Macroprudential Commentaries

Issue No: 9
November 2025

ESRB
European Systemic Risk Board
European System of Financial Supervision

Navigating tail risks: assessing euro area economic growth and equity market vulnerabilities

*Frederik Andersen, Magnus Andersson, Stephen Cecchetti, Raffaele Giuliana, Ugo Pelca, Jonathan Rice, Cecilia Sarchi and Arthur de Vries**

Technical annex: summary of the approaches adopted in the commentary

This technical annex outlines the analytical methodologies employed in "Navigating tail risks: assessing euro area economic growth and equity market vulnerabilities", a commentary published on 9 December 2025. The analysis draws on two complementary approaches to constructing one-year-ahead real GDP growth distributions and uses market-implied information to derive corresponding distributions for equity returns. This annex provides a brief description of the underlying techniques, assumptions and data sources supporting the commentary's assessment of macro-financial tail risks.

Specifically, this annex documents the following methodologies:

a. *Growth-at-Risk*
b. *Survey-based distribution from the ECB Survey of Professional Forecasters*
c. *Implied probability density functions for equity returns*

a. Growth-at-risk

The growth-at-risk (GaR) framework is part of the ESRB's toolkit, providing quantitative metrics that can be used by policymakers to assess the adequacy of policies against risks and resilience. Developed by the ESRB's Expert Group on Macroprudential Stance (2021) and the ESRB's Contact Group on Macroprudential Stance (2024), GaR employs panel quantile regressions to forecast the distribution of real GDP growth, focusing on the left tail to illustrate downside risks. This approach offers a parsimonious and robust method that can be used to estimate the differentiated effects of financial conditions, systemic risks and resilience factors on GDP growth, covering short-term and medium-term

* ESRB Secretariat.
+ Acknowledgement: This analysis has benefited from valuable input and comments provided by Stephan Fahr, Mikael Hagelstam, Lorenzo Mangiola, Francesco Mazzaferro, Tuomas Peltonen and Frauke Skudelny. The illustration on page 2 was created by Inka Cordero Garcia.

horizons to capture the materialisation of risks as well as the build-up of vulnerabilities. While economic growth is not a direct objective of macroprudential policy, the absence of financial stability manifests itself in a higher likelihood of deep recession. GaR thus provides a forward-looking tool that can be used to assess GDP growth risks in the context of an evolving macro-financial environment and to proactively manage financial stability policies.

The specification is an unbalanced quantile panel regression with country fixed effects for euro area countries, estimated from the first quarter of 1999 to the third quarter of 2025. The variables[1] in the ESRB framework include a systemic risk indicator (SRI)[2], a financial stress indicator (CLIFS)[3], a cumulative macroprudential policy index[4], lagged real GDP growth, interaction terms (e.g. SRI x CLIFS) and structural factors (e.g. trade openness and bank sector concentration). Unlike the Expert Group's stance indicator, which focuses on the gap between the 50th and the 10th percentiles of GDP growth, the approach followed in this commentary note examines the full distribution and employs an adjusted explanatory variable set. The debt service ratio and the European Commission's economic sentiment indicator are incorporated, with the latter shown to improve the quantification of downside risks relative to median predictions (see Lang et al., 2022). SRI, CLIFS and lag terms are retained, while variables measuring capital, policies and structural factors are excluded. The estimation embeds country fixed effects to account for time-invariant heterogeneity and evaluates results across a range of quantiles, with data standardised to ensure greater stability. Notably, euro area-level outcomes are constructed ex post as a GDP-weighted aggregation of estimated country-level contributions. The estimation emphasises near-real-time monitoring with a forecasting horizon of four quarters. For further details, please refer to the specification and notation used in the exercise.

Country-level quantile regression

$$Q_{\mathcal{T}}(y_{i,t+h}|F_t) = a_{\mathcal{T}} + FE_{i,\mathcal{T}} + \beta_{\mathcal{T}}^{GDP} GDP_{i,t-h} + \beta_{\mathcal{T}}^{SRI} SRI_{i,t} + \beta_{\mathcal{T}}^{CLIFS} CLIFS_{i,t} + \beta_{\mathcal{T}}^{ESI} ESI_{i,t} + \beta_{\mathcal{T}}^{DSR} DSR_{i,t}$$

EU-level quantile regression

$$\begin{aligned} Q_{\mathcal{T}}(y_t^{EU}|F_t) = a_{\mathcal{T}} &+ \sum_{i=1}^{N} \omega_{i,t}\, FE_{i,\mathcal{T}} + \beta_{\mathcal{T}}^{GDP} \sum_{i=1}^{N} \omega_{i,t}\, GDP_{i,t-h} + \beta_{\mathcal{T}}^{SRI} \sum_{i=1}^{N} \omega_{i,t}\, SRI_{i,t} + \beta_{\mathcal{T}}^{CLIFS} \sum_{i=1}^{N} \omega_{i,t}\, CLIFS_{i,t} \\ &+ \beta_{\mathcal{T}}^{ESI} \sum_{i=1}^{N} \omega_{i,t}\, ESI_{i,t} + \beta_{\mathcal{T}}^{DSR} \sum_{i=1}^{N} \omega_{i,t}\, DSR_{i,t} \end{aligned}$$

where:

[1] The chosen indicators must be available for a wide range of countries over a long period of time and must score a good prediction performance for the GDP left tail.

