EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Giacomo Carboni, Luís Fonseca,
Fabio Fornari, Leonardo Urrutia

Structural drivers of growth at risk: insights from a VAR-quantile regression approach

No 3171

Abstract

We investigate the impact of structural shocks on the joint distribution of future real GDP growth and inflation in the euro area. We model the conditional mean of these variables, along with selected financial indicators, using a VAR and perform quantile regressions on the VAR residuals to estimate their time-varying variance as a function of macroeconomic and financial variables. Through impulse response analysis, we find that demand and financial shocks reduce expected GDP growth and increase its conditional variance, leading to negatively skewed future growth distributions. By enabling this mean-volatility interaction, demand and financial shocks drive significant time variation in downside risk to euro area GDP growth, while supply shocks result in broadly symmetric movements. For inflation, supply shocks drive instead a positive mean-volatility co-movement, where higher inflation is associated with increased uncertainty, causing time variation in upside risk.

Keywords: Quantile Regressions, Tail Risk, Stochastic Volatility, Vector Autoregression (VAR), Mean-Variance Correlation, Structural Shocks, Downside Risk, Euro Area
JEL Codes: C32, C58, E32, G17.

Non-technical Summary

In recent decades, monitoring risks to the economic outlook has become a priority for policymakers around the world. In advanced economies, this growing emphasis has been driven by a series of extraordinary events, namely the global financial crisis and the European sovereign debt crisis, which brought significant macroeconomic disruptions and deflationary pressures, and more recently the COVID-19 pandemic and a subsequent surge in inflation. In response to the elevated uncertainty resulting from these events, researchers and policymakers have dedicated significant efforts to better characterise the risks surrounding central forecasts when evaluating economic prospects and formulating policy.

A growing body of research, known as the "growth-at-risk" literature, has focused on modelling these risks by estimating how current economic and financial variables influence the full distribution of future macroeconomic outcomes, particularly GDP growth. Since the seminal work of Adrian et al. (2019), two key insights have emerged from this literature. First, deteriorating financial conditions tend to be associated with both a lower conditional mean and a higher conditional volatility of GDP growth. Second, this mean-volatility interaction induces an asymmetry in the GDP growth distribution, as changes in financial conditions predict significant movements in the lower tail (downside risk), while the upper tail (upside risk) remains relatively stable.

While much of the existing literature has investigated how specific indicators help to forecast tail risks, this paper extends the analysis by identifying the underlying structural forces that drive these risks. Specifically, we investigate two questions: First, do specific structural shocks drive the mean-volatility interaction and, as a result, cause the asymmetric response of downside versus upside risks? Second, how do these shocks affect the joint risks to inflation and growth? To answer these questions, we model the conditional mean of macroeconomic variables using a standard vector autoregression (VAR) and perform quantile regressions on the VAR residuals, using lagged endogenous variables as predictors. This approach allows us to estimate the time-varying co-movement as a function of the state of the economy. The main implication is that current economic and financial conditions, and hence their underlying structural shocks, can influence the future distributions of GDP growth and inflation by affecting their mean, variance, or both.

Our findings reveal a significant time variation in the downside risk to GDP growth, whereby the lower quantiles are more volatile than the upper ones, consistent with the findings of Adrian et al. (2019). This asymmetry arises from the dominance of demand and financial shocks, which are estimated to reduce expected GDP growth while simultaneously increasing its future variance. Through this mean-volatility interaction, these shocks can lead the distribution of future

GDP growth to become negatively skewed. This effect is state-dependent, being particularly pronounced at times of heightened market distress. In contrast, supply shocks do not generate the same interaction. Instead, they affect the mean and variance less systematically, leading to rather symmetric movements in the lower and upper quantiles of GDP growth. For inflation, supply shocks drive instead a positive mean-volatility co-movement, where higher inflation is associated with increased uncertainty, causing a significant time variation in upside risk. Finally, our framework makes a substantial contribution by formalizing the analysis of joint tail risks for GDP growth and inflation. This approach enables the identification of (model-consistent) structural shocks driving these risks and provides policymakers with critical insights to the aim of managing them effectively.

1 Introduction

In recent decades, monitoring risks to the economic outlook has become a key element in the conduct of monetary policy for major central banks worldwide. This growing emphasis has been driven by a series of extraordinary events, such as the global financial crisis and the European sovereign debt crisis, with their profound macroeconomic disruptions and disinflationary pressures, and more recently, the COVID-19 pandemic and the subsequent surge in inflation. The heightened uncertainty brought forward by these events has prompted researchers and policymakers to devote substantial effort to characterise the risks surrounding central forecasts when assessing the economic outlook and formulating policy responses.

A growing body of research, commonly referred to as the "growth-at-risk" literature, has focused on evaluating these risks by estimating how financial and economic variables predict the entire distribution of future macroeconomic outcomes, particularly GDP growth. Since the seminal work of Adrian et al. (2019), two key insights have emerged. First, worsening financial conditions portend a lower conditional mean and higher conditional volatility of GDP growth. Second, this mean-volatility interaction induces an asymmetry in the GDP growth distribution, as changes in financial conditions predict significant movements in the lower quantiles while the upper quantiles remain relatively unaffected.

Our research is guided by two questions. First, do specific structural shocks drive the mean-volatility interaction and, consequently, cause the asymmetric response of downside versus upside risks? Second, how do these shocks shape the joint risks to the inflation and growth outlooks? Although these questions have received mixed attention in the literature, they remain highly relevant, as the appropriate policy response depends critically on the nature of the shocks driving these risks.

In this paper, we address these questions by examining how structural shocks affect the joint distribution of future real GDP growth and inflation. We model the conditional mean of these variables using a standard vector autoregression (VAR) and perform quantile regressions on the VAR residuals, using lagged endogenous variables as predictors. This approach allows for the estimation of a time-varying covariance matrix as a function of the state of the economy while conveniently retaining the assumption of conditional Gaussianity for the error terms. Additionally, this multivariate approach enables us to exploit cross-variable movements for structural identification. The implication is that current macroeconomic and financial conditions, along with their underlying structural shocks, can influence the future distribution of GDP growth and inflation by affecting their mean, variance, or both.

Our model reveals a significant time variation in the downside risk to GDP growth, with the lower quantiles moving more strongly than the upper ones, consistent with Adrian et al.

(2019). This asymmetry is not assumed, but rather stems from the prominence of demand and financial shocks, which are estimated to reduce expected GDP growth while simultaneously increasing its future variance. Through impulse response analysis, we document how the mean-volatility interaction induced by demand and financial shocks causes the distribution of future GDP growth to be negatively skewed. In contrast, supply shocks do not generate this type of interaction, as their effects on the mean and variance of GDP growth are less systematic, resulting in more symmetric tail movements. For inflation, supply shocks drive a positive mean-volatility co-movement, where higher inflation is associated with increased uncertainty, causing the upper tail of its distribution to become more volatile than the lower tail. Although more modestly, demand shocks qualitatively resemble supply shocks in terms of inflation risk, as they also generate a significant time variation in upside risk; in contrast, financial shocks drive a negative mean-volatility co-movement, leading to greater downside risk to inflation.

Our paper relates to two distinct strands of the literature. The first involves quantile regressions, which provide direct estimates of the quantiles of the distribution of a macroeconomic variable at specific horizons (e.g., Giglio et al., 2016; Adrian et al., 2019).[1] However, this approach has limitations in delivering path forecasts for quantiles and poses challenges for structural identification. Quantile vector autoregressions (Chavleishvili and Manganelli, 2024) partially address these limitations by modelling the simultaneous interaction between quantiles within and between different variables. However, the ordering of variables in this framework is crucial, the estimation process rapidly becomes computationally expensive, and achieving structural identification (beyond assuming a Cholesky decomposition of the covariance matrix) remains challenging.[2]

Given these limitations of the quantile regression approach, a second strand of the literature relies on VAR models with stochastic volatility, inspired by the initial contributions of Primiceri (2005) and Cogley and Sargent (2005). In their simplest form, these models feature a variance process that evolves due to exogenous uncertainty shocks.[3] More recent work incorporates observable variables into the variance law of motion, so to capture a direct mean-volatility interaction (e.g., Carriero et al., 2018; Caldara et al., 2021; Carriero et al., 2024).[4] However, a drawback of models of this type is that estimation results can be sensitive to the ordering of variables in the VAR, leading to different estimates for the conditional variances and for density forecasts (e.g., Arias et al., 2023; Chan et al., 2024). Although there are methods that allow

[1] This work spurred substantial follow-up research applying the methodology to different regions (see, e.g., Figueres and Jarociński (2020) and Adrian et al. (2022)). A recent strand has begun to explore machine learning approaches (see, e.g., Clark et al., 2023; Goulet Coulombe et al., 2023; Chronopoulos et al., 2024).

[2] Other applications of quantile vector autoregression can be found in Iseringhausen et al. (2023) and Huang et al. (2024). See also Korobilis and Schröder (2025) for an extension based on factors.

[3] See also Delle Monache et al. (2024), Wolf (2021), and Iseringhausen (2024) for univariate models with stochastic volatility.

[4] See also Montes-Galdón and Ortega (2022) for VARs with time-varying higher moments.

to overcome the ordering issue (Carriero et al., 2016; Chan et al., 2022), they have primarily been proposed for models where observables do not affect the variance processes.[5] Without a mean-variance interaction, these models cannot generate asymmetric patterns in downside and upside risks in out-of-sample forecasts. To our knowledge, no stochastic volatility model is both order-invariant and accounts for endogenous variables in the volatility process.

