; in such an environment, earlier and stronger monetary tightening may reduce inflation with relatively lower output costs. This suggests that the appropriate policy response depends on the joint assessment of uncertainty, inflation persistence, the size of the shock, the slope of the Phillips curve and the risk that expectations become de-anchored.

For policy communication, the key implication is that central banks should explain state dependence without appearing passive. In high-uncertainty periods, the absence of an immediate output or inflation response should not necessarily be interpreted as evidence that monetary policy is not working; it may reflect a slower and less predictable transmission process. But precisely because the effects of policy are weaker, slower and more uncertain, communication needs to make clear that the central bank remains committed to keeping inflation expectations anchored. This requires explaining that policy may need to be sufficiently forceful when inflation risks are persistent, while also remaining agile and data-dependent as the transmission of past decisions becomes clearer. In this sense, monitoring uncertainty is not only a diagnostic exercise: it is part of assessing the appropriate intensity, timing and communication of monetary policy.

References

[1] Alessandri, P. and H. Mumtaz (2019), "Financial regimes and uncertainty shocks", *Journal of Monetary Economics*, 101, 31-46.

[2] Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto and G. Ragusa (2019), "Measuring euro area monetary policy", *Journal of Monetary Economics*, 108, 162-179.

[3] Baker, S. R., N. Bloom and S. J. Davis (2016), "Measuring economic policy uncertainty", *Quarterly Journal of Economics*, 131(4), 1593-1636.

[4] Bondarenko, Y., M. Grothe, T. Larsen and A. O. Vermeulen (2024), "Euro area geopolitical risk", ECB Working Paper.

[5] Caggiano, G., E. Castelnuovo and G. Nodari (2022), "Uncertainty and monetary policy in good and bad times: A replication of the VAR investigation by Bloom (2009)", *Journal of Applied Econometrics*, 37(1), 210-216.

[6] Caldara, D. and M. Iacoviello (2022), "Measuring geopolitical risk", *American Economic Review*, 112(4), 1194-1225.

[7] Cloyne, J., Ò. Jordà and A. M. Taylor (2023), "State-dependent local projections: Understanding impulse response heterogeneity", NBER Working Paper No. 30971.

[8] Holló, D., M. Kremer and M. Lo Duca (2012), "CISS - A composite indicator of systemic stress in the financial system", ECB Working Paper No. 1426.

[9] Ioannou, D., R. Prioriello and A. Durrani (2026), "Measuring geoeconomic tension: A large language model approach for the euro area", forthcoming ECB Working Paper Series.

[10] Jarociński, M. and P. Karadi (2020), "Deconstructing monetary policy surprises: The role of information shocks", *American Economic Journal: Macroeconomics*, 12(2), 1-43.

[11] Jordà, Ò. (2005), "Estimation and inference of impulse responses by local projections", *American Economic Review*, 95(1), 161-182.

[12] Jurado, K., S. C. Ludvigson and S. Ng (2015), "Measuring uncertainty", *American Economic Review*, 105(3), 1177-1216.

[13] Pellegrino, G., E. Castelnuovo and G. Caggiano (2023), "Uncertainty and monetary policy during the Great Recession", *International Economic Review*, 64(3), 1115-1145.

Appendix

A. Underlying Indexes to Construct Uncertainty Factors

A.1 Euro area uncertainty indicators

Composite Indicator of Systemic Stress (CISS)

The Composite Indicator of Systemic Stress (CISS), developed by the ECB, measures stress in the euro area financial system. It combines information from money, bond, equity, foreign exchange and financial intermediary markets. The indicator assigns larger weights to periods in which stress occurs simultaneously across multiple market segments and is therefore designed to capture systemic financial stress rather than isolated market disruptions. Higher values indicate elevated financial stress.

VSTOXX

VSTOXX is the implied volatility index of the EURO STOXX 50 stock market index. It is computed from option prices and reflects market expectations regarding future equity market volatility over approximately the next thirty days. Like the VIX for the United States, it is commonly interpreted as a forward-looking measure of financial market uncertainty.

B. European Economic Policy Uncertainty Index (EPUEU)

The European Economic Policy Uncertainty (EPU) index is developed by Baker, Bloom and Davis and is based on newspaper coverage of policy-related economic uncertainty. It aggregates information from major European newspapers and increases when media references to uncertainty regarding economic policies become more frequent. The index is intended to capture uncertainty related to fiscal, monetary, regulatory and other economic policies.

