MASTERARBEIT | MASTER'S THESIS

Titel | Title

Belief Shocks and Growth-at-Risk in Eurozone
The Role of Expectations in Macro-Financial Risk

verfasst von | submitted by

lic. Kacper Sebastian Skura

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 913
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Applied Economics
Betreut von | Supervisor:	Univ.-Prof. Mag. Mag. Dr. Harald Fadinger Privatdoz.

Kurzfassung

Diese Arbeit untersucht, wie sogenannte *Belief Shocks* – definiert als systematische Prognosefehler in professionellen Erwartungen – die makroökonomischen Ergebnisse im Euroraum beeinflussen. Aufbauend auf Enders et al. (2021) und Boeck and Pfarrhofer (2025) erweitere ich den Growth-at-Risk-(GaR)-Rahmen über US-Daten hinaus, indem sowohl Output- als auch Inflationserwartungen aus der *ECB Survey of Professional Forecasters* einbezogen werden. Anschließend identifiziere ich mithilfe eines strukturellen VAR-Modells und einer Bayes'schen zeitvariierenden Quantilregression *Belief*- und *Non-Belief-Shocks* und untersuche deren Einfluss auf die Verteilung des Wirtschaftswachstums.

Die Ergebnisse zeigen, dass *Belief-Shocks* etwa die Hälfte der Output-Schwankungen auf dem Einjahreshorizont erklären, allerdings mit einer ausgeprägten Regimeabhängigkeit. In ruhigeren Phasen dominieren fundamentale Faktoren, während in Krisenzeiten die Bedeutung von Erwartungen zunimmt und ihr Beitrag auf bis zu 70 % ansteigt. Im Gegensatz dazu weisen Inflations-*Belief-Shocks* eine bemerkenswerte Stabilität auf, wobei *Belief*- und *Non-Belief-Shocks* über die Teilstichproben hinweg nahezu gleich stark beitragen. Darüber hinaus zeigt sich, dass der GaR-Rahmen die Analyse der realwirtschaftlichen Auswirkungen lediglich nuanciert, anstatt sie grundlegend zu verändern.

Abstract

This thesis examines how belief shocks, defined as systematic forecast errors in professional expectations, shape macroeconomic outcomes in the euro area. Building on Enders et al. (2021) and Boeck and Pfarrhofer (2025), I extend the Growth-at-Risk (GaR) framework beyond U.S. data by incorporating both output and inflation expectations from the ECB Survey of Professional Forecasters. Then I use structural VAR and Bayesian time-varying quantile regression to identify belief and non-belief shocks and examine their effect on the distribution of growth.

The results show that belief shock accounts for roughly half of output fluctuations at the one-year horizon, but with stronger regime dependence. During calmer periods, fundamentals dominate, while crises dominated times amplify the role of beliefs, raising their contribution to as much as 70%. By contrast, inflation belief shocks display remarkable stability, with belief and non-belief shocks contributing almost equally across subsamples. Additionally, proving that the GaR framework adds only nuance to the examination of the real impact.

Acknowledgements

I would like to express my deepest gratitude to Professor Harald Fadinger for his guidance, constructive advice, and patient support throughout the writing of this thesis.

I am also deeply thankful to my family for their constant support during my studies, and to my girlfriend, Barbara Podstawczyńska, for her endless motivation and belief in me.

Contents

List of Graphs

List of Tables

1 Introduction

In the landscape of constant changes in macro-financial world, the fundamental aspect necessary to achieve a sustainable growth and investments is managing risks. Value-at-Risk (VaR) became increasingly popular in the financial world as it provides an answer to one simple question "How bad can things get?". The answer for this at a given time may be "There is a 5% chance this portfolio will lose more than $10 million over the next day". However, the same question can be stated for the economy in general, where one would ask "What is the worst expected GDP growth over the next year, given today's macro conditions?" the framework that provides a key to understanding this is Growth-at-Risk (GaR). The approach that links current macro-financial conditions with the probability distribution of future real GDP growth was established by Adrian et al. (2019). They showed that specifically the left tail of the distribution is positively correlated with a decrease in financial conditions. More specifically, for US data, deterioration of macro-financial variables significantly lowers the left tail in the short-run (1-year) real GDP growth. As argued by Adrian et al. (2019), this rationalizes the notion pushed by Peek et al. (2016), that the FOMC's is increasingly interested in financial conditions.

There is a strong rationale that policymakers can only benefit significantly from having information about the entire growth distribution. As explained by Prasad et al. (2019), one can imagine an inflation-targeting central bank that adjusts interest rates to meet its inflation goals. Given that policymakers communicate risks surrounding their point forecasts, which are central scenarios, however with an entire distribution at hand, they may choose the policy stance that minimizes risks rather than optimizing central scenario outcome. Growth-at-Risk in this case would allow policymakers to quantify the risk scenarios, which would also help communicate the nature of risks. As expressed further by Adrian et al. (2022), the GaR measure could be a promise of a practical tool for macroprudential policy design, as it offers an objective measure of downside risks to expected growth, informing both the need for and the effectiveness of policy interventions. In particular, it is aimed at enhancing financial system resilience.

Theoretical work such as Kozeniauskas et al. (2016) links uncertainty to belief formation, showing how disagreements about future conditions can amplify macroeconomic fluctuations through uncertainty and misperception of aggregate risk. Building on similar idea, Enders et al. (2021) provide empirical evidence that systematic forecast errors, which are named belief shocks, capture such misperceptions in professional expectations and can explain a significant share of short-run output fluctuations. Using the data from the Survey of Professional Forecasters for the United States, authors show that waves of optimism or pessimism generate strong co-movements across output, investment and consumption, acting as an independent source of business cycle fluctuations. Extending this further, Boeck and Pfarrhofer (2025) embed belief shocks into a Growth-at-Risk framework, setting a connection between expectation-driven fluctuations and the conditional distribution of future growth. Especially investigating the downward risk. Their results confirm that belief shocks show a persistent real effect, however they do not significantly alter downward risk, or is amplified at different percentiles. Both works provide the conceptual and empirical foundations for a further studies on belief shocks.

Although the work of Enders et al. (2021) and Boeck and Pfarrhofer (2025) has focused on GDP beliefs, inflation belief shocks may also play a significant role in shaping macroeconomic risk. Furthermore, these studies focused primarily on the United States. This thesis aims to expand the literature on belief shocks following the methodology of Boeck and Pfarrhofer (2025), incorporating both GDP and HICP inflation expectations. In particular, I construct belief shocks by using nowcast errors from the ECB's Survey of Professional Forecasters. The key question of this thesis is then: "To what extent belief shocks in inflation and output expectations drive macroeconomic fluctuations in the Eurozone?". To do so, I follow the TVP-Quantile Regression and Bayesian SVAR identification strategies, but also explore the simplified model, where only median of nowcasts is taken.

The analysis carried out yields noteworthy findings on how belief shocks shape economic fluctuations in the Eurozone. The results show that belief shocks play a significant role in short-run output dynamics, accounting for roughly half of the variance in forecast error over one-year. This importance increases sharply during periods of elevated uncertainty, as during crises such as financial, pandemic, and energy shocks, belief can explain up to 70% of the variation of output. Whereas, in calmer times, fundamentals dominates. By contrast, inflation seems to be remarkably stable, with belief and non-belief shocks contributing almost equally across time subsamples. These findings reveal that the transmission of belief shocks is dependent on the regime, as during times the importance is amplified while remaining balanced otherwise. Additionally, the Growth-at-Risk framework refines the findings and proves them further, rather than overturning them.

The remainder of this thesis is structured as follows. Section 2 examines the relevant literature and theoretical background. Section 3 introduces the data, Section 4 outlines the methodology, Section 5 presents and interprets the results, whereas Section 6 discusses the robustness of study and limitations. The conclusion with a discussion of the implications of the policy can be found in Section 7.

2 Literature Review

The Growth-at-Risk framework provides a link between financial conditions and the future output growth distribution. Existing research has developed many approaches across countries, horizons, and indicators, providing a broad understanding of how macro-financial conditions shape risks. This section first reviews in more detail the literature on GaR, to turn to the role of beliefs, expectation formation, and uncertainty and their impact on business cycles.

Building on Adrian et al. (2019), Adrian et al. (2022) use panel estimates for 11 advanced economies and 10 emerging economies that financial conditions have strong forecasting power for the distribution, not only its mean. Interestingly, signs of the coefficients may alter depending on the short- to medium-term horizons, in particular fragile financial conditions imply higher GaR in the near term, however for the medium term, they reverse, implying higher downside risk. This trade-off in short implies that loosening conditions may help today, but increases future vulnerability.

The idea of the medium term was developed further by Lang et al. (2023). They studied how much information can different financial stability indicators provide for GaR on the basis of the euro area, implementing also varying projection horizons. The distinction between financial stress indicators (spreads and market volatility) and financial vulnerability indicators (credit and asset price imbalances), allowed the authors to find that both of them contain information in the short-term for Growth-at-Risk. However, at horizons of three years or more, only vulnerability indicators seems to contain any information for GaR.

The financial conditions may prove to provide a better signal of negative shocks to the economy. Thus, the models focusing on Bayesian learning and its impact on the economy should evolve to incorporate the formation of beliefs on the basis of macro-financial conditions (Adrian et al., 2019). In their work Kozeniauskas et al. (2016) give an explanation of three types of uncertainty: macro - when agents believe the entire economy becomes less predictable; micro - an increase in the variance of idiosyncratic shocks to firms; higher-order - uncertainty about others' beliefs. All of them strongly co-move, due to the disaster risk. Since such distress is rare, it is difficult to assess properly and the scope to disagree about it is large. Thus, when such risk comes to fruition, it decreases the accuracy of forecasts raising macro uncertainty, the higher variance in forecasts pushes higher-order uncertainty, and it makes firms respond differently, creating micro uncertainty. This domino process is related to disaster risk, which may be evaluated using the aforementioned GaR.