Our work relates to several other studies investigating the effects of identified shocks on predictive distributions.[6] One approach, used by Loria et al. (2025), is to employ pre-identified shocks within a local projections quantile regression framework. However, this method inherits a key limitation of standard quantile regressions: it cannot produce path forecasts for quantiles over time and across variables, making it unsuitable for the scenario and counterfactual analyses that our framework is designed to handle.[7]

Similarly to our approach, Boire et al. (2021) and Forni et al. (2023) integrate quantile regressions into a structural VAR framework to assess the impact of specific shocks on the risk surrounding key macroeconomic variables.[8] However, important differences remain. Boire et al. (2021), for instance, limit their analysis to a single variable and one structural shock with time-varying distribution, whereas our approach is more general, accommodating time-varying volatility in all shocks and variables. Our assumption of conditionally Gaussian errors also provides technical simplifications over non-parametric methods (as employed in Boire et al. (2021)), enabling the use of techniques which are standard in the VAR literature.

Another crucial distinction lies in the application of the quantile regression itself. Forni et al. (2023) consider a framework in which the dependent variable of quantile regressions is the level of macroeconomic variables, with predictors governed by a standard linear VAR. Consequently, although their framework may capture asymmetric patterns in-sample, the underlying data-generating process is linear and does not produce asymmetries in out-of-sample forecasts or in simulations. Our approach, instead, uses the VAR residuals as the dependent variables in the quantile regressions. This choice is crucial because it allows the VAR error terms to have time-varying volatility as a function of macroeconomic and financial variables, while their conditional mean is governed by the VAR model. As a result, by enabling a mean-volatility interaction, our framework can generate asymmetric patterns in downside versus upside risk that depend on prevailing economic conditions and the underlying structural shocks.

The paper is organised as follows. Section 2 introduces our modelling approach, after drawing

[5] See also Caldara et al. (2024) and Guerrón-Quintana et al. (2023) for factor-based approaches. In the latter, state dependence enters the mean equation rather than the shock variances.

[6] See Gertler (2020) on the relevance of structural analysis in assessing the importance of financial conditions for risks to economic activity.

[7] Based on the identification scheme of Lewis (2021), Carriero et al. (2021) propose a structural stochastic volatility VAR with mean shocks and endogenous components in the volatility process. However, their analysis focuses on the uncertainty shock and their implemented approach relies on a Cholesky decomposition to identify the other structural shocks economically.

[8] In a similar vein, see also Duprey and Ueberfeldt (2020) and Forni et al. (2025).

motivation from key stylized growth-at-risk facts, which we replicate for the euro area. Section 3 describes our empirical application. In Section 4, we estimate the model and identify the structural shocks driving the various quantiles of the distribution of future real GDP growth and inflation. Section 5 illustrates some policy applications. Section 6 concludes.

2 A VAR-Quantile Regression Approach

2.1 Informing Our Modelling Approach: Key Growth-at-Risk Facts

Our contribution is motivated by two key insights from the growth-at-risk literature. First, worsening financial conditions portend a lower conditional mean and higher conditional volatility of GDP growth. Second, this mean-volatility interaction generates an asymmetry in the distribution of GDP growth, as changes in financial conditions predict significant movements in the lower quantiles of GDP growth while the upper quantiles remain relatively unaffected. Although these insights were initially derived from U.S. data, they also hold for a broader range of advanced economies (Adrian et al., 2019; Figueres and Jarociński, 2020; Adrian et al., 2022).

In Figure (1), we replicate this evidence for the euro area by following the methodology of Adrian et al. (2019) and conducting linear and quantile regressions (QR) to characterise the conditional relationship between future GDP growth and current financial and economic conditions. In these regressions, our dependent variable is one-quarter-ahead real GDP growth, while the explanatory variables include current GDP growth, inflation, changes in the short-term interest rate[9], and the Composite Indicator of Systemic Stress (CISS) of Kremer et al. (2012). We use the CISS as a broad proxy for financial conditions and financial stress and, for brevity, will often refer to it simply as "financial conditions".

Figure (1a) shows a scatterplot of the conditional mean of the GDP growth distribution (y-axis) against its conditional dispersion (x-axis), which is proxied by the difference between the upper and lower quantiles. Each dot represents a single quarter in our sample. Overall, the figure illustrates a strong negative correlation between the conditional mean and volatility of GDP growth, confirming a key stylised fact established by the literature. The regression results, not shown here for brevity, indicate that as financial and economic conditions deteriorate, expected growth declines, causing the entire growth distribution to shift downward. The accompanying increase in conditional volatility implies that this downward effect is amplified for the lower quantiles and dampened for the upper quantiles. Consistently, Figure (1b) shows the 0.1 and 0.9 quantiles of the conditional distribution of GDP growth, confirming that the lower quantile exhibits greater variability than the upper quantile.

Building on these insights, the model we present formalises the interaction between mean

[9]Specifically, we use the shadow interest rate from Krippner (2013), which addresses the challenge of using nominal rates during the effective lower bound period.

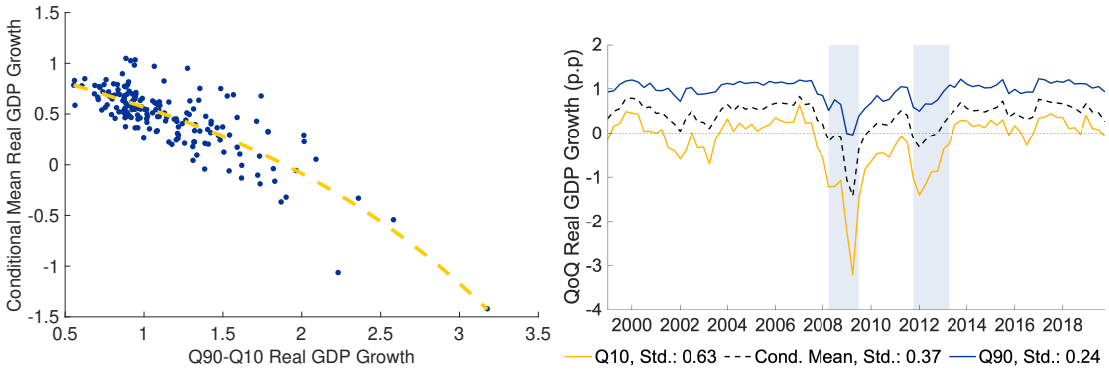

(a) Conditional mean and volatility of GDP growth

(b) Expected GDP growth and Quantiles

Figure 1: Stylized Growth-at-Risk Facts for euro area real GDP growth (Quantile Regressions)

Note: The results are derived from quantile regressions (for the 0.1 and 0.9 quantiles) and an OLS regression (for the mean) for quarter-on-quarter real GDP growth. The covariates include lagged GDP growth, inflation, changes in the Shadow Short-Term Rate (SSR), and the CISS. The estimation sample spans from 1973:Q1 to 2019:Q3. Shaded areas denote recession dates.

and volatility by modelling them as functions of a common set of observable variables. It then identifies the structural shocks that drive time-varying tail risks by exploiting standard identification schemes. To do so, it incorporates quantile regression techniques within a VAR framework, effectively addressing some limitations of standalone quantile regressions. These limitations include difficulties in characterising path forecasts for quantiles over time and across variables, as well as challenges in integrating structural identification.

2.2 Model Set-up

The empirical insights documented in the previous section inform our approach to modelling both the conditional mean and variance as functions of financial and economic variables. First, we model the conditional mean of $\mathbf{y}_t$ using a linear VAR specification. Second, we model the conditional quantiles of the VAR residuals as a function of lagged endogenous variables by running quantile regressions on the reduced-form VAR residuals.

For the conditional mean, consider the following reduced-form linear VAR(p) specification:

$$\mathbf{y}_t = c + A_1\mathbf{y}_{t-1} + \ldots + A_p\mathbf{y}_{t-p} + \mathbf{u}_t \tag{2.1}$$

where $\mathbf{y}_t$ is an $n\times 1$ vector of endogenous variables, c is an $n\times 1$ vector of constants, $A_1, \ldots, A_p$ are $n \times n$ coefficient matrices, and $\mathbf{u}_t$ is an $n \times 1$ vector of serially uncorrelated, reduced-form innovations (white noise) with $\mathbb{C}\text{ov}(\mathbf{u}_t) = E[\mathbf{u}_t\mathbf{u}_t'] = \Sigma_{u,t}$. We can derive consistent estimates of the A_k parameters using ordinary least squares (OLS), whether $\Sigma_{u,t}$ is constant (homoscedastic) or time-varying (heteroscedastic).[10]

[10]In the latter case, the OLS estimator is less efficient but remains unbiased and consistent. At the same time, the OLS approach may encounter scalability challenges, as model estimation becomes increasingly complex with

For the conditional variance, we proceed as follows. Let $\mathbf{X}_{t-1} = [1, \mathbf{y}'_{t-1}, \ldots, \mathbf{y}'_{t-q}]'$ be a $(1 + nq) \times 1$ vector of regressors, where $q \leq p$. The conditional τ-quantile of the i-th residual, $u_{i,t}$, is given by:

$$Q^{\tau}_{u_{i,t}}(\mathbf{X}_{t-1}) = \mathbf{X}'_{t-1}\beta_{i,\tau} \tag{2.2}$$

where $\beta_{i,\tau}$ is the corresponding $(1 + nq) \times 1$ vector of coefficients estimated for each residual series ($i = 1, \ldots, n$) using the quantile regression estimator from Koenker and Bassett Jr. (1978). This specification allows the conditional quantiles of the residuals to be driven by the recent history of all variables in the system. The goal is to use these quantile regressions to estimate $\Sigma_{u,t}$ as a function of lagged endogenous variables. To do this, we fit a parametric Gaussian distribution to the estimated quantiles for each reduced-form residual, a method we describe in more detail below.[11]

Intuitively, our approach jointly models the conditional mean and volatility of the variables as functions of a common set of underlying drivers, $\mathbf{y}_{t-1}, \ldots, \mathbf{y}_{t-q}$, rather than treating these two moments as independent processes. By explicitly modelling this mean-volatility interaction, the model can generate non-Gaussian, asymmetric conditional distributions for $\mathbf{y}_{t+j}$ for horizons $j > 1$, while conveniently retaining the assumption of conditional Gaussianity for the one-period-ahead residuals.[12] Consequently, any estimated asymmetry in the distribution of the endogenous variables arises from this interaction rather than from assumptions about the distribution of the error terms. Shocks that induce a correlation between the conditional mean and volatility of these variables will naturally lead to asymmetry in the distribution of $\mathbf{y}_{t+j}$.