C. SPF disagreement

The SPF disagreement measure is derived from the ECB Survey of Professional Forecasters. It is computed as the cross-sectional dispersion of forecasts across survey participants and captures heterogeneity in expectations regarding future macroeconomic developments. Greater forecast dispersion is interpreted as reflecting higher uncertainty about the economic outlook.

D. Euro Area Geopolitical Risk Index (GPREA)

The euro area geopolitical risk index (EA-GPR) is a newspaper-based measure of geopolitical risk constructed using major newspapers from Germany, France, Italy, Spain and the Netherlands. The index tracks references to geopolitical tensions, military conflicts and terrorism and aggregates national measures using GDP weights. It is designed to capture geopolitical uncertainty from a specifically euro area perspective.

Geoeconomic Risk Index (GEOECON)

The geoeconomic risk indicator measures uncertainty arising from international economic tensions, including trade conflicts, sanctions, export restrictions, strategic dependencies and broader geoeconomic fragmentation. The indicator is intended to capture risks associated with disruptions to international economic and financial integration.

A.2 Global uncertainty indicators

VIX

The VIX is the Chicago Board Options Exchange volatility index derived from S&P 500 option prices. It represents market expectations of future stock market volatility and is widely used as a measure of global financial uncertainty.

Jurado Financial Uncertainty (JURF)

The Jurado-Ludvigson-Ng financial uncertainty measure is based on the forecast uncertainty associated with a large set of financial variables. Unlike volatility-based indicators, it measures uncertainty as the conditional unpredictability of future outcomes rather than the level of market fluctuations.

Global Economic Policy Uncertainty (EPUG)

The global EPU index extends the Baker-Bloom-Davis methodology to a broad set of countries and captures policy-related economic uncertainty at the global level. It increases during periods of heightened uncertainty regarding economic policy decisions worldwide.

Jurado Macroeconomic Uncertainty (JURM)

The Jurado-Ludvigson-Ng macroeconomic uncertainty measure is constructed from forecast errors across a large set of macroeconomic indicators. It captures the common component of uncertainty regarding future macroeconomic conditions.

Geopolitical Risk Index (GPR)

The geopolitical risk index of Caldara and Iacoviello is based on the frequency of newspaper articles discussing geopolitical tensions, wars, military conflicts and terrorism. It is one of the most widely used measures of global geopolitical risk and serves as the benchmark global geopolitical uncertainty indicator in this study.

B. First-stage residualization of uncertainty measures

To isolate euro-area-specific uncertainty dynamics from global co-movement, each euro-area uncertainty indicator is orthogonalized with respect to a pre-assigned set of global counterparts. Concretely, for each index $(y_{i,t})$, we estimate a separate projection on contemporaneous and lagged global controls:

$$y_{i,t} = \alpha_i + \sum_{k \in \mathcal{G}i} \beta\, i,k g_{k,t} + \sum_{\ell=1}^{L} \sum_{k \in \mathcal{G}i} \delta\, i,k,\ell g_{k,t-\ell} + u_{i,t},$$

and retain the residuals $(u_{i,t})$ as the orthogonalized euro-area component. In the baseline specification, we demean each $y_{i,t}$, we set L=13 and we include a constant. Missing observations are interpolated before estimation (linear, then nearest-neighbor), while lag-induced boundary gaps are filled by nearest-neighbor extrapolation after residualization to preserve a balanced sample.

C. Factors extraction

The latent uncertainty factors are extracted from the purged euro-area index panel after standardization. Let $x_t \in \mathbb{R}^N$ denote the vector of purged indexes at time t, with $\mathrm{N} = 6$. Each series is centered and scaled as

$$z_{it} = \frac{x_{it} - \mu_i}{\sigma_i}, \qquad i = 1, \dots, N,; t = 1, \dots, T.$$

Stacking the standardized data in Z, we estimate a three-factor PCA representation,

$$z_t \approx \Lambda_0 f_t^{(0)},$$

where $\Lambda_0 \in \mathbb{R}^{N\times K}$, $f_t^{(0)} \in \mathbb{R}^K$, and $K = 3$. Because PCA factors are not uniquely identified up to orthogonal rotations, we impose economic interpretability by rotating Λ_0 and $f_t^{(0)}$ with randomly drawn orthogonal matrices Q such that $Q'Q = I_K,\ \det(Q) = +1$, and $\Lambda(Q) = \Lambda_0 Q,\ f_t(Q) = f_t^{(0)} Q$.