Whereas Kozeniauskas et al. (2016) provides a theoretical foundation for the role of beliefs and the uncertainty they generate, the work of Enders et al. (2021) offers an empirical validation. Belief shocks are measured by systematic forecast errors, that is, forecast errors in expectations about current or future output growth, in particular, they measure whether agents are overly optimistic or pessimistic compared to actual realisation. Thus,

$$\text{Nowcast Error} = \text{Realised Output - Forecasted Output}$$

if the value is positive, the forecasters were too pessimistic, conversely, if negative, they were too optimistic. Enders et al. (2021) used a Survey of Professional Forecasters (SPF) and compared the median forecast of contemporaneous output with the realised value.

The idea behind this statement is that if forecasters systematically miss what is already happening, it reflects misperception, making it a belief shock proxy. As the errors may be contaminated by other structural shocks (e.g. TFP, monetary, fiscal shocks), the authors decided to purify it in two distinct ways. Firstly, regressed NE on a wide set of known structural shocks, where the residuals are interpreted as purified nowcast error. The second method utilised sign-restricted vector autoregression, imposing structural assumptions that belief shocks cause output and nowcast errors to move in opposite directions and non-belief shocks (TFP or demand) cause to move together in the same direction. The restriction originates from the idea that overly optimistic agents expect higher growth, but actual growth is lower, inducing negative co-movement. Conversely, favourable belief shocks may cause perceived growth to overshoot actual (a negative nowcast error), boosting economic activity. The authors found that the belief shocks generate strong co-movement across output, investment, consumption, and hours worked, with persistent and significant drops in GDP. For a 1-year horizon, about one-third of output fluctuations can be explained using belief shocks, and for a 2-year horizon up to 40%. This was proven by evaluating the Great Recession, where belief shocks contributed more than structural shocks.

Building on this foundation, Boeck and Pfarrhofer (2025) extend the work of Enders et al. (2021) while adopting the same identification strategy for belief shocks, but depart from the conventional focus on point forecasts. They extended the analysis by embedding belief shocks into a Growth-at-Risk framework. Despite this quantile-based extension, the results show that belief shocks identified via different parts of the predictive distribution yield statistically similar impulse responses, indicating that for the U.S. the mean-based approach may be sufficient to capture the business cycle impacts. However, their methodology still remains valuable, as they ensemble the entire distribution of both SPF forecasts and growth outcomes, the latter are estimated using Bayesian Time-Varying Parameter Quantile Regression (TVP-QR) developed by Pfarrhofer (2022). Collectively, these studies show that expectations-driven shocks can impact the economy in persistent and significant way.

The relevance of belief shocks is also deeply grounded in the New Keynesian literature. It embeds expectation formation and information frictions at the core of macroeconomic fluctuations. Lorenzoni (2009) shows that when agents rely on noisy public signals about productivity, they make expectational mistakes, producing ”news shocks” that temporarily increase output, employment, and inflation, in a manner similar to demand shocks. Milani (2011) extends this reasoning in an estimated NK model and finds that expectational shocks, which are formed through waves of optimism or pessimism (not connected to the fundamentals), account for nearly half of U.S. business cycle fluctuations. Empirically, Blanchard et al. (2013) find that distinguishing news from noise is in and of itself hard in VARs, yet they find that noise shocks explain a sizeable fraction of short-run consumption fluctuations.

Recent work on the New Keynesian framework combines information frictions with price rigidities. Hellwig and Venkateswaran (2024) show that with dispersed information, flexible prices give a ”Hayekian benchmark”, indicating that imperfect information alone is irrelevant. However, once prices are sticky and there are strategic complementarities or persistence differences across shocks, informational and nominal frictions interact, leading to slow price adjustment and amplified real effects of nominal disturbances (such as belief shocks). This idea fits into the setting of this thesis, as belief shocks shift demand in the short-run and nominal rigidities make these effects persist. Angeletos and La'O (2013) supplement this view by showing that ”sentiment” shocks can influence cycles of optimism and pessimism in models with rational expectation and a unique equilibrium. Altogether, the works reinforce the view that in the New Keynesian framework, beliefs should be taken into account, and be viewed as an integral part of the economy, with measurable effects on output and inflation.

While works so far have focused solely on GDP beliefs, inflation belief shocks may also play a significant role in shaping macroeconomic risk. The debate on the impact of inflation on economic growth is still not resolved (Gillman and Kejak, 2005), the general agreement is that inflation has a globally negative impact on growth (excluding short-term) (Barro, 2001; Chari, 1996). However, the relationship is not linear. The effect is negative or insignificant between inflation and growth when inflation levels are high, but it changes to be positive for low levels of inflation (Ghosh and Phillips, 1998; Fischer, 1993). The threshold level was examined by López-Villavicencio and Mignon (2011), and the authors found that it differs strongly among advanced and developing countries. For the industrialised economies the estimate is about 2.7%, with a positive inflation-growth link at around 3%. However, for the emerging ones the threshold was estimated to be 17.5%, with the inflation-growth link non-significant below this level. The effects of inflation on growth are mixed, as it depends on how money is introduced in the model. Whether money is a substitute for capital or required for purchasing capital goods the situations differ. In the former scenario monetary growth boosts capital accumulation, thus inflation has a positive effect on long-run growth (Tobin, 1965). In the latter, Stockman (1981) assumes money is needed to carry out investment. Here, higher inflation may reduce the real value of money balances, limiting the ability of firms to invest, thus lowering growth.

In endogenous growth frameworks, economic growth is driven by internal factors - such as capital accumulation - rather than by external. In the AK model, where emphasis is put on physical capital, marginal product of capital is constant ($Y = AK$; $MPK = A$). Thus, inflation by eroding the real returns of the capital, acts as a tax that reduces the incentive to invest, thereby reducing long-run growth rate. As opposed to AH models, which focus on human capital to be the factor of growth. When inflation increases, it reduces the return to education and labour effort (reducing the marginal product of human capital), which, in fact, lowers growth (López-Villavicencio and Mignon, 2011).

The mechanisms are also different in general monetary endogenous growth models, such as one introduced by Gillman and Kejak (2005), where they incorporate both physical and human capital with money demand elasticity. They offer an explanation that in the scenario where the effect of inflation on growth is non-linear, while inflation is at lower levels, the economy is more sensitive to inflation due to inelasticity of money demand. However, when inflation is at higher levels, the substitution of money reduces growth. These theoretical foundations highlight the importance of inflation belief shocks, especially near to the threshold moment. In this light incorporating inflation-based belief shocks into Growth-at-Risk may allow for a better understanding of how misjudgment in inflation expectations may influence the distribution of output growth.

3 Data

The Survey of Professional Forecasters lies at the centre of interest of this paper. The SPF is a dataset maintained by the European Central Bank, it follows a similar structure to a longer running SPF maintained by the Federal Reserve Bank of Philadelphia. Namely, it is a quarterly released dataset, where the respondents of the survey are experts employed by financial or non-financial institutions. It is carried out four times a year: January, April, July, and October. A key distinction between the two lies in their running period. While the U.S. SPF started in 1968, the ECB's SPF only began in 1999, together with the formation of the Euro Area. The shorter time window of the ECB SPF imposes a limitation on the length of available data for empirical analysis in the Eurozone.

Each survey round collects individual point and density forecasts for key indicators, such as real GDP growth, inflation (measured via HICP) and unemployment, across multiple forecast horizons. The focus of this paper is real GDP growth and HICP inflation.

Another difference between the two lies in the forecast horizons. While U.S. SPF includes a true nowcast, which is a forecast for the current quarter before any official data is released, the ECB does not include this. However, the current-year forecasts in the ECB SPF display high dispersion, which suggests it reflects real-time uncertainty about the current state of the economy. Thus, this paper uses it as an approximation of nowcast.

For the usage of this paper the availability of density forecasts in SPF is crucial, as it enables to identify full predictive distribution, and thus allows identification of nowcast errors across quantiles. However, not all respondents submit density forecasts in each round, while providing point forecasts, with probabilities reported in bins for a specific output. This limits the number of consistent responses and making it less suitable for a robust panel structure. Following Boeck and Pfarrhofer (2025), the nowcast distribution will be retrieved from point predictions, as detailed in Section 3.

To evaluate the forecast accuracy and compute nowcast errors, real-time GDP data vintages are used. The vintages are drawn from Eurostat's macroeconomic dataset, and each vintage represents the state of published GDP data at a given release date. These vintages allow us to construct a real-time data series, reflecting the information available to forecasters at the time of the SPF, and a final (or most recent series) for benchmark comparison. In contrast, inflation data (HICP) are not published with historical vintages. Thus, the analysis of inflation forecasts is based solely on the latest available data from Eurostat.

In addition to forecast and vintage data, the empirical framework relies on a constructed panel of macro-financial variables, which are used to develop a set of latent factors that summarise the state of the economy. These factors are then used as covariates in the Time-Varying Parameter Quantile Regression (TVP-QR). Following Boeck and Pfarrhofer (2025), they ensure that key dimensions of economic activity are captured. The full table can be found in Appendix A, Table A. The dataset draws from several institutional sources, including Eurostat, the ECB Statistical Data Warehouse, BIS, FRED and STOXX. The dataset is composed of both monthly and quarterly series, covering credit conditions, financial market volatility, spreads, sentiment indicators, the variables that have been found to be predictive of the tail risk (Lang et al., 2023; Adrian et al., 2019; Clark et al., 2024). Most of them have a starting point of the first quarter of 1999 or before. The variables that are not expressed in quarterly frequency are aggregated either by taking the average value across months (mostly for flow and sentiments variables) or the last value of the quarter (for stock and rate series). Additionally, I use an auxiliary indicator Composite Indicator of Systematic Stress (CISS) which is a Euro equivalent of National Financial Conditions Index (NFCI). It serves to capture episodes of financial stress, which may not be captured by constructed macro-financial factors.

4 Methodology

The baseline framework of this paper is similar to that of Boeck and Pfarrhofer (2025), i.e. construct the nowcast distributions, estimate Growth-at-Risk quantiles using TVP-QR, compute both nowcast errors for the median SPF and for the entire distribution. Lastly, identify belief shocks via a sign-restricted Vector Autoregression.