Before describing how we derive the properties of $\Sigma_{u,t}$, it is important to highlight three key aspects of our methodology. First, while the VAR residuals are uncorrelated with lagged endogenous variables by construction, this lack of correlation does not extend to the quantiles of the residuals. A systematic relationship between these quantiles and lagged variables can reflect dependencies through higher-order moments. We focus specifically on the relationship between the variance of the residuals and lagged endogenous variables.

Second, performing quantile regressions on the reduced-form VAR residuals is a key inno-

a growing number of variables, a limitation not easily overcome even by Bayesian methods when variances are allowed to be time-varying. Nonetheless, our analysis adopts a deliberate approach to variable selection, aiming to strike a balance between a parsimonious set of relevant structural drivers and the inclusion of key predictors of tail risk.

[11] Our two-stage estimation procedure has similarities with Koenker and Zhao (1996), who also estimate a univariate ARCH model using a combination of OLS and QR. However, our direct estimation of quantiles, as opposed to scaling parameters, allows for a better assessment of the fit of subsequent parametric distributional approximations.

[12] In our model, the assumption of a Gaussian distribution for one-period-ahead errors is convenient, enabling direct specification of the parameters of the multivariate normal distribution, $\Sigma_{u,t}$, which facilitates model simulation. Crucially, this assumption does not limit the model's ability to generate non-Gaussian conditional and unconditional distributions for the endogenous variables. Indeed, as noted by Carriero et al. (2024) and Caldara et al. (2021), Gaussian heteroscedastic models are capable of generating skewness in forecast distributions.

vation. For example, Forni et al. (2023) run quantile regressions of y_{t+j} on its lagged variables while assuming that y_t follows a linear VAR process. Consequently, although their model can estimate in-sample asymmetries, it cannot generate them out-of-sample because the underlying data-generating process is linear. By contrast, our framework generates such asymmetries by allowing macroeconomic and financial conditions to influence the mean-volatility interaction. Moreover, it provides predictive paths for the full distributions of the endogenous variables.

Third, our approach allows us to uncover the drivers of changes in tail risk, a task not feasible in frameworks where stochastic volatility arises from exogenous processes. In our model, uncertainty is explained by observed economic conditions, as represented by the variables in the VAR. By linking uncertainty to the endogenous variables and, by extension, to the underlying structural shocks, enables us to decompose movements in tail risks according to the contributions of those shocks, as we discuss in Section 3.

2.3 Estimation Procedure

As outlined previously, the initial step of the estimation involves running a VAR on the selected endogenous variables to derive OLS parameter estimates for $A_1, \ldots, A_p$ and calculating the resulting reduced-form residuals. The quantile regressions on these residuals form the foundation for constructing their conditional distributions and, notably, the $\Sigma_{u,t}$ matrices.

The Variance-Covariance Matrix of Reduced-Form Residuals, $\Sigma_{u,t}$

To derive the diagonal elements of $\Sigma_{u,t}$, we first estimate selected quantiles of the distribution of the reduced-form residuals. We then fit a univariate Gaussian distribution to these quantiles using a least squares method.[13] Since we assume that the reduced-form errors are Gaussian with zero mean, we would only require a single quantile to estimate the variance $\sigma^2_{u_{i,t}}$. For robustness, we base our estimate on two percentiles, the 10^{th} and 90^{th}, although more could be used. While we are interested in tail risk, we avoid more extreme quantiles for which estimates may be less precise. For similar reasons, we use fewer lags in the quantile regressions (q) than in the VAR (p) to avoid over-fitting.

To complete the estimation, we must specify the off-diagonal terms in the covariance matrix $\Sigma_{u,t}$. We derive these terms as a by-product of the structural identification. The relationship between the reduced-form residuals $\mathbf{u}_t$ and the structural shocks ϵ_t is given by:

[13] We use least squares to determine the variance parameter that minimises the squared distance between the estimated and theoretical quantiles:

$$\min_{\sigma_{u_{i,t}}} \sum_{\tau \in \{\tau_1, \ldots, \tau_K\}} [Q^{\tau}_{u_{i,t}} - \sigma_{u_{i,t}} \cdot \Phi^{-1}(\tau)]^2 \tag{2.3}$$

where Φ^{-1} is the inverse CDF of the standard normal distribution. The solution is the OLS estimator, which is a convenient linear function of the endogenous variables. This means we can decompose the stochastic variance into contributions from the lagged values of all endogenous variables.

$$\mathbf{u}_t = B_0 \Lambda_t^{1/2} \epsilon_t \tag{2.4}$$

where $\epsilon_t \sim \mathcal{N}(0, I_n)$ are the i.i.d. structural shocks, B_0 is the $n \times n$ constant impact matrix describing how shocks contemporaneously affect the endogenous variables, and $\Lambda_t^{1/2}$ is a diagonal matrix containing the time-varying standard deviations of the structural shocks. Any time variation in the covariance of the reduced-form errors arises solely from changes in the volatility of the structural shocks. This specification leads to the time-varying covariance matrix:

$$\Sigma_{u,t} = B_0 \Lambda_t B_0' \tag{2.5}$$

where Λ_t is the diagonal matrix containing the variances of the structural shocks.

Structural Identification

As a final step, we outline the identification strategy that delivers estimates of B_0 and Λ_t. Further details are provided in appendix A. Broadly, we apply standard sign and narrative restrictions to identify a set of valid impact matrices $\{B_0(m)\}_{m=1}^{M}$. The reduced-form and structural variances are linked via:

$$\text{diag}(\Lambda_t) = (B_0(m) \odot B_0(m))^{-1} \text{diag}(\Sigma_{u,t}) \tag{2.6}$$

where the reduced-form variances, $\text{diag}(\Sigma_{u,t})$, are based on the quantile regression results. We discard any candidate matrix that implies negative structural shock variances at any point in time, ensuring that $\text{diag}(\Lambda_t)$ remains positive. Finally, we exploit the model's heteroscedasticity to identify a unique impact matrix B_0 from the set of suitable candidates. As eq. (2.6) shows, alternative impact matrices lead to different in-sample structural variances. Unlike in homoscedastic VARs, this time variation in variances makes the likelihood function sensitive to the parameters in B_0. This feature allows us to select a unique impact matrix using a maximum likelihood approach, similar to methods used in the SVAR-GARCH literature (e.g., Luetkepohl and Milunovich, 2016).

Summary of Estimation Steps

The model can be summarized by the following system of equations:

$$\mathbf{y}_t = c + A_1 \mathbf{y}_{t-1} + \ldots + A_p \mathbf{y}_{t-p} + \mathbf{u}_t, \qquad \mathbf{u}_t \sim \mathcal{N}(0, \Sigma_{u,t}) \tag{2.7}$$

$$\mathbf{u}_t = B_0 \Lambda_t^{1/2} \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, I) \tag{2.8}$$

where $\Lambda_t = \text{diag}(\lambda_{1t}, \ldots, \lambda_{nt})$ is the diagonal matrix of time-varying structural shock variances. The dynamics of the reduced-form variances are determined by quantile regressions. For each

reduced-form residual $u_{j,t}$, we first estimate a set of K conditional quantiles, $\{\tau_k\}_{k=1}^K$:

$$Q_{u_{j,t}}^{\tau_k}(\mathbf{X}_{t-1}) = \mathbf{X}'_{t-1}\beta_{j,\tau_k} \tag{2.9}$$

From these estimated quantiles, we obtain the conditional standard deviation by fitting them to the theoretical quantiles of a standard normal distribution:

$$\hat{\sigma}_{u,j,t} = \frac{\sum_{k=1}^K Q_{u_{j,t}}^{\tau_k} \cdot \Phi^{-1}(\tau_k)}{\sum_{k=1}^K [\Phi^{-1}(\tau_k)]^2} \tag{2.10}$$

The estimated variances form the diagonal of the reduced-form covariance matrix, $\text{diag}(\Sigma_{u,t}) = [\hat{\sigma}^2_{u,1,t}, \ldots, \hat{\sigma}^2_{u,n,t}]'$. Finally, the structural variances are recovered via:

$$\text{diag}(\Lambda_t) = (B_0 \odot B_0)^{-1}\text{diag}(\Sigma_{u,t}) \tag{2.11}$$

Algorithm 1 Estimation Procedure for the VAR-QR Model

1: Estimate the VAR parameters $(c, A_1, \ldots, A_p)$ via OLS and compute the reduced-form residuals, $\mathbf{u}_t$.
2: For each residual series $u_{j,t}$ and for each required quantile τ_k, estimate the coefficient vector β_{j,τ_k} using quantile regression.
3: Using the estimated quantiles, compute the conditional variances $\hat{\sigma}^2_{u,j,t}$ for each residual series, forming $\text{diag}(\Sigma_{u,t})$.
4: Identify the impact matrix B_0 and the structural variances Λ_t using sign restrictions and the maximum likelihood method described in Section 2.3.