Identification is based on block-consistent loading patterns. Let $\mathcal{B}_{FIN}$, $\mathcal{B}_{MACRO}$, and $\mathcal{B}_{GEO}$ denote the indicator sets assigned to the financial, macro, and geopolitical factors. For each candidate rotation, the target factor for block j must have sufficiently positive own-block loadings:

$$\lambda ij^{own} > \tau, i \in \mathcal{B}j,$$

where τ is the stricter lower bound between the minimum loading cutoff and the positivity tolerance. Under stricter implementation, we also impose that own loadings must also exceed cross-loadings by a margin:

$$\lambda ij^{own} \geq \max_{k\neq j} \lambda_{ik} + \text{margin}.$$

Feasibility is evaluated under progressively relaxed criteria (all variables in the block satisfy restrictions; then majority; then mean-loading criterion). The algorithm accepts the first feasible rotation encountered in the random search.

When no strictly feasible draw is found, the selected rotation is the one maximizing a soft objective that rewards large own-block loadings and penalizes cross-loadings and restriction violations: $\mathcal{J}(Q) = \sum_{j=1}^{K} \left(\overline{\lambda} j^{own} - 0.25, \overline{|\lambda|} j^{cross}\right) - \mathcal{P}(Q)$, with $\mathcal{P}(Q) = 50 \cdot \text{(sign violations)} + 25 \cdot \text{(dominance violations)} + 1000 \cdot \text{(empty-block failures)}$.

This fallback guarantees an economically interpretable rotation even when exact hard restrictions are not jointly satisfiable in finite samples. After a rotation is selected, factor signs are oriented towards economic readability. For each factor j, let

$$\bar{z}_{j,t} = \frac{1}{|\mathcal{B}j|} \sum i \in \mathcal{B}_j\, zit.$$

If $\operatorname{corr}(f_{j,t}, \bar{z}_{j,t}) < 0$, both the factor and the corresponding loading column are multiplied by (-1): $f_{j,t} \leftarrow -f_{j,t},\ \Lambda_{\cdot j} \leftarrow -\Lambda_{\cdot j}$. The final identified factor vector is therefore $F_t = (FIN_t,\ MACRO_t,\ GEO_t)'$. To quantify uncertainty in the latent uncertainty factors, we implement a moving-block bootstrap on the purged indicator panel, using 12-month blocks to preserve short-run serial dependence and cross-series co-movement (for a total of 1000 draws). For each bootstrap replication, we re-estimate the FIN, MACRO, and GEO factors under the same extraction and sign-restriction scheme as in the baseline sample, and we construct pointwise 68% bands from the empirical bootstrap distribution. Crucially, these bootstrap factor paths are not used only for descriptive bands: they are fed into the subsequent Bayesian local-projection step as alternative latent state paths. Posterior draws of regime-dependent impulse responses are then formed by integrating over both coefficient uncertainty and latent-state uncertainty, so that the reported credible sets reflect uncertainty both in the regression parameters and in the estimated state variables that determine regime classification.

D. State-dependent local projections

The state-dependent local projections are estimated separately for the financial, macroeconomic and geopolitical uncertainty factors constructed in Appendix C. All variables are monthly and the estimation sample covers the period 2003M1-2025M11. The response variables are euro-area industrial production growth and HICP inflation. Monetary policy shocks are measured using the target factor of Altavilla et al. (2019), aggregated at the monthly level using the weighted averaging of Jarociński and Karadi (2020). Responses are scaled to a 25-basis-point tightening.

For each uncertainty factor f_t, we define a binary high-uncertainty indicator, S_t, as $S_t(\gamma, d) = \mathbf{1}\{f_{t-d} > \gamma\}$, where γ is the threshold and d is the threshold delay. In the baseline specification, d=1. For each response variable and horizon $h = 0, \ldots, 13$, we estimate the state-dependent local projection

$$y_{t+h} = \alpha_h + \beta_h^L m_t(1 - S_t) + \beta_h^H m_t S_t + \delta_h S_t + \kappa_h f_{t-d} + \Gamma_h{}' X_{t-1} + u_{t+h},$$

where m_t is the monetary-policy target shock and X_{t-1} contains thirteen lags of macroeconomic controls (both y own lags and lags of other dependent macroeconomic aggregates). Notably, the baseline control set includes lags of output and inflation as well as of monetary shock. The regime-specific impulse responses at horizon h to a 25-basis-point tightening are $\Phi_h^L = 0.25\beta_h^L$ in the low-uncertainty regime, $\Phi_h^H = 0.25\beta_h^H$ in the high-uncertainty one, the difference being $\Delta\Phi_h = 0.25(\beta_h^H - \beta_h^L)$.