The first step involves reconstructing nowcast distributions using point forecasts from the SPF. Each forecaster i has a predictive density $x_t^{f,i}$ that has a mean, which in the SPF is presented as a point forecast μ_{it} for a variable x at time t, where $x \in \{\text{GDP Growth, HICP Inflation}\}$. Each of these point forecasts has a variance σ_{it}^2 that is not observable. The forecasts are stored as vectors of values for each forecast round. In a set of N_t of forecasters, predictions vary across the individuals due to heterogenous beliefs, information, or interpretation of conditions in the economy. Thus, the cross-sectional mean, which is the consensus forecast is:

$$\mu_t = N^{-1} \sum_{i=1}^{N_t} \mu_{it} \tag{1}$$

and the cross-sectional variance, which may be interpreted as a disagreement among forecasters is:

$$s_t^2 = \frac{1}{N_t - 1} \sum_{i=1}^{N_t} (\mu_{it} - \mu_t)^2 \tag{2}$$

This disagreement s_t^2 captures the dispersion of forecasters' beliefs and may inform about potential uncertainty as in Kozeniauskas et al. (2016).

Following Boeck and Pfarrhofer (2025), I move from individual point forecasts to a full predictive distribution of x, thus, I construct an ensemble method similar to Krüger and Nolte (2016), as they used a cross-section of economic survey forecasts to predict the distribution of US macro variables in real time. They found that "micro-level" forecasting method based on the individual SPF point forecasts performed well, and, as mentioned, is trivial since it is based on a single parameter. For this Micro-level method the full predictive distribution of x, I model the ensemble distribution as a mixture of Gaussian distributions:

$$x_t^{(f)} = \frac{1}{N_t} \sum_{i=1}^{N_t} \mathcal{N}(\mu_{it}, \sigma^2) \tag{3}$$

Where each individual forecast is treated as the mean of a normal distribution with a common but unknown variance σ^2 to be estimated. The final distribution is the equally weighted average of these Gaussians. The unknown parameter σ^2, that represents the unpredictable part of the variable x, common to all individuals. The estimate for this parameter is recovered by minimising the continuous ranked probability score (CRPS) which evaluates the accuracy of probabilistic forecasts and the observed realization:

$$\mathrm{CRPS}(F, y) = \int_{-\infty}^{\infty} \left(F(z) - 1(z \geq y)\right)^2 \, dz \tag{4}$$

Where F is a CDF and y is a realised value of the variable (in our case actual GDP growth or inflation) and z represents all possible values that the forecasted variable could take (Krueger, 2017). In simpler terms, $1(z \geq y)$ jumps from 0 to 1 at the realization point y. In our context, the predictive distribution F is assumed to follow normal distribution with known mean and unknown variance. The optimal variance $\hat{\sigma}_\tau^2$ is then constructed by minimising the CRPS over a rolling window of $\tau = 12$ quarters, across all survey participants:

$$\hat{\sigma}_\tau^2 = \arg\min_{\sigma^2} \sum_{t=\tau-w}^{\tau-1} \sum_{i=1}^{N_t} \mathrm{CRPS}\left(\mathcal{N}(\mu_{it}, \sigma^2), x_t^{(r)}\right) \tag{5}$$

Where $x_t^{(r)}$ denotes a realisation of the target variable at a given time t, and μ_{it} is the point forecast from a given individual i. This estimation recovers time-varying variances, and forecast uncertainty. The rolling window of 12 quarters implies that the time for our empirical analysis using GaR frameworks shrinks to the period from Q1 1999 to Q1 2002.

With the estimated $\hat{\sigma}_t^2$ for each quarter, one can define the full distribution:

$$x_t^{(f)} = \frac{1}{N_t} \sum_{i=1}^{N_t} \mathcal{N}(\mu_{it}, \hat{\sigma}_t^2) \tag{6}$$

To obtain these SPF distribution quantiles, I simulated 10,000 draws from the mixture distribution for each quarter.

The next step is to recover the true quantiles of our unobserved variables. Each observation of our variables of interest is just a single point, within an ongoing stochastic process. The idea is to estimate the quantiles, and for this I follow Boeck and Pfarrhofer (2025), who used a version of quantile regression that features time-varying parameters (TVPs) by Pfarrhofer (2022). This means that the model allows not only the coefficients to vary across quantiles p, but also their changes over time. Let y_t be a dependent variable, which in my case is either GDP or HICP, x_t be a vector of K predictors and $q_p(x_t) = x_t^\top \beta_{p,t}$ be a pth quantile function. The form of the model is then:

$$y_t = x_t^\top \beta_{p,t} + \varepsilon_t \tag{7}$$

Where the error term satisfies:

$$\int_{-\infty}^{0} f_p(\varepsilon_t)\, d\varepsilon_t = p \tag{8}$$

Equation (8) means that the probability that the error is less than or equal to zero under a given distribution $f_p(\cdot)$ is equal to p. In the TVP-QR according to Pfarrhofer (2022), it is assumed that the error term ε_t follows an asymmetric Laplace distribution. The regression coefficients $\beta_{p,t}$ are not fixed over time, instead, they tend to evolve. The evolution is modelled as a random walk, reflecting that some relationships in the economy may change over time.

In my case the target variable y_t is either GDP growth or HICP inflation, where GDP is taken as the last vintage released before the participants of the Survey were questioned about forecasts. For HICP, due to the lack of vintages, the most recent released values are used. The vector of predictors x_t in my case is a self-constructed dataset of 30 variables, in a way that they not only resemble information a forecaster may possess but are also proven to be both capable of explaining quantiles of growth (Lang et al., 2023; Adrian et al., 2019; Boeck and Pfarrhofer, 2025; Clark et al., 2024) and inflation (Korobilis et al., 2021; Manzan and Zerom, 2013). For the variables, I used a Principal Components Analysis (PCA), which is used to reduce dimensionality from a large set of variables. It finds new axes that capture as much variation in the data, which are the linear combinations of the original variables. It is useful for dataset simplification while preserving most of the information (James et al., 2023). Due to missing data points for quarters, where not every data series was available, the PCA is not a feasible solution. As a standard PCA cannot handle missing observations directly. On the other, hand Bayesian Principal Component Analysis (BPCA) provides a probabilistic framework, that not only estimates factors, but is capable of imputing missing entries. It has proven to be an appropriate tool when the number of predictors is larger and some series contain gaps due to inconsistency in publication frequencies or measurement (Oba et al., 2003). Watson (2002) showed, that extracting four factors is usually sufficient. It was robustly checked by Rapach and Zhou (2018), concluding that the four-factor model performs as well as, or even better than, leading multifactor models. Additionally, following Boeck and Pfarrhofer (2025) further, I include an intercept, lags of the dependent variable $y_{t-1}, \ldots, y_{t-P}$, lags of the auxiliary macro-financial stress indicator (Composite Indicator of Systematic Stress -CISS), specifically using $P = 4$ lags. Thus, the vector of predictors

$$x_t = \left(1, y_{t-1}, \ldots, y_{t-P}, z_{t-1}, \ldots, z_{t-P}, f'_{t-1}, \ldots, f'_{t-P}\right)' \tag{9}$$

where z_t denotes the CISS variable and f_t the $Q = 4$ factors.

The model is built on expanding-window Bayesian framework, so that while estimating it ensures that only information available at a given time t is used. For each forecast time point and quantile p, the distribution of the quantile regressions coefficients β_{pt} is recovered and used to compute the fitted quantile:

$$\hat{y}_{pt}^{(r)} = x'_t \beta_{pt} \tag{10}$$

These are explained as the best available real-time estimate of the conditional quantiles of the variable y.

After assembling the full distribution of growth and nowcasts, following the approach of Boeck and Pfarrhofer (2025), we construct and define nowcast errors (NEs), both using mean-based (Enders et al. (2021) style) and quantile-based measures. For the mean-based approach the sample is full and runs from 1999Q1 to 2024Q1, however, for quantile based, due to the rolling window of CRPS needed to retrieve the distribution of nowcast the running period shrinks to period from 2002Q1 to 20024Q4. Firstly, we can distinguish between different types of realisation: $x_t^{(r)}$ is actual realised value of variable, $\hat{x}_{p,t}^{(r)}$ is an estimated p-th quantile of x from TVP-QR model, $x_t^{(f)}$ is a consensus SPF nowcast (either mean or median), $x_{p,t}^{(f)}$ is a SPF nowcast at quantile p (from ensemble forecast distribution).

Thus, we can define the nowcast errors as:

$$\mathrm{NE}_t = x_t^{(r)} - x_t^{(f)} \tag{11}$$

$$\mathrm{NE}_{p,t} = \hat{x}_{p,t}^{(r)} - x_{p,t}^{(f)} \tag{12}$$

These measures compare the realised outcome to the SPF forecast at both the central and distributional level. NEs tend to exhibit persistence. To ensure any remaining autocorrelation in the series is eliminated, I purge the errors using an ARIMA model with automatic lag selection.

In order to examine the dynamic causal effects of belief shocks I use bivariate Structural Vector Autoregression. Bivariate SVAR can offer a meaningful structural insight, when carefully identified. As an example, it was found successful in estimating the effect of technological shocks in case of the U.S., finding that positive technological shocks generate an increase in real wages, consumption, investment and output in data (Dedola and Neri, 2007). It was also proven to be of use across different time horizons (Fève and Guay, 2016). This approach follows the framework of Boeck and Pfarrhofer (2025), as we denote the vector of endogenous variables as $x_t = \left[\mathrm{NE}_t\ , x_t^{\mathrm{obs}}\right]$, where NE_t is the nowcast error at a given quarter and x_t^{obs} is the observed target variable. Thus, the reduced form of the VAR is:

$$x_t = \sum_{l=1}^{P} A_l x_{t-l} + B d_t + u_t, \quad u_t \sim \mathcal{N}(0, \Sigma) \tag{13}$$

A_l are lag coefficient matrices, $d_t = [1, t, t^2]'$ contains deterministic trends and $u_t = (u_{ne,t}, u_{y,t})$ is the reduced-form errors, which have a normal distribution and zero mean and covariance matrix Σ. The assumption is that reduced-form errors are a linear combination of two structural shocks $\varepsilon_{nb,t}$- non-belief shocks and $\varepsilon_{b,t}$- belief shock. Thus $u_t = A_0 \varepsilon_t$. For this we need to identify the the element A_0, which maps structural shocks, with an identification condition $\Sigma = A_0 A_0'$. To achieve identification, the imposition of sign restrictions on the impulse response NE_t and x_t^{obs} is required.