3 Empirical Application

3.1 Data

We estimate our model using euro area data, focusing on the distributions of real GDP growth and inflation. Our dataset covers the period from January 1980 to December 2019, ending before the COVID-19 pandemic.[14] The VAR includes four variables: the log of the real Gross Domestic Product (GDP) index, the log of the headline consumer price index (HICP), the shadow short-term interest rate (SSR) from Krippner (2013), and the Composite Indicator of Systemic Stress (CISS) from Kremer et al. (2012). We choose the SSR over the nominal short-term interest rate as the latter is a poor indicator of the monetary policy stance during the effective lower bound period. The CISS captures financial conditions by incorporating volatility and spread components from financial markets and has been shown to be useful for forecasting tail risk (e.g., Figueres and Jarociński, 2020).[15] Since GDP data are only available quarterly, we disaggregate

[14] We stop the sample before the pandemic for practical reasons. Lenza and Primiceri (2022) provide a potential avenue to integrate the large outliers in macroeconomic data generated by the pandemic in a standard VAR framework.

[15] Following Kremer (2016), we take the square root of the CISS to account for potential non-linearities.

the series to a monthly frequency with the method of Santos Silva and Cardoso (2001) using data on Industrial Production excluding the construction sector. Additional details about the data and their transformations can be found in appendix B. Based on the likelihood and information criteria for different lag specifications, we estimate the model using five lags in the VAR part ($p = 5$) and two lags for the endogenous variables in the quantile regressions based on the VAR residuals ($q = 2$). This more parsimonious specification for the latter regressions is intended to mitigate the risk of over-fitting the tails.

3.2 Identification Scheme

After estimating the VAR coefficients, the quantile regression parameters, and the conditional standard deviations, we complete the final step of our procedure by estimating the impact matrix B_0 using standard structural identification methods. Specifically, we employ a combination of sign (e.g., Uhlig, 2005), magnitude (e.g., Peersman, 2005), and narrative restrictions (e.g., Antolín-Díaz and Rubio-Ramírez, 2018).[16]

As detailed in Table 1, our sign restrictions require that a positive monetary policy shock increases both the shadow rate and the CISS, while causing a subsequent decline in GDP. In addition, the price level is persistently lowered, consistent with the literature. Positive demand shocks lead to increases in GDP, HICP, and the shadow rate. Supply shocks result in a negative co-movement between GDP and inflation. Finally, financial shocks are assumed to cause a decrease in output as well in the price level and the shadow rate, while being accompanied by an increase in the CISS.[17]

Variable/Shock	Monetary	Demand	Supply	Financial
GDP	$-_{(6)}$	$+_{(6)}$	$+_{(6)}$	$-_{(6)}$
HICP	$-_{(6\text{-}12)}$	$+_{(6)}$	$-_{(12)}$	$-_{(6)}$
SSR	$+_{(0\text{-}4)}$	$+_{(6)}$		$-_{(0\text{-}3)}$
CISS	$+_{(0\text{-}3)}$			$++_{(0\text{-}3)}$

Table 1: Sign and Magnitude Identification Restrictions

Note: The +/- signs indicate the direction of the restriction, while the double signs (++) specify that the shock is restricted to have the largest contemporaneous impact on the respective variable (magnitude restriction). Numbers in parentheses denote the time horizons, measured in months following the shock, for which the restrictions are imposed.

[16] We use code from Adrian et al. (2020) for the quantile regressions, the VAR Toolbox from Ambrogio Cesa-Bianchi for estimation, the 'ZeroSignVar' package from Breitenlechner et al. (2019) for structural identification, and routines from Kevin Sheppard's MFE Toolbox for cross-correlation estimation. The algorithm used also allows for exclusion restrictions.

[17] The literature discusses two types of financial shocks—financial uncertainty and financial conditions shocks—which are highly correlated (e.g., Caldara et al., 2016; De Santis and Van der Veken, 2022). We adopt a broad definition to encompass both types. Huang et al. (2024) conduct a structural analysis to identify how these alternative shocks influence the quantiles of US GDP growth.

Regarding magnitude and narrative restrictions, we require that financial shocks (i) have the largest absolute contemporaneous impact on the CISS among all shocks, and (ii) played a more significant role in the 2009 financial crisis than supply and monetary policy shocks combined.[18] While these restrictions typically result in set identification, we uniquely identify the structural shocks by exploiting the time-varying volatility of our model, as described in appendix A.[19] Finally, to enhance identification, we impose restrictions at specific horizons, as indicated in Table 1. Most of these restrictions apply for up to six months (two quarters). In our identification scheme, only the restrictions on monetary and supply shocks extend beyond six months.

4 Results

4.1 Asymmetric Patterns in Downside Versus Upside Risks

First, we document that our model replicates the stylized facts from the growth-at-risk literature shown in Section 2.1, notably the asymmetry between downside and upside risks to future macroeconomic variables.

Significant time variation in the downside risk to GDP growth The model generates significant time variation in downside risks to GDP growth, whereas upside risks are more stable. While this pattern is evident when examining the quantiles of future GDP growth, we introduce more comprehensive measures of tail risk, i.e. *Expected Shortfall* (ES) and *Expected Longrise* (EL), drawing on Adrian et al. (2019) and Caldara et al. (2021), among others. These measures represent the expected value of a random variable conditional on it being below or above a given percentile, respectively. Using the 5^{th} and 95^{th} percentiles as thresholds, ES and EL provide a more detailed representation of higher-order moments than simple quantiles, as they capture changes across the entire tail of the distribution.[20]

Figure (2) illustrates the behavior of the one-year-ahead ES and EL for GDP growth and inflation. The downside tail of GDP growth is significantly more volatile than its upside tail. Interestingly, the evidence for inflation is less clear; if anything, the EL for inflation appears

[18] Specifically, we require that the drag on year-on-year GDP growth in January 2009 from financial shocks was greater than the drag from the combined effects of supply and monetary policy shocks.

[19] We perform the structural analysis on the subsample starting in January 1999. The mathematical restrictions on the parameters discussed in appendix A make potential changes in the structural parameters more important for our model than for linear models. Evidence in the literature suggests that reduced-form parameters change little compared to the structural side (see, e.g., Primiceri, 2005; Carriero et al., 2018). Therefore, we leverage the full dataset to maximise the information available for the reduced-form estimation while allowing for flexibility in the structural identification by not constraining the full sample to a single impact matrix, with the beginning of the euro area serving as a reasonable structural break.

[20] ES and EL are defined as: $ES_{t+h} = \frac{1}{\pi}\int_0^{\pi} \hat{F}^{-1}_{y_{t+h}|x_t}(\tau)d\tau$ and $EL_{t+h} = \frac{1}{\pi}\int_{1-\pi}^{1} \hat{F}^{-1}_{y_{t+h}|x_t}(\tau)d\tau$, where $\pi = 0.05$, $\hat{F}^{-1}_{y_{t+h}|x_t}$ is the estimated inverse conditional CDF of the variable of interest y given information x_t, and h is the forecast horizon.

more volatile than the ES.

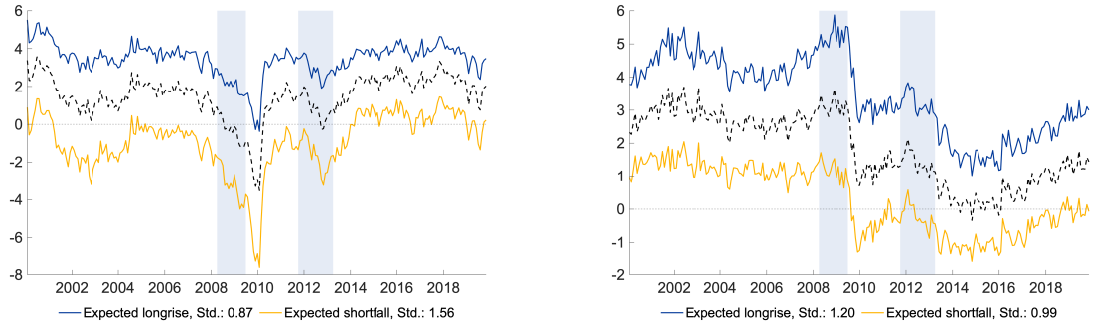

(a) GDP Growth One Year Ahead (YoY) (b) Inflation One Year Ahead (YoY)

Figure 2: Expected Shortfall and Longrise for Real GDP Growth and Inflation

Note: The figures plot one-year-ahead in-sample conditional mean (black) for real GDP growth and inflation derived from the VAR-QR model, along with the expected shortfall (yellow) and expected longrise (blue), using the 5^{th} and 95^{th} percentiles are respective thresholds. Shaded areas denote recession dates.

Mean-volatility interaction As anticipated in Subsection 2.1, the asymmetry in the behaviour of the upper and lower tails of GDP growth distribution can be attributed to an interaction between its conditional mean and volatility. In particular, a negative co-movement emerges: a lower conditional mean of GDP growth coincides with higher volatility. This pattern amplifies the effects of negative GDP growth outturns on the lower tail while dampening its effects on the upper tail. Figure (3) displays the relationship between the in-sample conditional mean and the difference between the conditional 10^{th} and 90^{th} quantiles for monthly, quarterly, and annual real GDP growth rates. It is important to note that we estimate a single monthly VAR-QR model and derive results for all frequencies from it. Therefore, the dynamics of the mean and uncertainty are driven by the economic patterns in the data, and do not result from smoothing across frequencies. Overall, our model successfully reproduces the negative, and possibly non-linear, relationship between the conditional mean and volatility of GDP growth across horizons.