[2] Weighted average of indicators with early warning property for financial crises: change in bank credit/GDP ratio, current account/GDP ratio, change in residential real estate price/income ratio, growth rate of real equity price, change in debt service ratio and growth rate of real total credit. SRI is from Lang et al. (2019).

[3] Based on three key market segments: equity (stock price index), bond (ten-year government yields) and foreign exchange markets (real effective exchange rate). CLIFS is from Duprey et al. (2017).

[4] From the Macroprudential Policies Evaluation Database (MaPPED; Budnik and Kleibl, 2018) and the ESRB's Notifications Database, cumulating the net change of dummies on policy direction: +1/-1 for tightening/loosening.

$Q_\tau(.|F_t)$: τ – th conditional quantile of the outcome, given the information set F_t at time t.

y_t^{EU}: EU – level oucome variable at time t.

$y_{i,t}$: Country – level outcome variable for country i at time t.

a_τ : Intercept term, specific to quantile τ but not country – specific.

$FE_{i,\tau}$: Country fixed effect for country i, specific to quantile τ.

$\omega_{i,t}$: GDP weight of country i at time t.

In this commentary note, the specification outlined above is augmented with a geopolitical risk indicator to quantify systemic risks arising from geopolitical uncertainty. This approach aligns with the ECB/ESRB workstream on financial stability risks from geoeconomic fragmentation (report forthcoming in 2026). In particular, we build on the insights from this workstream and use the Economic Policy Uncertainty (EPU) index, which is identified as a promising indicator based on multiple quantitative findings illustrated in the ECB/ESRB (2026) report.

b. Survey of Professional Forecasters

Methodology

The analysis of the ECB's Survey of Professional Forecasters (SPF) constructs probability distributions for expected real GDP growth using individual forecaster data from the survey. The SPF is conducted quarterly and collects information on expected rates of inflation, real GDP growth and unemployment in the euro area at multiple forecast horizons, both as point estimates and as probability distributions. For this exercise, we focus on the "one-year-ahead" real GDP forecasts published in the third quarter of 2024 and the first three quarters of 2025. Each SPF respondent reports probabilities of year-on-year changes in real GDP materialising within a fixed set of outcome intervals (bins). In the periods considered, these bins span 0.5 percentage-point intervals between -1% and +4%, with open-ended tails below -1% and above +4%.

Only respondents providing a complete set of probabilities across the bins are included, while those providing point forecasts alone are excluded. For the periods analysed, this corresponds to 34, 33, 32, and 30 forecasters respectively. To obtain the consensus forecast of the expected distribution of future GDP growth, we adopt a non-parametric averaging approach. Each respondent's probability distribution is treated as a discrete probability vector, with vectors then averaged across all qualifying forecasters to produce a pooled histogram representing the panel's mean subjective distribution. Hence, each bin's mean probability reflects the proportion of total probability mass that the panel has assigned to that interval.

To transform the pooled discrete histogram into a continuous Probability Density Function (PDF), we apply kernel smoothing. Each bin is represented by a Gaussian kernel centred at the midpoint of its interval; the lower-tail kernel is centred at -1.5% and the upper-tail kernel at +4.5%. The standard deviation of each kernel is proportional to the bin width. The individual kernels are weighted by their mean probabilities, summed to form a continuous density and normalised so that the total area under the curve equals 1. The resulting smoothed PDFs aim to depict the SPF panel's collective expectations (and perceived uncertainty) about the next 12 months' GDP growth.

Smoothing assumptions/data issues

The kernel-based smoothing procedure aims to produce interpretable PDFs from the discrete data. Bandwidths are set at 70% of each bin's width for regular intervals and a fixed 0.7 percentage points for the open-ended tails. A secondary Gaussian smoothing filter is applied to the aggregated density to mitigate small discontinuities between bins. This procedure generates smoother distributions but may slightly redistribute probability near bin edges, as some mass from the open-ended tails spreads into adjacent higher (or lower) growth bins. This effect is rather minor.

It is important to note is that the SPF is effectively based on information available at the start of each quarter, as the survey is conducted at that time. Consequently, the reported point estimate and distribution for a given quarter largely reflect conditions at the beginning of that quarter rather than the quarter as a whole. This timing effect can create an apparent lag in the SPF data (which helps explain why the point estimate for the third quarter of 2025 is substantially lower than the value for the second quarter of 2025, as the latter captures expectations formed at the very start of the second quarter).