In this paper, the distinction between GDP and inflation versions of the model is important, as sign restrictions reflect different economic interpretations of shocks: For GDP, the sign restrictions are as follows:

$$\begin{bmatrix} u_{\mathrm{NE},t} \\ u_{x,t} \end{bmatrix} = \begin{bmatrix} + & + \\ + & - \end{bmatrix} \begin{bmatrix} \varepsilon_{\mathrm{nb},t} \\ \varepsilon_{\mathrm{b},t} \end{bmatrix} \tag{14}$$

As explained by Enders et al. (2021), belief shocks arise when agents misjudge the economic situation, meaning that NE_t and GDP need to move in opposite directions. Suppose forecasters overestimated growth $\mathrm{NE}_t < 0$, implying they were too optimistic, expecting higher output. This misperception affects decisions similarly to Kozeniauskas et al. (2016), firms may over-invest, households overspend, causing an outward shift of the demand curve, thus increasing output. Although actual fundamentals do not justify it, the demand increase causes real economic activity to rise temporarily. On the other hand, Non-belief shocks has a positive co-movement. This is a structural shock such as a TFP or policy shock, for example, when the economy experiences a positive productivity shock, forecasters may underestimate ($\mathrm{NE}_t > 0$), but the realised value still increases.

For the inflation version of the model, the restriction logic differs, because of the mechanisms through which inflation affects output.

$$\begin{bmatrix} u_{\mathrm{NE},t} \\ u_{x,t} \end{bmatrix} = \begin{bmatrix} + & - \\ - & - \end{bmatrix} \begin{bmatrix} \varepsilon_{\mathrm{nb},t} \\ \varepsilon_{\mathrm{b},t} \end{bmatrix} \tag{15}$$

Here a positive NE_t indicates that inflation was higher than forecasters expected, whereas a negative indicates overestimation. As discussed by Blinder and Rudd (2013), non-belief shocks in the inflation context may stem from cost-push or supply-side shocks. This is when a commodity price spikes or an energy shock, which drive inflation unexpectedly ($\mathrm{NE}_t > 0$), while at the same time lowering output. This co-movement reflects a possible stagflation. The view that cost-pull shocks (in particular oil price shock) influence the movement in the opposite direction was proven by Bernanke et al. (1997). In contrast, belief shocks arise when agents misjudge inflation to be high but the realised value is lower (or opposite). According to Ha and So (2024), an increase in inflation uncertainty leads to tightening of financial conditions, a decline in investment, and a general slowdown in economic activity. When agents perceive a higher uncertainty about inflation, even without an actual inflation realisation, they adjust expectations and behaviour. Increased uncertainty itself can induce precautionary behaviour, which lowers demand even without the realisation of inflation.

The structural decomposition of inflation is mathematically equivalent if we flip the sign restriction for the belief shock to represent "good inflation". In this understanding, the cost surprise (e.g. energy shocks) are interpreted as "bad inflation", which raises inflation unexpectedly and lowers output, matching the non-belief shock signs. It reflects the classic cost-push inflation as observed during stagflation Bernanke et al. (1997); Primiceri (2005). In this concept, "good inflation" may arise from expansion in demand or loosening financial conditions, thereby raising both inflation and output.

Although belief shocks are structurally represented by negative nowcast errors and a decreasing output ($(-,-)$ pattern), one can flip the sign vector, giving the impulse response functions observationally equivalent to a $(+,+)$ co-movement. This reevaluation is consistent with the theory of demand-driven inflation coinciding with strong real activity. Empirically supported by Andrle et al. (2013), showing that tight co-movement of inflation and output in business cycles in the euro area cannot be precisely explained by cost-push or technology driven models. They argue that demand factors dominate the euro area cycles. Similarly, Del Negro et al. (2020) emphasized that demand shocks, particularly in an environment with a flat Phillips curve, generate a positive correlation between inflation and output gaps. Finally, the cross-country empirical study finds that *"There is a positive and statistically significant relationship between demand-driven inflation and the output gap across 28 countries"* Firat and Hao (2023). In aggregate, the results validate the reinterpretation of belief shocks with originally negative co-movements as reflecting demand-driven "good inflation", where output and inflation may both rise. Without flipping the belief shocks appear as negative co movement, where pessimism or inflation uncertainty lowers demand, but with flipping, forecasters misjudge a demand boom, inducing positive nowcast errors and output.

5 Results

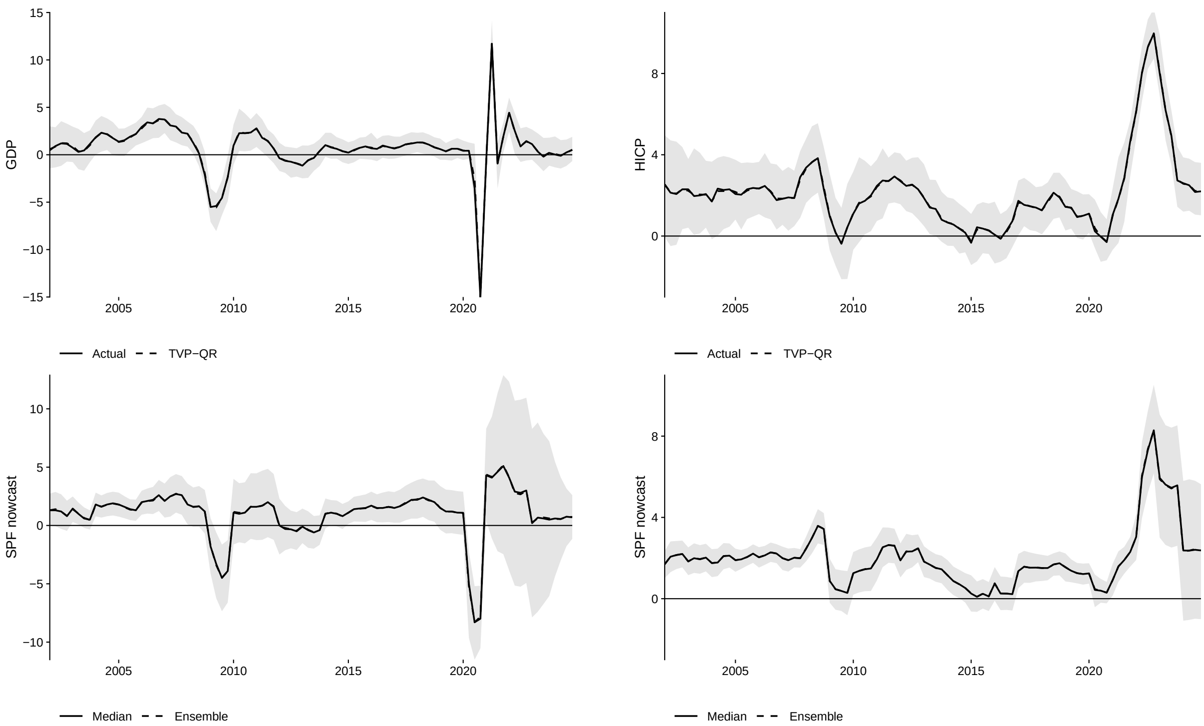

(a) GDP and GDP Nowcast distribution

(b) HICP and HICP Nowcast distribution

Figure 1: Distribution of (a) GDP growth and its nowcast, and (b) HICP and its nowcast.

Notes: Constructed using ensemble forecasts from the ECB SPF and TVP-QR. Shaded areas denote 90% predictive intervals.

Figure 1 presents our first results using fan charts to depict the distribution for both GDP growth and HICP. The top panel of each graph represents the recovered distribution of a variable x using TVP-QR, while the bottom displays the nowcast distribution for the corresponding variable. The shaded areas in the GDP fan chart show increased uncertainty during the Great Recession (2009) and the COVID-19 era (from 2020 onwards), with bands widening considerably. In contrast, inflation uncertainty remained stable until 2021, after which the distribution was drastically increased, especially peaking during the energy price surge 2022-2023. The charts show that the forecast uncertainty is time-varying and particularly sensitive to the two largest crises of the past two decades. The widening of the predictive distribution corresponds to periods where belief shocks are most likely to matter. It is worth noting that while GDP has a smoother pattern response to shocks, HICP inflation shows visible "jumps". This likely reflects the nature of the variables as inflation is dominated in the short run by volatile components (energy and food) and is measured in real time, while GDP is subject to multiple revisions and vintage updates *ex post*. Due to the large spikes observed around and after the COVID-19 era, the visualisation of the full predictive distribution is limited, thus, a reduced chart up to 2020 can be found in the Appendix B.

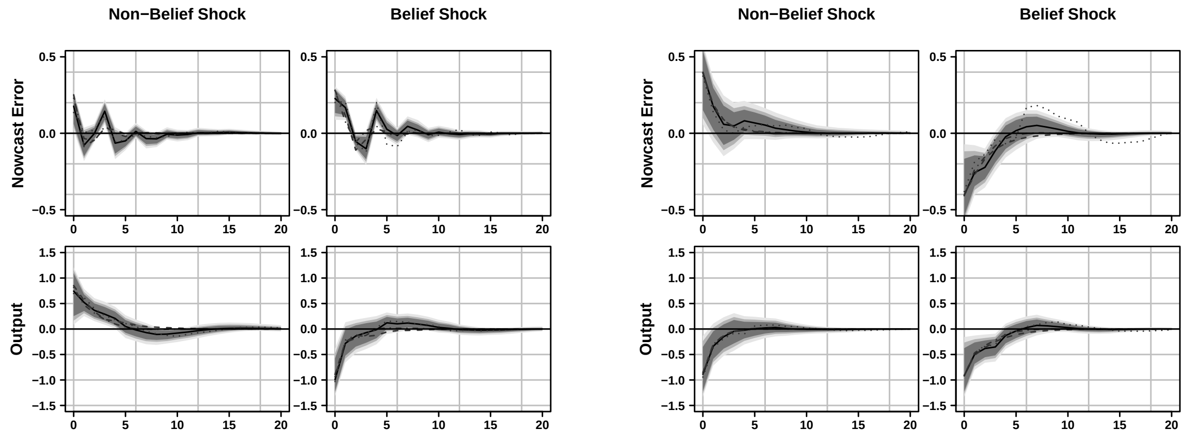

(a) GDP shock (b) HICP shock

Figure 2: Impulse response functions to a non belief shocks for GDP (a) and HICP (b).