To illustrate the underlying mechanism, we examine the response of the future real GDP growth distribution to a reduced-form CISS shock, calibrated to match the magnitude observed during the 2008-2009 financial crisis. Figure (4) compares the distribution of average GDP growth over two quarters in scenarios with the shock (yellow densities) versus without (blue densities) for both our VAR-QR model (left) and a standard linear homoscedastic VAR (right).

In both models, the CISS shock induces a leftward shift of the distribution, driven by a change in the mean (red arrow). The linear model incorporates only this mean effect, causing both tails to shift by the same amount. In our VAR-QR model, the distribution undergoes an additional change: the deterioration in financial conditions not only lowers the mean but

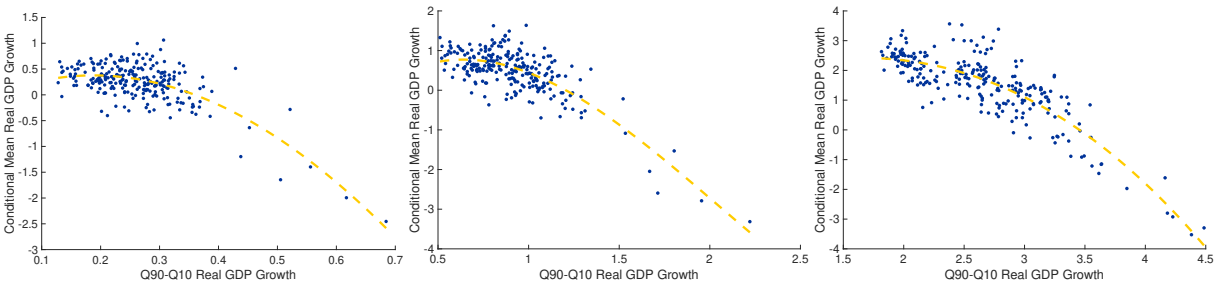

(a) Month on Month (b) Quarter on Quarter (c) Year on Year

Figure 3: Correlation between Conditional Mean and Volatility of GDP Growth (VAR-QR)

Note: The figure shows in-sample conditional predictions from the VAR-QR, depicting the relationship between the mean of real GDP growth and the Q90-Q10 quantile difference (a proxy for volatility). The panels show forecasts for (a) one-month-ahead monthly growth, (b) one-quarter-ahead quarterly growth, and (c) one-year-ahead annual growth. The yellow dashed line represents a cubic function fitted to the data points.

also simultaneously increases the variance (green arrows). This variance effect amplifies the impact of the mean shift on the lower tail while dampening it for the upper tail. Consequently, the ES declines substantially more than the EL. Through this mean-volatility interaction, our model generates skewed forecast distributions for the endogenous variables while retaining the assumption of Gaussian innovations.[21] Importantly, our model allows for these interactions, but does not impose them; they are a result of the estimation.

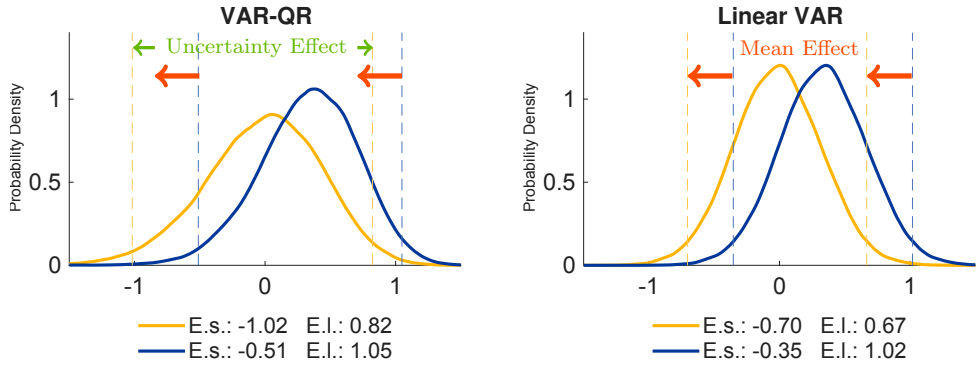

Figure 4: Conditional Forecast Distributions: VAR-QR vs. Linear VAR

Note: The figures depicts the densities of the distribution of average quarter-on-quarter real GDP growth over the two quarters ahead. The yellow densities are conditional on a 0.5 units shock to the CISS, while the blue densities are the baseline with no shock. The left panel shows the VAR-QR model, and the right panel shows the linear VAR.

4.2 Structural Drivers of Growth-at-Risk

Having established that our model replicates key stylized facts from the growth-at-risk literature, we now turn to our main objective: identify the structural drivers of macroeconomic tail risk. We first assess the effects of individual structural shocks on the conditional distribution of macroeconomic variables before turning to their historical contributions over time.

[21] Even the unconditional distribution (blue density in the left panel of Figure (4)) is slightly skewed to the left, reflecting the underlying mean-volatility interaction.

Impulse responses We begin by extending the analysis from Figure (4) to the identified structural shocks. To isolate the impact of each shock on the predicted distribution of GDP growth, we derive a "baseline" predicted density (blue) from Monte Carlo simulations of our model, conditional on the state of the economy, and allowing for all structural shocks to occur. We then generate a "shocked" density (yellow) by incorporating a one-time realization of a specific shock at the beginning of the simulation horizon. The difference between these densities captures the impact of that shock. Figure (5) displays the resulting predicted densities of real GDP growth for the four shocks we consider, i.e. monetary policy, demand, financial, and supply. For each shock, we compare the densities from our VAR-QR model (left panels) with those from a linear homoscedastic VAR (right panels). The main finding is that demand and financial shocks generate pronounced asymmetries in the predicted distributions from the VAR-QR model, with the lower tail responding much more strongly than the upper tail. This is evident from the substantial decline in the ES compared to the EL. This asymmetric pattern contrasts sharply with the linear model's predictions, where these shocks induce a parallel shift of the conditional distribution. Monetary policy shocks also induce a similar, though smaller, asymmetric response in our VAR-QR framework. Supply shocks, on the other hand, have symmetric effects, with no major differences emerging between the two models.

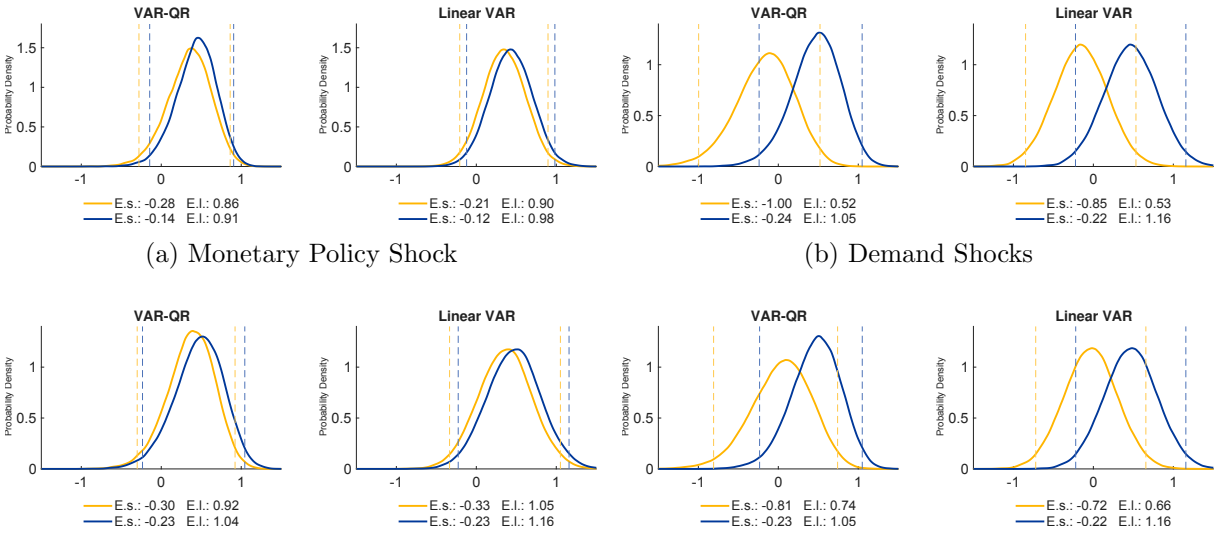

(a) Monetary Policy Shock (b) Demand Shocks

(c) Supply Shock (d) Financial Conditions Shock

Figure 5: Conditional Distributions Following Structural Shocks

Note: The figures show the densities of the distributions of average quarter-on-quarter real GDP growth over the two quarters ahead (and over the four quarters ahead for the monetary policy shock). The yellow densities are conditional on a one-time realization of a specific shock at the beginning of the simulation horizon while the blue densities correspond to the baseline scenario with no (one-time) shocks. In each sub-figure, the left panel represents the VAR-QR model, while the right panel corresponds to the linear VAR. The system is initialized during a calm period (June 2005).

Building on this exercise, we adopt a complementary perspective to illustrate the asymmetric response of the predictive densities to various shocks. Figure (6a) illustrates the mean-volatility

interaction associated with each shock type for GDP growth. Consistent with our findings, we find that demand and financial shocks, and to a lesser extent monetary policy shocks, generate a negative mean-volatility co-movement. In contrast, supply shocks do not appear to trigger a significant mean-volatility interaction.

Figure (6b) shows the results of the same exercise for inflation. In this case, supply shocks are the key driver of a positive mean-volatility co-movement, as higher inflation is associated with wider uncertainty, resulting in a more volatile upper tail compared to the lower tail. A key difference from the GDP growth results lies in the behaviour of demand shocks, which, in the context of inflation, act differently from financial and monetary policy shocks. In fact, demand shocks behave more similarly to supply shocks, as they also generate a positive mean-volatility interaction. A possible explanation is that the systematic monetary policy response effectively curtails the risk of very low inflation outcomes following an adverse demand shock. However, this policy response appears to be less effective against financial shocks, which tend to generate more downside than upside risk to inflation. It is also worth noting that, aside from supply shocks, the magnitude of these mean-volatility interactions is generally more modest for inflation compared to GDP.