Some data-related limitations are also worth noting. Firstly, the SPF data are reported in relatively wide discrete bins, which may constrain the precision of individual forecast densities. Secondly, the open-ended tails require assumptions to be made as to how to centre and spread the tail probabilities (set at -1.5% and +4.5% in this exercise). Thirdly, we decided to manually exclude forecaster 108 from the third quarter of 2024 one-year-ahead real GDP forecast, as this submission appeared unrealistically low and likely the result of an error. Finally, as discussed with the survey organisers, there is strong heterogeneity in reporting practices across respondents. For example, some respondents tend to assign rounded probabilities to a few central bins, while others rely on model-based procedures and distribute probabilities more finely across the available bins.

c. Estimating probability density functions of one-year-ahead EURO STOXX 50 index returns

Our approach aims to extract the risk-neutral probability density of the index level at a one-year horizon and then express it as a distribution of percentage returns from the spot price at each evaluation date.

For each evaluation date, the inputs consist of: the spot index level S_0, the corresponding one-year-forward price F and an option smile (strikes K_i with associated implied volatilities σ_i, with an expiry exactly one year ahead). All data points are retrieved from Bloomberg. The strike levels are equivalent to the 5-delta, 10-delta, 15-delta, 25-delta, 35-delta, 40-delta, 45-delta, 47.5-delta and 50-delta calls and puts.

Implied volatilities are converted to total variance $w_i = \sigma_i^2 T$, where T is the time to maturity.

The strike domain is generated by a log-moneyness grid $k \epsilon [Fe^{-0.8}, Fe^{0.8}]$ and evaluated on a dense grid.

The smile is smoothed by fitting the raw SVI parameterisation of total variance (Gatheral, 2004):

$$w(k) = a + b[\rho(k-m) + \sqrt{(k-m)^2 + \sigma^2}],$$

where parameters (a, b, ρ, m, σ) control the level, slope, skewness and curvature of the smile. These parameters are estimated by minimising the sum of squared errors between the market total variances and those produced by the SVI function, using a Nelder-Mead numerical optimisation routine. The resulting fit provides a smooth estimate of implied volatility across strikes.

Under the Black model, the price of a European call option on a forward F with strike K, volatility σ, maturity T, and risk-free rate r is given by

$$C(F, K, T, r, \sigma) = e^{-rT}[F\Phi(d_1) - K\Phi(d_2)],$$

where $d_1 = \left[\ln\left(\frac{F}{K}\right) + 0.5\sigma^2 T\right] / (\sigma\sqrt{T})$ and $d_2 = d_1 - \sigma\sqrt{T}$.

This pricing function is used to compute call prices from the fitted implied volatilities and can be applied across the strike space since the smile is put-call parity consistent.

Once the volatility smile is fitted and the call price curve is obtained, the risk-neutral PDF is recovered using the Breeden-Litzenberger identity, which states that under no-arbitrage conditions the discounted second derivative of the call price with respect to the strike price yields the risk-neutral density in price space:

$$f(K) = e^{rT} \frac{\partial^2 C(F, K, T, r, \sigma(K))}{\partial K^2}.$$

In practice, we compute call prices for a fine grid of strikes using the Black formula (as shown above) and differentiate these numerically with a cubic spline interpolation. The resulting density is clipped to ensure non-negative values and then normalised so that the total probability integrates to one. To interpret the results in terms of percentage returns rather than absolute prices, the PDF over strike prices $f(K)$ is converted to a PDF of returns from spot $R = \left(\frac{K}{S_0} - 1\right) * 100$.

The Jacobian $\frac{dK}{dR} = \frac{S_0}{100}$ implies $f_R(R) = f_S(K) * \frac{S_0}{100}$. A second renormalisation is applied in return space to remove any residual numerical drift. From f_R we can compute a range of metrics such as the standard derivation of R and tail probabilities by integrating over the corresponding region.

References

Budnik, K. and Kleibl, J. (2018), "Macroprudential regulation in the European Union in 1995-2014: introducing a new data set on policy actions of a macroprudential nature", *Working Paper Series*, No 2123, ECB.

Duprey, T., Klaus, B. and Peltonen, T. (2017), "Dating systemic financial stress episodes in the EU countries", *Journal of Financial Stability*, Vol. 32, Issue C, pp. 30-56.

ESRB (2021), "Report of the Expert Group on Macroprudential Stance – Phase II (implementation)" European Systemic Risk Board, December.

ESRB (2024), "Improvements to the ESRB macroprudential stance framework", European Systemic Risk Board, January.

ECB & ESRB (2026), "Financial stability risks from geoeconomic fragmentation", European Central Bank and European Systemic Risk Board, *Occasional Paper Series*, January (forthcoming).

Gatheral, J. (2004), "A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives", presentation at Global Derivatives & Risk Management Conference Madrid, May.

Lang, J. H., Izzo, C., Fahr, S. and Ruzicka, J. (2019), "Anticipating the bust: a new cyclical systemic risk indicator to assess the likelihood and severity of financial crises", *Occasional Paper Series*, No 219, ECB.