The aforementioned descriptive evidence motivates the next step, which is the structural analysis, where I separate the impact of belief and non-belief shocks on GDP and inflation. To accomplish this, I simulated shocks for both GDP and HICP nowcast errors. The results can be seen in Figure 2, where on the left side we can see the shock induced on GDP, and on the right side, the shock for HICP. For each, the black lines are the posterior median responses, while shaded areas represent 68, 80, and 90 percent credible intervals, and dashed lines refer to posterior medians from alternative specifications. Both nowcast errors and output are expressed as deviations from the trend in percentage points.

It is worth noting again that the sign of the output response for both shocks is a direct result of the sign restrictions, which were constructed and explained in the previous chapter. However, the path beyond the initial impact is not restricted, and the persistence of the output decline is an empirical outcome. By construction, GDP belief shocks are defined as having a negative co-movement between nowcast error and output, whereas non-belief shocks induce a positive co-movement in accordance with Enders et al. (2021); Boeck and Pfarrhofer (2025). The path of the output response is hump-shaped. For the non-belief shock, output rises by about 0.5 to 0.7 percentage points above trend in the first quarters, gradually declines back toward zero, even overshooting below trend for a few quarters. This pattern is characteristic for fundamental shocks, where initial gains are gradually absorbed as expectations adjust (Dedola and Neri, 2007). This adjustment is also visible in the top panels, where the nowcast error rapidly alternates between overshooting and undershooting. This impact is moderate but economically significant, as for the slowly growing economies of the euro area, it can provide a meaningful temporary push. On the other hand, output falls by a similar magnitude in the case of belief shocks. Both findings are in line with the U.S. evidence presented by Enders et al. (2021), supporting the view that pessimistic beliefs spread through investment and consumption channels, thereby reducing real output.

Moving to HICP, the initial response is again imposed using the economic theory and empirical evidence outlined in the Methodology part. For non-belief shocks, output initially falls by up to 1 percentage point below the trend, while the NE rises. The effect is sharper, but short-lived, with output recovering after a few quarters and the NE fading just as quickly as expectations adjust. This is in line with the stagflationary nature of energy or commodity price shocks (Blinder and Rudd, 2013). For belief shocks, both NE and output decline, displaying positive co-movement. Output drops by roughly 1 percentage point in the first quarter and then returns to baseline within about five quarters. As explained earlier, the situation, in which both variables are negative corresponds to uncertainty-driven pessimism (Ha and So, 2024). The most striking result is that inflationary shocks are less persistent.

So far, the focus has been on average responses, the Growth-at-Risk framework allows us to look at how shocks affect the entire distribution of outcomes. In particular, the interest lies in the change in magnitude of the tails. To do so, I extend the VAR from the previous analysis to a quantile specification, relying on the full predictive distribution of

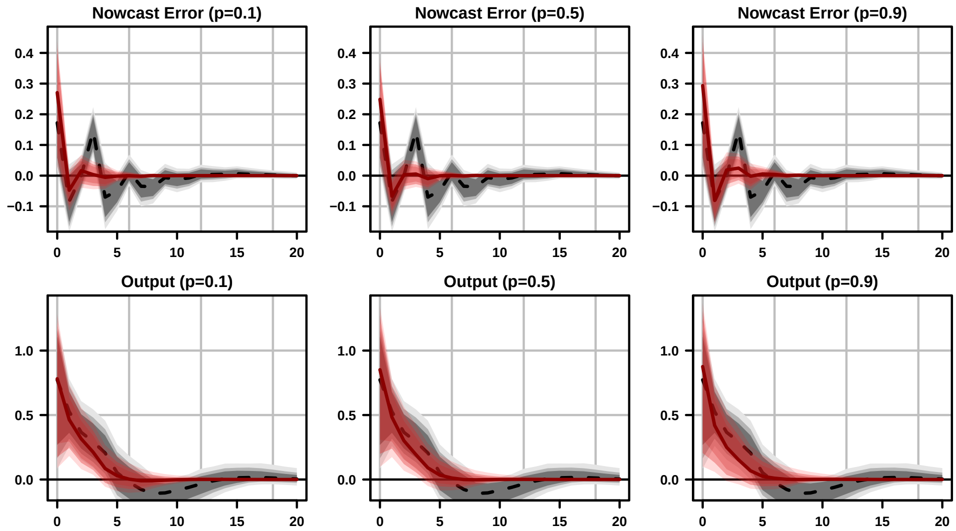

(a) Non-Belief Shocks

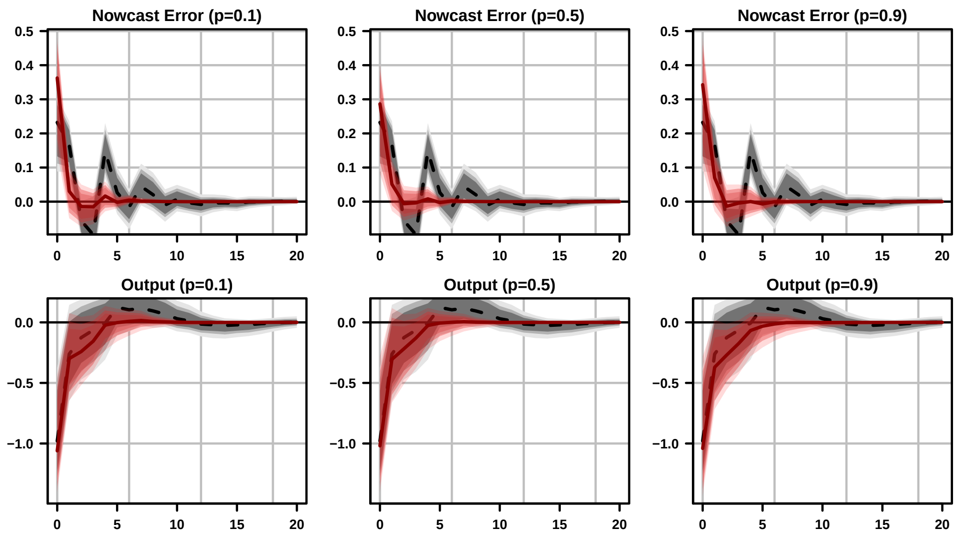

(b) Belief Shocks

Figure 3: Impulse response functions from the belief and non-belief shocks on GDP within Growth-at-Risk framework.

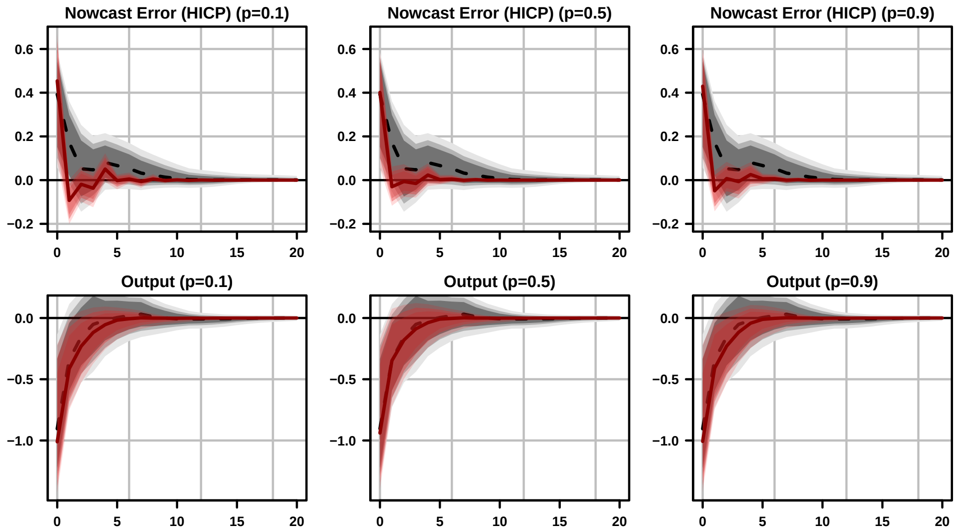

(a) Non-Belief Shocks

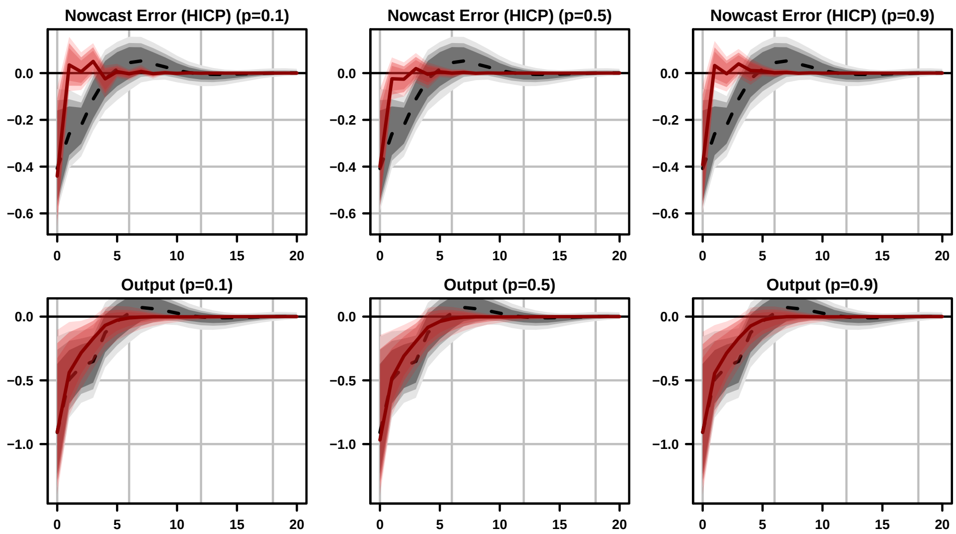

(b) Belief Shocks

Figure 4: Impulse response functions from the belief and non-belief shocks on HICP within Growth-at-Risk framework.

the SPF and the nowcast errors defined and shown in Figure 2. The shocks are identified through the sign restriction at three points of the distribution, $p \in \{0.1, 0.5, 0.9\}$. Figures 3 and 4 demonstrate the resulting impulse responses, where the red lines depict the GaR estimates across quantiles, while the grey curves reproduce the mean-based specification explained before, with the dashed line representing the posterior median.