(a) GDP growth (b) Inflation

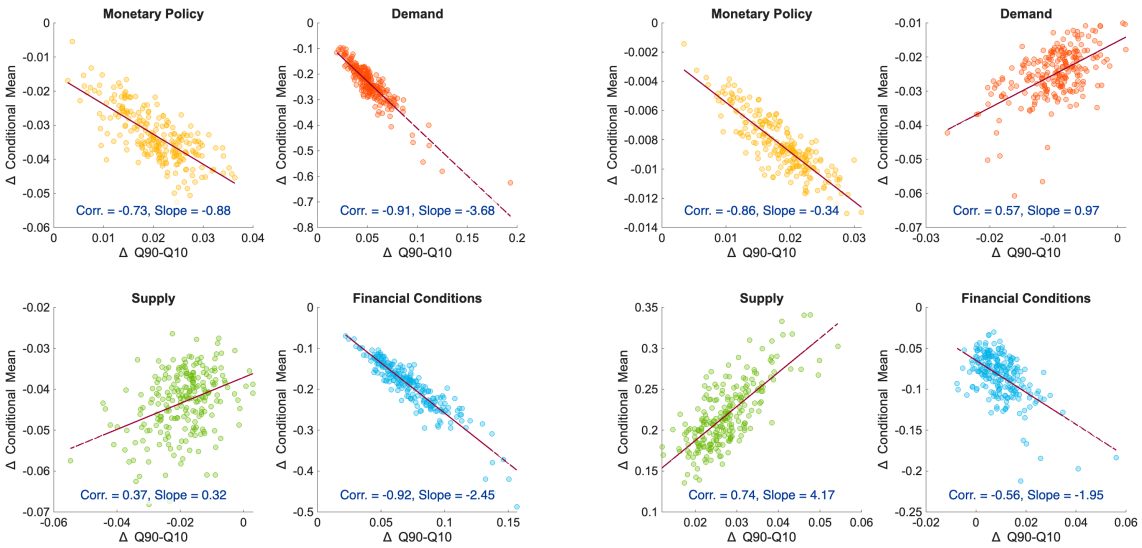

Figure 6: Historical Decomposition of the Mean-Volatility Correlation for GDP Growth and Inflation

Note: The figures generalize the analysis in Figure (5) by replicating the simulation across all points in the sample. At each point, an additional structural shock is simulated, and the resulting changes in the conditional mean are plotted against the corresponding changes in the inter-quantile range. The frequencies and forecast horizons are identical to those in Figure (5).

While these density plots document the effect of shocks at a specific horizon, they do not capture the dynamic evolution of tail risk. To address this, we use Generalized Impulse Response

Functions (GIRFs) for the ES and EL.[22]

To illustrate how these mechanisms vary under different economic conditions, we simulate the impact of shocks originating at two dates: June 2005, a period of calm (upper panels), and January 2009, the peak of the global financial crisis (lower panels).

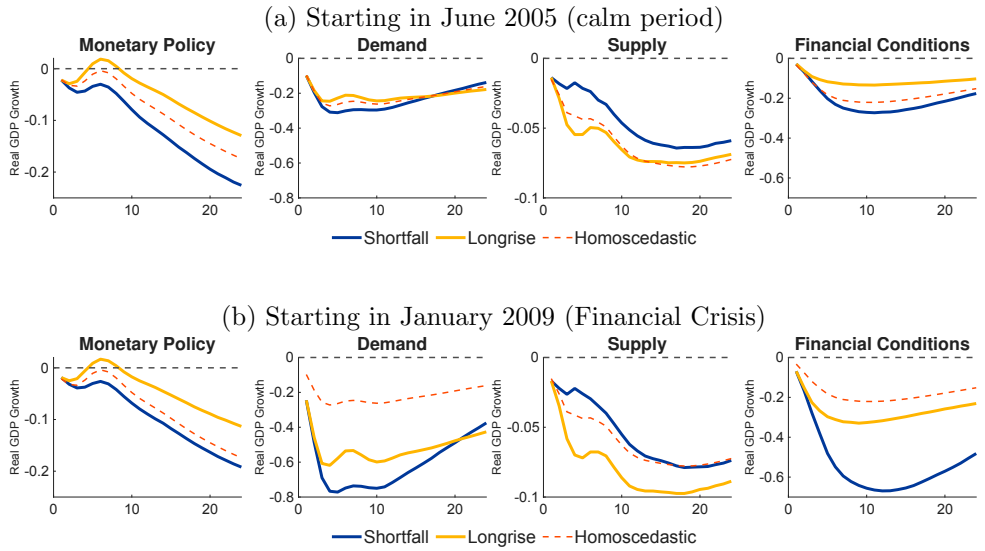

Figure 7: Cumulative GIRFs of ES and EL of GDP Growth to Structural Shocks

Note: The figures plot the cumulative GIRFs for the Expected Shortfall (ES) and Expected Longrise (EL) of GDP growth at a horizon of h months ahead. For instance, at $h = 12$, the lines represent the cumulative GDP growth between $t = 0$ and $t = 12$. The analysis incorporates one-standard-deviation shocks in June 2005 and January 2009.

In our model, the responses of ES and EL diverge significantly following demand-type shocks (monetary policy, demand, and financial). These differences become evident as early as three months after the shock. For instance, after a financial shock, the ES declines much more than the EL and remains depressed for an extended period. In contrast, after supply shocks, the EL initially experiences a more pronounced negative impact than the ES.

The results from our model contrast with the GIRFs produced by the homoscedastic model, the latter typically remaining between the ES and EL values that our model generates. This suggests that while the mean dynamics of the two specifications are broadly comparable, the dynamics in the tails of the distribution differ due to variations in the amount of uncertainty, which can occur either in tandem with or independently of the changes in the mean. As illustrated in Figure (C3), we find that the divergence between the ES and EL responses is statistically significant for monetary, demand, and financial shocks, but not for supply shocks.

Finally, the importance of this state-dependent mechanism is underscored when we evaluate the impulse responses during a period of financial distress (Figure (7), panel b). While the

[22]The non-linearities in our model require the use of GIRFs, which can be understood as an average IRF conditional on a given starting point. There is a distribution of possible IRFs around this average, which we refer to as "economic uncertainty" (implied by the model's nature), distinct from "statistical uncertainty" (rooted in estimation precision).

responses are qualitatively similar to those prevailing in the calm period, their magnitudes are quantitatively much larger. For instance, the responses to demand and financial shocks are more than twice as large as those observed during the calm period. This amplification implies that the differences between our model and its homoscedastic counterpart become even more pronounced during crisis episodes, as the latter produces IRFs that are not state-dependent.

5 Historical Structural Decomposition and Counterfactuals

5.1 Historical Structural Decomposition

We now conduct a historical decomposition to evaluate how different structural shocks have shaped the evolution of macroeconomic variables over time. Our analysis departs from the conventional focus on the mean of the distribution to rather emphasize the behaviour of the tails. Historical decompositions in non-linear models pose greater technical challenges than in their linear counterparts. The primary difficulty is that a shock at time t not only has a direct linear effect on the variables but also alters the variance of future shocks. This raises the question of whether observed changes should be attributed solely to current shocks or also to previous shocks creating the volatile conditions that amplify or dampen the effects of current shocks.[23]

Here, we adopt the *change in forecast function* (CFF) approach from Balke (2000) to examine how a structural shock at time t alters the forecast for future variables.[24] The CFF isolates the linear and non-linear effects of a single shock type, while interactions between different shock types are absorbed into a residual term.

Using this approach, Figure (8) shows the historical decomposition of the ES for one-quarter-ahead euro area GDP growth. Our structural framework offers a more insightful decomposition of growth-at-risk than reduced-form quantile regressions because it directly attributes downside risk to specific economic shocks. Our decomposition identifies financial and demand shocks as significant drivers of the ES for GDP growth in the euro area. In the lead-up to the 2008 financial crisis, adverse financial shocks, compounded by negative demand shocks, pushed the left tail of GDP growth to historical lows. Financial shocks also played a pivotal role during the onset of the sovereign debt crisis in 2012.

Table (2) systematically quantifies the contribution of each shock to the ES of GDP growth on average over our sample period. The results confirm that financial and demand shocks were

[23] Attempts at historical decomposition in this context include Carriero et al. (2018) and Chan et al. (2022). However, in the first, only the impact of uncertainty and mean shocks as a composite are disentangled, while the latter focuses only on the decomposition of the mean variables. Since the model lacks the GaR channel, it is also not useful for decomposing downside risk, which is independent of the state of the economy in this model.

[24] More formally, the CFF is computed as $\text{CFF}(\Omega_{t-1}, k, i) = \mathbf{E}[Y_{t+k}|\Omega_{t-1}, \epsilon_t^i, \ldots, \epsilon_{t+k}^i] - \mathbf{E}[Y_{t+k}|\Omega_{t-1}]$, where Ω_{t-1} is the information set at $t-1$, k is the forecast horizon, and ϵ^i is a structural shock of type i. The final term captures the contribution of initial conditions. Notably, in linear models, the CFF aligns with the exact historical decomposition.