Estimation is Bayesian and is conducted separately for each horizon, response variable and uncertainty factor. Conditional on a regime classification, the local projection can be written as:

$$y_h = X_h\theta_h + u_h, \ with \ u_h \sim N(0, \sigma_h^2 I).$$

We use a conjugate normal-inverse-gamma prior, $\theta_h \mid \sigma_h^2 \sim N(\theta_0, \sigma_h^2 V_0)$, $\quad \sigma_h^2 \sim IG(a_0, b_0)$, with $a_0 = b_0 = 0.01$. The prior mean is set to zero, including for the own first lag. The prior variance matrix V_0 is diagonal. We assign to the intercept a diffuse prior variance of 10^6. The prior variance is 100 for the monetary-policy shock coefficients, and the regime and state terms, 1 for the lagged macroeconomic controls. The baseline specification uses Minnesota-type shrinkage, with overall tightness $\lambda = 0.5$, cross-variable tightness 0.5, state-variable tightness 0.5, monetary-policy-shock tightness 1, lag-decay parameter 1, and scaling by regressor variances.

Given a candidate regime classification, posterior moments are computed using standard conjugate updates. Let V_0^{-1} denote the prior precision. Then $V_h = (V_0^{-1} + X_h{}'X_h)^{-1}$, $\quad \theta_h^{post} = V_h(X_h{}'y_h + V_0^{-1}\theta_0)$, $a_h^{post} = a_0 + T_h/2$, and $b_h^{post} = b_0 + \frac{1}{2}\left(y_h{}'y_h + \theta_0{}'V_0^{-1}\theta_0 - (\theta_h^{post})'V_h^{-1}\theta_h^{post}\right)$. Posterior draws are then drawn from $\sigma_h^{2,(r)} \sim IG\left(a_h^{post}, b_h^{post}\right)$, $\quad \theta_h^{(r)} \sim N\left(\theta_h^{post}, \sigma_h^{2,(r)} V_h\right)$.

The uncertainty threshold is selected endogenously. For each factor, candidate thresholds are defined as empirical percentiles, $\gamma_g = Q_f(g), with\ g \in \{20, 25, \ldots, 80\}$. For each candidate threshold and delay, the high-uncertainty indicator $S_t(\gamma_g, d)$ is constructed and the full set of local projections is estimated across both response variables and all horizons. Candidate regimes are retained only if $0.30 \leq \frac{1}{T}\sum_t S_t\left(\gamma_g, d\right) \leq 0.70$. Among admissible candidates, the selected threshold maximizes the penalized log marginal likelihood, $(\hat{\gamma}, \hat{d}) = \arg\max_{\gamma_g, d}\left\{\log p\left(Y \mid \gamma_g, d\right) + \mathcal{P}\left(\bar{S}_{\gamma_g,d}\right)\right\}$, where $\bar{S}_{\gamma_g,d} = T^{-1}\sum_t S_t\left(\gamma_g, d\right)$. The penalty term $\mathcal{P}\left(\bar{S}_{\gamma_g,d}\right)$ discourages unbalanced regime splits by favoring a high-uncertainty share close to one half of the sample. The baseline delay is fixed at d=1.[9] If no admissible threshold is found, the fallback rule uses the 75th percentile threshold with the baseline delay. Threshold uncertainty is incorporated directly into the posterior simulation. After evaluating the threshold grid, marginal likelihoods are converted into posterior model probabilities over admissible threshold candidates, $\pi_g = \frac{\exp\left(\ell_g - \max_j \ell_j\right)}{\sum_j \exp\left(\ell_j - \max_k \ell_k\right)}$, where ℓ_g denotes the penalized log marginal likelihood for candidate g. Posterior draws then jointly sample over regression coefficients, thresholds and factor paths. Factor-path uncertainty is incorporated using the bootstrap factor paths generated in Appendix C: up to 500 bootstrap paths are matched to the local-projection sample and combined with the baseline factor path. Each posterior draw first selects a mixture component (b, g), defined by factor path

[9] In unreported results, we also allow for a search over $d \in \{0,1,2,3\}$, and the resulting optimal thresholds are identical to our baseline specification.

b and threshold candidate g, with probability $p(b,g) = p(b)\pi_g$. Conditional on that component, the high-uncertainty indicator is rebuilt, posterior moments are recomputed, the residual variance is drawn from its inverse-gamma posterior, and coefficients are drawn from the associated normal posterior.