For GDP, the responses are highly similar across quantiles. Non-belief shocks raise output, while belief shocks reduce it, with a persistence of about two years. The practically overlapping quantile responses point to the conclusion that the transmission of both shocks is nearly identical across the distribution. Turning to HICP, the distributional effects seem to be similarly uniform. Non-belief shocks display the expected stagflationary pattern, with nowcast errors rising and output falling across quantiles. Belief shocks likewise shift both series downward, but the effects remain almost perfectly parallel across the tails. There are minor differences in the sharpness of comeback to the baseline in nowcast errors, but those can be assumed to be insignificant nuances, while the central dynamic remains unchanged. From this perspective, the GaR framework does not overturn the mean-based evidence, but rather reinforces it while at the same time allowing to formally declare that the impulse responses are largely quantile-invariant. This strengthens the replication of the original evidence of Boeck and Pfarrhofer (2025) in euro area setting.

Although impulse responses point to invariance, the decomposition of the forecast error variance in Table 1 tells us to which shock,and at which quantiles the risk can be assigned. For GDP, belief shocks account for around half of output fluctuations at the mean, but rise to even 55-57% in the lower tails, indicating that downside risks are more belief-driven. By comparison, the inflation decomposition is flatter across quantiles, suggesting that both belief and non-belief shocks affect the output more uniformly. These results echo with the fundamentals of Growth-at-Risk presented by Adrian et al. (2019), that the macro-financial conditions disproportionally shape the lower tail of the distribution. In contrast, the euro area results presented here suggest that beliefs explain roughly half of output fluctuations at the one-year horizon and across quantiles. This points to a balanced 50/50 role between both shocks. This stands in contrast to Enders et al. (2021), who showed for U.S. that belief shocks account for one third at the shorter horizon and rise to around 40% at two years. The more even attribution in the euro area may be due to the distinct crises covered in the sample, namely the 2008 financial crisis, COVID-19 pandemic, and the 2021-23 energy shock. This may have amplified the role of expectations in both downside and median outcomes. However, those findings are in line with other papers on the expectation shocks in the business cycle, such as Milani (2011), who shows that waves of optimism and pessimism can account for roughly half of US business cycle fluctuations.

Table 1: Forecast Error Variance Decomposition (1-Year Horizon)

	GDP				**HICP**			
	Nowcast Error		Output		Nowcast Error		Output	
	NB	B	NB	B	NB	B	NB	B
SVAR Baseline	44.7	55.3	49.8	50.2	45.7	54.3	43.9	56.1
GaR (p = 0.05)	36.7	63.3	43.5	56.5	53.6	46.4	51.8	48.2
GaR (p = 0.10)	40.5	59.5	44.8	55.2	53.6	46.4	50.2	49.8
GaR (p = 0.16)	41.3	58.7	46.9	53.1	53.0	47.0	49.6	50.4
GaR (p = 0.50)	45.0	55.0	49.8	50.2	49.6	50.4	46.6	53.4
GaR (p = 0.84)	44.4	55.6	47.6	52.4	51.7	48.3	49.2	50.8
GaR (p = 0.90)	45.7	54.3	45.7	54.3	53.0	47.0	51.3	48.7
GaR (p = 0.95)	45.6	54.4	46.4	53.6	51.7	48.3	53.2	46.8
GaR (Realized)	44.8	55.2	49.6	50.4	48.8	51.2	45.8	54.2

Notes: NB = Non-belief shock, B = Belief shock. Values represent percent of forecast error variance explained at a one-year horizon.

To conclude this study, I strengthen the identification strategy for the mean-based specification by moving from single to double sign restrictions. Instead of requiring each shock to affect only one variable in a given direction, I impose the requirement that both nowcast errors and output must respond simultaneously in a theoretically consistent way. In practice, it means that a non-belief shock must decrease the NE for HICP and increase the NE GDP and output on impact. This tighter identification intends to reduce the risk that movements in GDP nowcast errors are offset by opposing reactions from NE HICP, which would leave the impact ambiguous.

The similarity between the original single-sign model and the double-sign strategy

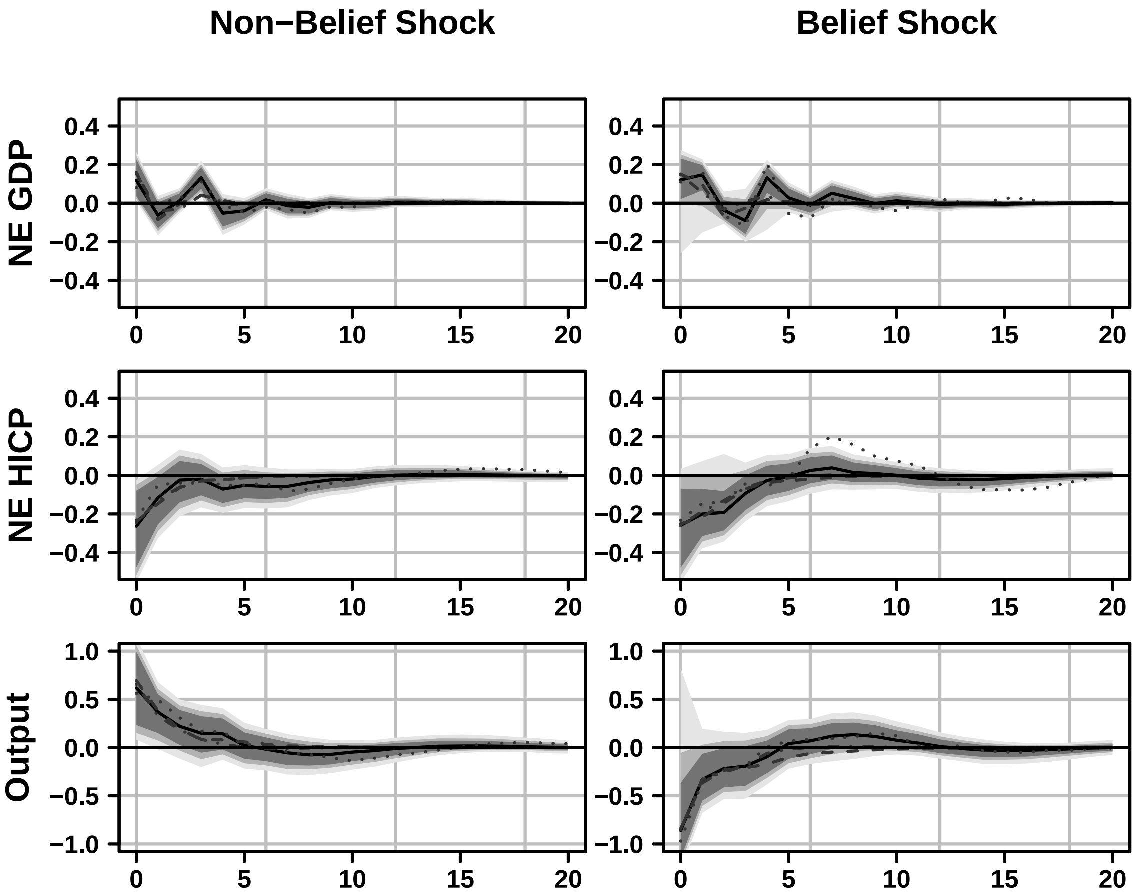

Figure 5: Impulse response functions from the double sign restricted SVAR

indicates that the imposed restrictions are almost redundant, as can be seen on Figure 5. When both criteria are applied, the shocks in fact behave in the same way as under single restriction. This suggests that the additional constraint does not materially change the identification. This finding is aligned with the work of Beaudry and Portier (2006), the authors show that alternative identification strategies deliver highly correlated shocks and impulse responses that are nearly identical. The reason for it is that the underlying variables already exhibit strong co-movement, so once one sign restriction is imposed, the other tends to be satisfied automatically.

6 Robustness check and limitations

To assess whether the baseline attribution (50/50 between belief and non-belief) is driven by extraordinary situations after 2019, I re-estimated the mean-based SVAR on three subsamples that either exclude the pandemic or focus on the crisis years. Identification and priors are unchanged, only the time windows vary. This follows Enders et al. (2021) who report stable dynamics and Boeck and Pfarrhofer (2025) who exclude the years after 2014 in their analysis (it is worth mentioning that their analysis focuses on the analysis until 2019, pandemic years are down weighted) and find the results to be robust.

Table 2: Forecast Error Variance Decomposition (FEVD): Subsample Comparisons

		1999–2019		2008–2024		2010–2024	
		Non-Belief	Belief	Non-Belief	Belief	Non-Belief	Belief
GDP	Nowcast Error	73.6	26.4	42.9	57.1	35.1	64.9
	Output	86.7	13.3	43.3	56.7	29.3	70.7
HICP	Nowcast Error	49.1	50.9	45.6	54.4	45.2	54.8
	Output	41.3	58.7	47.0	53.0	46.1	53.9

Notes: Table reports forecast error variance decomposition (FEVD) results for GDP and HICP across different subsamples.