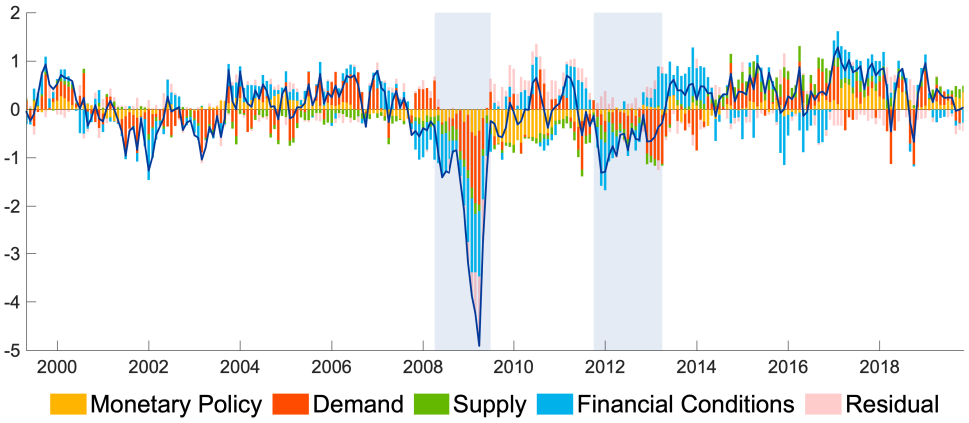

Figure 8: Structural Decomposition of the ES of GDP Growth (QoQ, one quarter ahead)

Note: The figure shows the structural decomposition of the Expected Shortfall (ES) for one-quarter-ahead quarter-on-quarter GDP growth, derived from the estimated VAR-QR. Both the contributions and the series are displayed as deviations from trend. The decomposition is carried out using the Change in Forecast Function (CFF) method. Shaded areas denote recession dates.

the primary drivers of downside tail risk in the euro area, jointly explaining approximately 60% of the variation in the ES of GDP growth, while supply shocks accounted for around 11%.

Contribution of structural shocks to the mean					
	Monetary	Demand	Supply	Financial	Residual
GDP Growth, QoQ	14	38	9	27	11
Inflation, QoQ	9	14	46	21	9
SSR	23	24	17	24	9
CISS	15	19	13	43	8
Contribution of structural shocks to Expected Shortfall of GDP growth					
One Quarter ahead, QoQ	15	30	11	29	13

Table 2: Shock Contributions (in Percent) to Historical Decompositions (CFF Method)

Note: Each shock's contribution is computed as $\text{Contribution}_j^i = \frac{\sum_{t=1}^{T} |\Gamma_{i,t}^j|}{\sum_{j=1}^{J} \sum_{t=1}^{T} |\Gamma_{i,t}^j|}$, where $\Gamma_{i,t}^j$ is the contribution of shock j to the deviation of variable i from its trend at time t. We use 1,000 paths for the forward simulation of the ES at each point in time. Contributions may not sum to 100% due to rounding.

Mean vs. Uncertainty Channel As previously discussed, tail dynamics are shaped by the interplay of mean and uncertainty effects. To disentangle their historical importance, we decompose the variance of the ES as follows:

$$\begin{aligned}\mathbb{V}\text{ar}(ES_{t+h}) &= \mathbb{V}\text{ar}(Mean_{t+h} + Uncertainty_{t+h}) \\ &= \overbrace{\mathbb{V}\text{ar}(Mean_{t+h})}^{\text{mean component}} + \underbrace{\mathbb{V}\text{ar}(Uncertainty_{t+h}) + 2\mathbb{C}\text{ov}(Mean_{t+h}, Uncertainty_{t+h})}_{\text{uncertainty component}}\end{aligned}$$

Figure (9) illustrates this decomposition for the ES of GDP growth at one-, three-, and twelve-month-ahead horizons. Within each panel, we break down the contributions of each shock into variations driven by the mean (solid bars) and those driven by uncertainty (patterned bars). The uncertainty component, indicated by the percentages above the bars, captures all variation not explained by the mean component and is thus derived as a residual. The findings indicate that while the mean channel dominates, the uncertainty channel plays a significant role, accounting for 14% to 52% of the variance in the ES, depending on the type of shock and length of the horizon. Moreover, the absolute and relative contributions of the uncertainty channel increase with the forecast horizon, as the mean-volatility interaction has more time to materialise. Financial and demand shocks are the dominant drivers of the overall effects, though the impact of different shocks tends to converge over longer horizons.

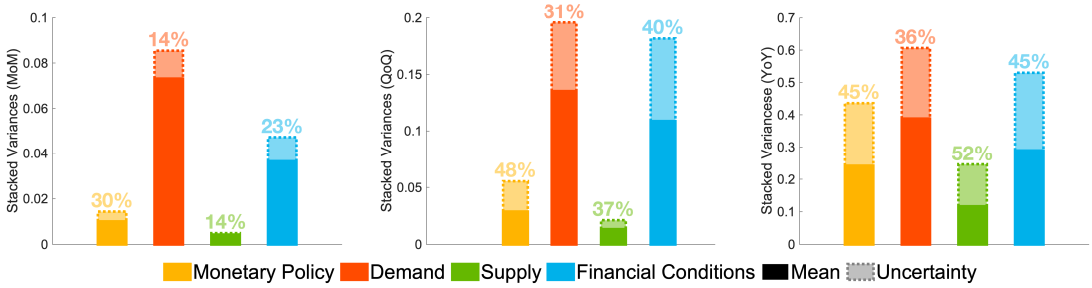

Figure 9: Mean and Uncertainty Channel Contributions to ES of GDP Growth

Note: The figure illustrates the variance decomposition of the Expected Shortfall (ES) by shock and forecast horizon, and specifically, for one month ahead (MoM), one quarter ahead (QoQ), and one year ahead (YoY). The mean component is defined as $Mean_{t+h} = (\mathbb{E}_t[\Delta GDP_{t+h}] - \Delta GDP^{Trend}_{t+h})$, while the uncertainty component is given by $Uncertainty_{t+h} = ES^{GDP}_{t+h} - Mean_{t+h}$, capturing the residual variation in the tail. Solid bars represent $\mathbb{V}\text{ar}(Mean_{t+h})$, whereas the more transparent bars with a dashed outline represent $\mathbb{V}\text{ar}(Uncertainty_{t+h}) + 2\mathbb{C}\text{ov}(Mean_{t+h}, Uncertainty_{t+h})$. The percentages indicate the share of the uncertainty component to the total variance of the ES. The decomposition is conducted using the CFF method.

The decomposition for inflation, shown in Figure (C4), reveals several key findings. First, inflation tail risk is primarily driven by the conditional mean, though the uncertainty channel remains significant. Second, the contribution of the uncertainty channel varies substantially across shocks and horizons (from 5% to 58%). Third, supply shocks are the main drivers of variation in inflation risk, with a strong contribution from the mean channel.

Structural Drivers of Joint Tail Risk We now analyse how structural shocks affect the joint distribution of inflation and GDP growth. While the notion of a "tail" is intuitive for

a single variable, its interpretation becomes more complex for a joint distribution. To shed light on this aspect, we adopt the following approach. First, we construct a baseline joint distribution for GDP growth and inflation by simulating the system under the combined impact of all shocks. Next, we incorporate a specific additional shock to the system and evaluate its effects. The resulting outcomes are categorized into four distinct regions, based on whether GDP growth and inflation values lie above or below the mean of the baseline distribution. Using this approach to highlight joint macroeconomic risks, Table 3 summarises the effects of each structural shock on the probability of falling into one of these four regions, starting from a baseline distribution evaluated as of January 2009.

The results indicate how contractionary demand and financial shocks move GDP and inflation in the same direction, thereby increasing the likelihood of simultaneous downward adjustments in both variables while conversely reducing the probability of combined upward movements. Monetary policy shocks, on the other hand, exhibit a relatively limited impact, reflecting their smaller magnitude during a period characterized by systematic and predictable monetary policy. In contrast, adverse supply shocks significantly increase the likelihood of a stagflationary trade-off, where recessionary forces are accompanied by inflationary pressures.

Figure (10) provides a graphical illustration on how financial shocks primarily drive positively correlated outcomes for GDP growth and inflation. The main panel of the figure contrasts the baseline joint probability distribution from January 2009 (blue contour lines) with the post-shock predicted distribution (solid yellow lines and filled density). The financial shock induces a significant shift of the probability mass toward a severe recession and a deflation, as indicated by the increased density in the bottom-left corner of the joint distribution. Conversely, the likelihood of an expansionary, inflationary outcome diminishes significantly. Meanwhile, the probability of a negatively correlated outcome, a stagflation (a combination of recession with inflationary pressures) declines slightly, reflecting the demand-like nature of the financial shock. The marginal distributions for each variable are displayed on the sides, with dashed lines marking the 5^{th} and 95^{th} percentiles.

5.2 Historical Counterfactual

The Global Financial Crisis without the Uncertainty Channel In our framework, a shock not only directly impacts the economy but also influences the variance of future shocks, amplifying their effects down the road. This mechanism is conceptually linked to a vast literature on the interaction between financial conditions and uncertainty, which highlights the amplifying role of risk-taking and binding constraints.[25] In our model, changes in uncertainty are proxied

[25]Recent contributions in this direction include Brunnermeier and Sannikov (2014), Guerrieri and Iacoviello (2017), and Aikman et al. (2024).

	Probability of Joint Scenarios			
Growth/Inflation	↓ & ↓	↑ & ↓	↓ & ↑	↑ & ↑
Before Shock: January 2009				
	30%	21%	18%	32 %
Effects of shocks on changes in probabilities (p.p.)				
Monetary	1	1	-1	-1
Demand	21	-14	12	-20
Supply	-13	-13	16	10
Financial	25	-2	-5	-18

Table 3: Structural Shock Effects on the Joint Distribution of Inflation and GDP Growth

Note: The table displays, in the first row, the probabilities associated with the baseline joint distribution for GDP growth and inflation, obtained by simulating the system under the combined impact of all shocks as of January 2009. Subsequent rows show the percentage point changes in these probabilities resulting from incorporating a specific additional two-standard-deviation shock. The outcomes for GDP growth and inflation are partitioned into four regions depending on whether they fall below(↓) or above (↑) the mean of the baseline distribution.