For each response, horizon and uncertainty factor, the algorithm generates 5,000 posterior draws and discards the first 1,000 as burn-in. The retained draws are used to compute posterior medians and 68% and 90% credible intervals for the low- and high-regime responses. Since the posterior integrates over coefficient uncertainty, factor-path uncertainty and threshold uncertainty, the reported credible sets reflect uncertainty both in the transmission estimates and in the latent regime classification.

As to the Bayesian Model Averaging (BMA) across models, let ℓ_s denote the best admissible log marginal likelihood for state specification $s \in \{FIN, MACRO, GEO\}$. Model weights are computed as $w_s = \frac{\exp\left\{\left(\ell_s - \max_j \ell_j\right)/\tau\right\}}{\sum_j \exp\left\{\left(\ell_j - \max_k \ell_k\right)/\tau\right\}}, where\ \tau = 500.$ BMA impulse responses are obtained by averaging posterior draws across the three uncertainty models, $\Phi^r_{h,BMA} = \sum_s w_s\, \Phi^r_{h,s}$. This produces the aggregate low- and high-uncertainty responses shown in the main text, while preserving the contribution of each uncertainty dimension to the final posterior distribution. In our baseline simulations, $w_{FIN} = 0.306$, $w_{MACRP} = 0.259$ and $w_{GEO} = 0.435$.

E. Additional Tables and Figures

Table E1. Descriptive statistics of output growth and inflation by uncertainty regime

Uncertainty factor	Output growth		Inflation	
	Low Unc.	**High Unc.**	**Low Unc.**	**High Unc.**
FIN	2.23 (7.28) [107]	-0.27 (5.78) [161]	1.74 (0.92) [107]	2.31 (2.27) [161]
MACRO	1.72 (6.16) [147]	-0.49 (6.76) [121]	1.71 (1.41) [147]	2.54 (2.23) [121]
GEO	0.87 (5.43) [81]	0.66 (6.96) [187]	1.67 (1.42) [81]	2.27 (2.01) [187]

Note: The table reports regime-specific means. Standard deviations are shown in parentheses and numbers of observations in square brackets. "Financial," "Macroeconomic," and "Geopolitical" refer to the uncertainty factor used to classify observations into low- and high-uncertainty regimes.

Table E2. IRFs statistics for a 25bps monetary policy shock

Model	Statistic	Output growth		Inflation	
		Low Unc.	**High Unc.**	**Low Unc.**	**High Unc.**
FIN	(1)	-0.50	-	-0.09	-
	(2)	-1.22	-	-0.34	-
MACRO	(1)	-0.31	-	-	-
	(2)	-0.58	-	-0.09	-0.06
GEO	(1)	-0.67	-	-0.12	-
	(2)	-1.65	-	-0.32	-0.07
Total	(1)	-0.37	-	-0.07	-
	(2)	-1.27	-0.19	-0.27	-0.05
Baseline	(1)	-		-	
	(2)	-0.20		-0.05	

Notes: (1) peak response; (2) cumulated response. A '-' indicates that coefficients are not significant coefficients at the 68% confidence level. Total indicates the overall IRFs obtained via BMA across the three models. Baseline row reports full-sample local-projection estimates over the entire sample, using the same monetary-policy shock, controls, lag structure, and horizon as the uncertainty-dependent specifications. Coefficients are in percentage points.

Figure E1. Monetary policy target shocks

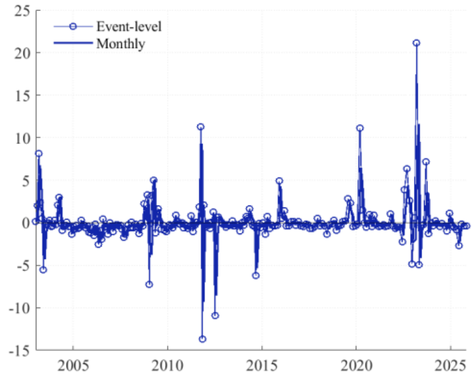

Notes: The markers show the event-level target shocks, while the solid line shows their monthly aggregation, which is used as the monetary policy shock in the local-projection analysis. The sample covers January 2003 to November 2025. The horizontal line at zero separates expansionary from contractionary shocks.