The most remarkable results of this thesis are revealed when I exclude post-2019, as visible in Table 2. Non-belief shocks dominate the FEVD for the one-year horizon with around 87%, with beliefs shocks explaining only 13%. However, in the subsample heavy in crisis (2010-2024), the pattern flips with beliefs accounting for even 70% and non-belief for 30%. A broader crisis window (2008-2024) sits in between as characterized with beliefs explaining up to 57%. This indicates regime dependence in the importance of the two shocks, where fundamentals dominate in calmer periods. The IRF shapes remain similar, with differences only in the magnitude of NE overshooting and undershooting relative to the baseline. It is worth noting that this finding does not affect the HICP, as the 50/50 attribution is firm across different subsamples.

This time dependence of results matches the evidence on the expectation formation. Milani (2011) demonstrates using the New Keynesian model, which allows for learning by agents, that expectation shocks can drive half of the U.S. business cycle, especially in turbulent episodes when agents revise beliefs more frequently. Binder (2020) shows how COVID-19 fears caused large downward revisions of macroeconomic expectations even before official data reflected the downturn (survey conducted on 5 and 6 March), and Dietrich et al. (2022) show that expectation revisions were unusually frequent and sharp during and after the pandemic, this is especially visible in the SPF Nowcast distribution as shown in Figure 1. Altogether, this provides a consistent explanation, the euro area evidence is not an anomaly but part of a pattern, where belief shocks gain prominence in regimes with higher uncertainty, while fundamentals dominate in calmer times.

Altogether, the robustness checks highlight a sharp contrast with the U.S. evidence. While Enders et al. (2021) and Boeck and Pfarrhofer (2025) find stable dynamics in the United States, the euro area results are strongly driven by the regimes. Belief shocks dominate only in times rich in crises. Whereas non-belief shocks dominate in calmer times. This division in the results likely reflects the shorter ECB SPF sample, which is dominated by crises relative to the longer US dataset and therefore biases the euro area evidence. The COVID-19 crisis and its aftermath were strong enough to spread across the entire sample, increasing the importance of belief shocks. This points to one of the central limitations, namely, the euro area SPF covers only a short period, beginning in Q1 1999. This makes the results highly sample dependent, the post-2019 period biased the dynamics toward "crisis dynamics", as shown before.A longer dataset, such as the

U.S. SPF beginning in 1968, dilutes individual crises with calm periods and allows for a larger sample of crisis episodes, producing more stable variance decompositions and results (Enders et al., 2021). By comparison, the euro area dataset provides fewer observations and places unreasonable weight on extreme events.

A further limitation relates to the forecast horizon, as explained in the data chapter. Unlike the US SPF, the ECB survey only reports current and next year forecasts. This thesis therefore uses the current-year forecast as a proxy for the nowcast, since it exhibits high variability. However, this possibly introduces measurement error and may blur the identification, particularly when expectations are anchored and revised infrequently.

Finally, the limitation stemming from the theory-driven identification itself. In this thesis, sign restrictions are used to interpret shocks as belief and non-belief, following Enders et al. (2021) and Boeck and Pfarrhofer (2025). While the results demonstrate that the data are consistent with the strategies used, they should not be taken as a unique or objective truth. As Baumeister and Hamilton (2014) note, *"the procedure actually delivers a set of possible inferences, each of which is equally consistent with both the observed data and the underlying restrictions."*. In short, different structural models might match the same empirical responses; thus, the results should be looked at as one plausible interpretation rather than definitive proof.

7 Conclusion

In this thesis, I have examined how belief shocks, which I recognized as systematic errors in professional forecasts, shape macroeconomic outcomes in the Eurozone. Building on the work of Enders et al. (2021) and Boeck and Pfarrhofer (2025), the analysis offers a novel view that extends the framework beyond US data, by incorporating both output and inflation expectations from the ECB Survey of Professional Forecasters. Using a time-varying quantile regression developed by Pfarrhofer (2022) and a structural VAR with sign restrictions, the study aimed to link belief shocks to Growth-at-Risk, capturing not only average effects, but also the influence of shocks on the distributional tail of output growth.

The results confirm that GDP belief shocks matter for output, but their impact is not time consistent. Baseline decompositions show that they can account approximately for half of short run output fluctuations, which is in line with US evidence. However, robustness checks reveal a different story, as there is a time-varying pattern. In calm periods, fundamentals dominate (non-belief shocks), while in more turbulent regimes, mainly Covid-19 crisis and its aftermath, belief shocks surge in importance, explaining up to 70% of the business cycle fluctuations. Contrary, the analysis of the inflation shows a remarkable stability across subsamples, both shocks contribute in almost uniformly throughout all the subsamples. This suggests that while GDP belief shocks intensify in crises, inflation belief shocks represent a steadier dynamics.

These finding hint that belief shocks amplify vulnerabilities when expectations are unanchored. Covid-19 not only generated extraordinary wide forecast distributions, but also shifted the euro area sample toward "crisis dynamics", increasing the weight of belief shocks relative to fundamentals. Rather than offering a universal predictor of business cycle, in my case the Growth-at-Risk framework combined with belief shocks have not enriched the findings significantly, the same as in the original work by Boeck and Pfarrhofer (2025). Altogether, in accordance with these results, policymakers should prioritize expectations management and communication in moments of heightened volatility, as to keep expectations anchored.

This thesis contributes in three main ways. First, it extends the belief shock literature from the US to the Eurozone. Second, it incorporates inflation-based belief shocks in addition to output, showing the stable 50/50 role across subsamples. And lastly, demonstrates that Eurozone results are strongly crisis driven, contrary to the more stable US evidence. Together, these findings emphasize that expectations matter most when they are unanchored.

At the same time, several limitations constrain interpretation. The lack of a true nowcast, the short sample of the ECB SPF, and reliance on a sign restriction strategy mean the results must be considered as one possible reading rather than definitive evidence. Different models could yield alternative decompositions. Yet, the evidences align with the literature showing that optimism and pessimism drive the business cycle during the uncertain times. Future research could extend this work by exploring the importance of inflation belief shocks on a more reliable dataset: US SPF, exploring the heterogeneity across Eurozone countries to the shocks, or studying interactions with the communication strategy of the policymakers. In doing so, it would increase the understanding of how belief shocks reshape growth distribution and risk assessment in Europe.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019, April). Vulnerable Growth. *American Economic Review 109*(4), 1263–1289.

Adrian, T., F. Grinberg, N. Liang, S. Malik, and J. Yu (2022, July). The term structure of growth-at-risk. *American Economic Journal: Macroeconomics 14*(3), 283–323.

Andrle, M., J. Bruha, and S. Solmaz (2013, August). Inflation and Output Comovement in the Euro Area: Love at Second Sight? Working Papers 2013/07, Czech National Bank, Research and Statistics Department.

Angeletos, G.-M. and J. La'O (2013). Sentiments. *Econometrica 81*(2), 739–779.

Barro, R. J. (2001, June). Economic Growth in East Asia Before and After the Financial Crisis. NBER Working Papers 8330, National Bureau of Economic Research, Inc.

Baumeister, C. and J. D. Hamilton (2014, December). Sign restrictions, structural vector autoregressions, and useful prior information. Working Paper 20741, National Bureau of Economic Research.

Beaudry, P. and F. Portier (2006, 02). Stock prices, news, and economic fluctuations. *American Economic Review 96*, 1293–1307.

Bernanke, B. S., M. Gertler, and M. Watson (1997). Systematic Monetary Policy and the Effects of Oil Price Shocks. *Brookings Papers on Economic Activity 28*(1), 91–157.

Binder, C. (2020, 10). Coronavirus fears and macroeconomic expectations. *The Review of Economics and Statistics 102*(4), 721–730.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013, December). News, noise, and fluctuations: An empirical exploration. *American Economic Review 103*(7), 3045–70.

Blinder, A. S. and J. B. Rudd (2013, October). The Supply-Shock Explanation of the Great Stagflation Revisited. In *The Great Inflation: The Rebirth of Modern Central Banking*, NBER Chapters, pp. 119–175. National Bureau of Economic Research, Inc.

Boeck, M. and M. Pfarrhofer (2025). Belief shocks and implications of expectations about growth-at-risk. *Journal of Applied Econometrics 40*(3), 341–348.

Chari, V V; Jones, L. E. M. R. E. H. G. D. S. A. C. (1996, May). Inflation, growth, and financial intermediation. *Review - Federal Reserve Bank of St.Louis 78*(3), 41–64.

Clark, T. E., F. Huber, G. Koop, M. Marcellino, and M. Pfarrhofer (2024, October). Investigating Growth-at-Risk Using a Multicountry Nonparametric Quantile Factor Model. *Journal of Business & Economic Statistics 42*(4), 1302–1317.

Dedola, L. and S. Neri (2007). What does a technology shock do? a var analysis with model-based sign restrictions. *Journal of Monetary Economics 54*(2), 512–549.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020, July). What's up with the Phillips Curve? Working Paper Series 2435, European Central Bank.

Dietrich, A. M., K. Kuester, G. J. Müller, and R. Schoenle (2022). News and uncertainty about covid-19: Survey evidence and short-run economic impact. *Journal of Monetary Economics 129*, S35–S51.

Enders, Z., M. Kleemann, and G. J. Müller (2021). Growth expectations, undue optimism, and short-run fluctuations. *The Review of Economics and Statistics 103*(5), 905–921.

Firat, M. and O. Hao (2023, October). Demand vs. Supply Decomposition of Inflation: Cross-Country Evidence with Applications. IMF Working Papers 2023/205, International Monetary Fund.

Fischer, S. (1993, December). The role of macroeconomic factors in growth. *Journal of Monetary Economics 32*(3), 485–512.

Fève, P. and A. Guay (2016, 05). Sentiments in svars. *The Economic Journal n/a*(n/a).

Ghosh, M. A. R. and M. S. T. Phillips (1998, May). Inflation, Disinflation, and Growth. IMF Working Papers 1998/068, International Monetary Fund.

Gillman, M. and M. Kejak (2005, February). Contrasting Models of the Effect of Inflation on Growth. *Journal of Economic Surveys 19*(1), 113–136.

Ha, J. and I. So (2024, April). The Economic Effects of Global Inflation Uncertainty. *International Journal of Central Banking 20*(2), 69–150.

Hellwig, C. and V. Venkateswaran (2024, November). Dispersed information, nominal rigidities and monetary business cycles: A hayekian perspective. Working Paper 33215, National Bureau of Economic Research.

James, G., D. Witten, T. Hastie, and R. Tibshirani (2023). *An Introduction to Statistical Learning: With Applications in R* (3 ed.). New York: Springer.

Korobilis, D., B. Landau, A. Musso, and A. Phella (2021). The time-varying evolution of inflation risks. ECB Working Paper 2600, Frankfurt a. M.

Kozeniauskas, N., A. Orlik, and L. Veldkamp (2016, July). The common origin of uncertainty shocks. Working Paper 22384, National Bureau of Economic Research.

Krueger, F. (2017, 08). Survey-based forecast distributions for euro area growth and inflation: ensembles versus histograms. *Empirical Economics 53*.

Krüger, F. and I. Nolte (2016). Disagreement versus uncertainty: Evidence from distribution forecasts. *Journal of Banking & Finance 72*(S), 172–186.

Lang, J. H., M. Rusnák, and M. Greiwe (2023). Medium-term growth-at-risk in the euro area. ECB Working Paper 2808, Frankfurt a. M.

Lorenzoni, G. (2009, December). A theory of demand shocks. *American Economic Review 99*(5), 2050–84.

López-Villavicencio, A. and V. Mignon (2011). On the impact of inflation on output growth: Does the level of inflation matter? *Journal of Macroeconomics 33*(3), 455–464.

Manzan, S. and D. Zerom (2013). Are macroeconomic variables useful for forecasting the distribution of u.s. inflation? *International Journal of Forecasting 29*(3), 469–478.

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *The Economic Journal 121*(552), 379–401.

Oba, S., M. A. Sato, I. Takemasa, M. Monden, K. Matsubara, and S. Ishii (2003). A bayesian missing value estimation method for gene expression profile data. *Bioinformatics 19*(16), 2088–2096.

Peek, J., E. Rosengren, and G. M. B. Tootell (2016, September). Does Fed policy reveal a ternary mandate? Working Papers 16-11, Federal Reserve Bank of Boston.

Pfarrhofer, M. (2022). Modeling tail risks of inflation using unobserved component quantile regressions. *Journal of Economic Dynamics and Control 143*, 104493.

Prasad, A., S. Elekdag, P. Jeasakul, R. Lafarguette, A. Alter, A. X. Feng, and C. Wang (2019, February). Growth at risk: Concept and application in imf country surveillance. IMF Working Paper WP/19/36, International Monetary Fund.

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *The Review of Economic Studies 72*(3), 821–852.

Rapach, D. and G. Zhou (2018, 01). Sparse macro factors. *SSRN Electronic Journal*.

Stockman, A. C. (1981). Anticipated inflation and the capital stock in a cash in-advance economy. *Journal of Monetary Economics 8*(3), 387–393.

Tobin, J. (1965). Money and economic growth. *Econometrica 33*(4), 671–684.

Watson, J. H. S. M. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics 20*(2), 147–162.

Appendices

A Macroeconomic Dataset

Table 3: Overview of Macroeconomic Variables Used for Factor Extraction

#	Mnemonic	Description	Source	Frequency	Start Date	Trans.
1	NEER	Nominal Effective Exchange Rate (EA)	BIS	Monthly	1999Q1	5
2	CREDITGAP	Credit-to-GDP Gap (EA)	BIS	Quarterly	2009Q1	1
3	HOUSEPRICE	Real Residential Property Prices	BIS	Quarterly	1999Q1	5
4	HICPEA	HICP, All items	ECB	Monthly	1995Q1	1
5	HICPCORE	HICP, Core (ex food & energy)	ECB	Monthly	1995Q1	1
6	M3EA	Broad Money (M3)	ECB	Monthly	1999Q1	1
7	EUROSTOXX50	Euro Stoxx 50 Index	ECB	Monthly	1999Q1	5
8	CREDITHH	Credit to Households (Stock)	ECB	Quarterly	1999Q1	5
9	CREDITNFC	Credit to Non-Financial Corporations	ECB	Quarterly	1999Q1	5
10	RGDPEA	Real GDP, Euro Area	ECB	Quarterly	1999Q1	1

Continued on next page

Table 3 – continued from previous page

#	Mnemonic	Description	Source	Frequency	Start Date	Trans.
11	NGDPEA	Nominal GDP, Euro Area	ECB	Quarterly	1999Q1	5
12	ECB_MRO	ECB Main Refinancing Rate	ECB	Irregular	1999Q1	1
13	YIELD10Y	10-year Govt. Bond Yield	ECB	Monthly	1999Q1	1
14	TERM_SPREAD	10Y minus 3M yield	ECB	Daily	2004Q3	1
15	CORP_HYSPREAD	Euro High-Yield Corporate Bond Spread	FRED	Daily	1999Q1	1
16	EURUSD	EUR/USD Exchange Rate	FRED	Daily	1999Q1	5
17	IPIEA	Industrial Production Index	Eurostat	Monthly	1995Q1	5
18	UNRATE	Unemployment Rate (EA19)	Eurostat	Monthly	2000Q1	1
19	ESI	Economic Sentiment Indicator	Eurostat	Monthly	1999Q1	1
20	CCI	Consumer Confidence Index	Eurostat	Monthly	1999Q1	1
21	BCI	Business Climate Indicator	Eurostat	Monthly	1999Q1	1
22	CONS	Private Consumption Expenditure	Eurostat	Quarterly	1998Q1	5
23	EMPLOY	Total Employment	Eurostat	Quarterly	2000Q1	5
24	GFCF	Gross Fixed Capital Formation	Eurostat	Quarterly	1999Q1	5

Continued on next page

Table 3 – continued from previous page

#	Mnemonic	Description	Source	Frequency	Start Date	Trans.
25	EXPORTS	Exports of Goods & Services	Eurostat	Quarterly	1999Q1	5
26	IMPORTS	Imports of Goods & Services	Eurostat	Quarterly	1999Q1	5
27	GOVCONS	Government Consumption	Eurostat	Quarterly	1998Q1	5
28	SAVE_RATE	Gross Household Saving Rate	Eurostat	Quarterly	1999Q1	1
29	VSTOXX	VSTOXX Volatility Index	STOXX	Daily	1999Q1	2
30	SOV_SPREAD	Sovereign Spread (EA vs. DE, 10Y)	WSJ	Daily	2006Q1	1

Notes: **Trans. - Transformation** indicates the transformation applied: 1 = level, 2 = first difference, 4 = natural logarithm, 5 = first difference of log. Transformations are chosen based on data stationarity and comparability.

B Figures

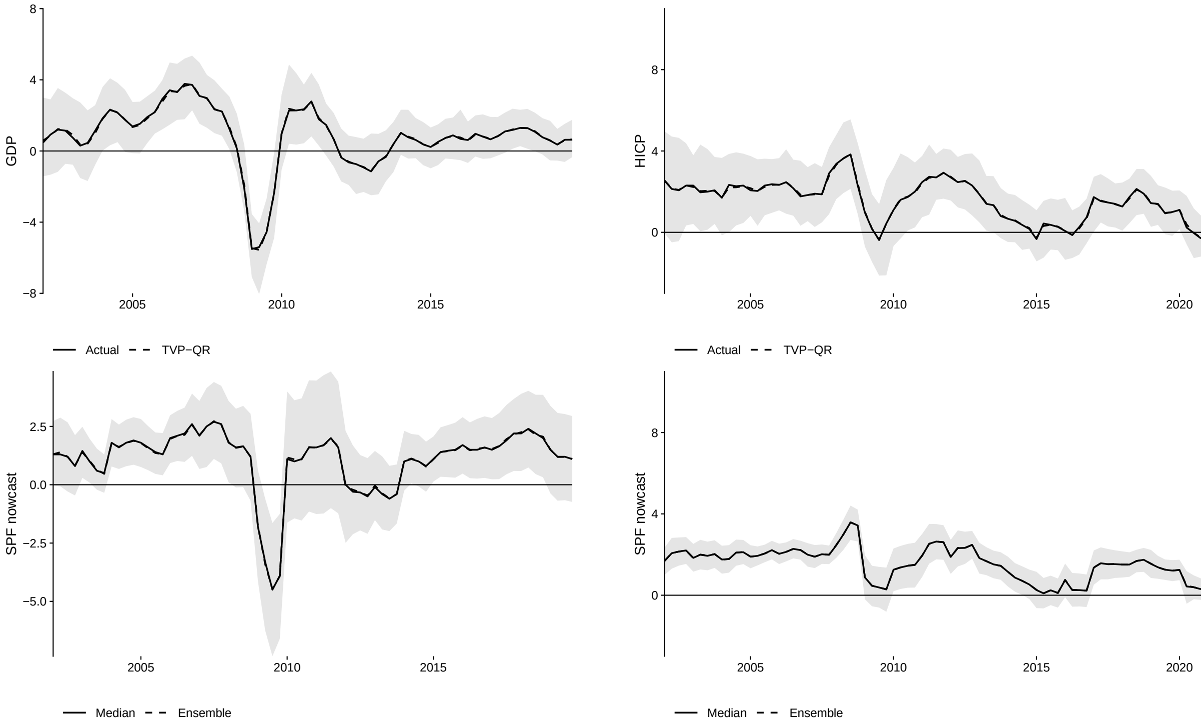

(a) GDP and GDP Nowcast distribution (b) HICP and HICP Nowcast distribution

Figure 6: Distribution of (a) GDP growth and its nowcast, and (b) HICP and its nowcast, within a time window 2001-2020.

Notes: Constructed using ensemble forecasts from the ECB SPF and TVP-QR. Shaded areas denote 90% predictive intervals.