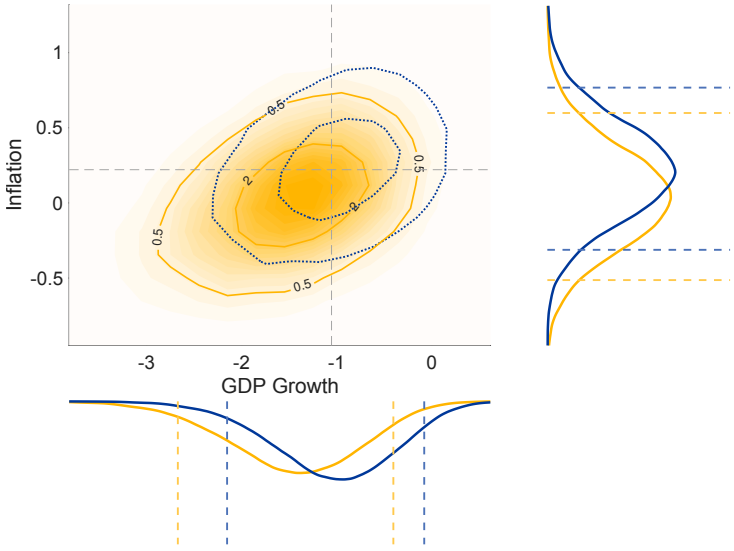

Figure 10: Financial Conditions Shock: Joint Distribution of Inflation and Real GDP Growth

Note: The figure displays the simulated joint and marginal distributions of average quarter-on-quarter real GDP growth and inflation two quarters ahead, starting from January 2009. The blue lines refer to the baseline distribution, obtained by simulating the system under the combined impact of all shocks. The yellow distribution is derived by incorporating an additional two-standard-deviation financial shock. Dashed coloured lines indicate the 5^{th} and 95^{th} percentiles of the marginal distributions shown along the axes. The dashed grey lines indicate the conditional mean of the baseline distribution.

by shifts in shock variances.

The structural nature of our framework allows us to construct a historical counterfactual to explore how the economy would have evolved without this uncertainty channel. Specifically, we simulate a counterfactual trajectory for the euro area economy during the 2008-2009 global financial crisis, keeping the variances of all shocks fixed at their unconditional mean.

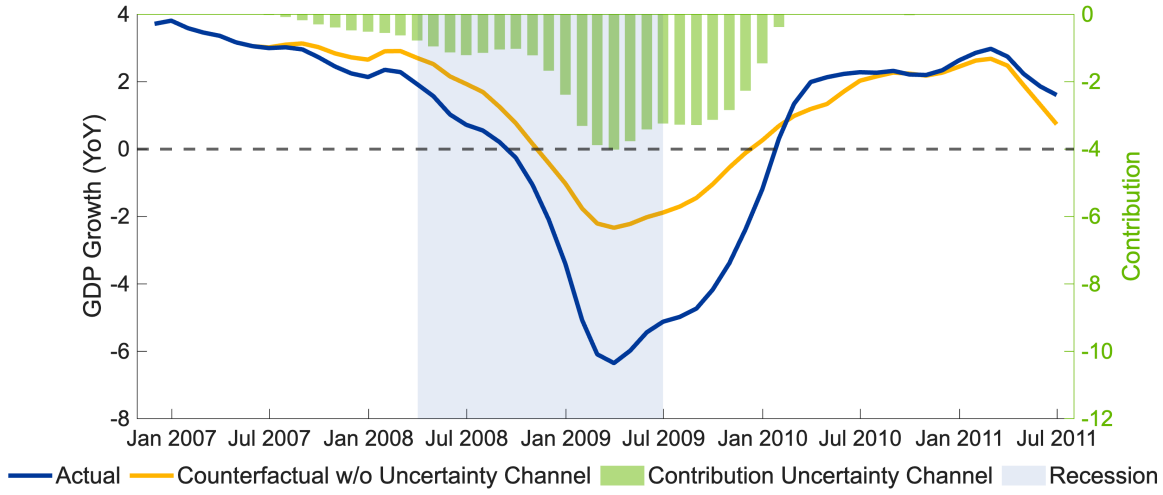

Figure 11: Counterfactual Global Financial Crisis

Note: The figure plots the actual trajectory of year-on-year real GDP growth (blue line) against a counterfactual path (yellow line) where the uncertainty channel is turned off from June 2007 to December 2010 by holding shock variances constant at their unconditional mean. The bars quantify the contribution of this channel, representing the difference between the two scenarios. Shaded areas denote recession dates.

Figure (11) depicts both the actual path of year-on-year GDP growth (blue line) and the counterfactual path without the uncertainty channel (yellow line). A substantial portion of the sharp decline in GDP growth during the financial crisis can be attributed to the amplification effect arising from this channel. Specifically, the uncertainty channel amplifies the decline in GDP growth by approximately 4 percentage points at the trough of the recession, accounting for roughly two-thirds of the overall contraction.

6 Conclusion

This paper presents a framework for identifying the structural drivers of risks to GDP and inflation. By shifting the focus from traditional macroeconomic indicators to structurally identified shocks, our analysis provides a deeper understanding of the dynamics of both downside and upside risks. The proposed approach integrates a VAR model, aimed at estimating the conditional mean of the distributions, with quantile regressions, capturing the time-varying volatility of the VAR residuals. Structural identification techniques are then employed to disentangle the contributions of fundamental shocks to movements in both the mean and various quantiles of the distributions. Our findings highlight the highly time-varying nature of downside risk to

GDP growth, primarily driven by demand and financial shocks. These shocks generate a strong interaction between the mean and volatility, resulting in a negatively skewed GDP growth distribution. This effect is particularly pronounced at times of heightened market distress. In contrast, supply shocks exhibit a more symmetric effect, influencing both the lower and the upper quantiles of GDP growth in a similar manner. For inflation, supply shocks drive instead a positive mean-volatility co-movement, where higher inflation is associated with increased uncertainty, causing a significant time variation in upside risk.

References

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (2019). "Vulnerable growth". In: *American Economic Review* 109.4, pp. 1263–1289.

Adrian, Tobias, Fernando Duarte, Nellie Liang, and Pawel Zabczyk (2020). "NKV: A new Keynesian model with vulnerability". In: *AEA Papers and Proceedings*. Vol. 110, pp. 470–476.

Adrian, Tobias, Federico Grinberg, Nellie Liang, Sheheryar Malik, and Jie Yu (2022). "The Term Structure of Growth-at-Risk". In: *American Economic Journal: Macroeconomics* 14.3, pp. 283–323.

Aikman, David, Kristina Bluwstein, and Sudipto Karmakar (2024). "A Tail of Three Occasionally Binding Constraints: A Modelling Approach to GDP-at-Risk". In: *IMF Economic Review*, pp. 1–37.

Antolín-Díaz, Juan and Juan F Rubio-Ramírez (2018). "Narrative sign restrictions for SVARs". In: *American Economic Review* 108.10, pp. 2802–2829.

Arias, Jonas E, Juan F Rubio-Ramirez, and Minchul Shin (2023). "Macroeconomic forecasting and variable ordering in multivariate stochastic volatility models". In: *Journal of Econometrics* 235.2, pp. 1054–1086.

Balke, Nathan S. (2000). "Credit and economic activity: credit regimes and nonlinear propagation of shocks". In: *Review of Economics and Statistics* 82.2, pp. 344–349.

Boire, Francois-Michel, Thibaut Duprey, and Alexander Ueberfeldt (2021). *Shaping the future: Policy shocks and the GDP growth distribution*. Staff Working Papers 21-24. Bank of Canada.

Breitenlechner, Maximilian, Martin Geiger, and Friedrich Sindermann (2019). *ZeroSignVAR: A zero and sign restriction algorithm implemented in MATLAB*.

Bruder, Stefan (2018). *Inference for structural impulse responses in SVAR-GARCH models*. ECON - Working Papers 281. Department of Economics - University of Zurich.

Brunnermeier, Markus K. and Yuliy Sannikov (2014). "A macroeconomic model with a financial sector". In: *American Economic Review* 104.2, pp. 379–421.

Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek (2016). "The macroeconomic impact of financial and uncertainty shocks". In: *European Economic Review* 88, pp. 185–207.

Caldara, Dario, Haroon Mumtaz, and Molin Zhong (2024). "Risk in a Data-Rich Model". Unpublished manuscript.

Caldara, Dario, Chiara Scotti, and Molin Zhong (2021). "Macroeconomic and financial risks: A tale of mean and volatility". In: *International Finance Discussion Paper* 1326.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2016). "Common drifting volatility in large Bayesian VARs". In: *Journal of Business & Economic Statistics* 34.3, pp. 375–390.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2018). "Measuring uncertainty and its impact on the economy". In: *Review of Economics and Statistics* 100.5, pp. 799–815.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2021). "Using time-varying volatility for identification in vector autoregressions: An application to endogenous uncertainty". In: *Journal of Econometrics* 225.1, pp. 47–73.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2024). "Capturing Macro-Economic Tail Risks with Bayesian Vector Autoregressions". In: *Journal of Money, Credit and Banking* 56.5, pp. 1099–1127.

Chan, Joshua, Eric Eisenstat, and Xuewen Yu (2022). "Large Bayesian VARs with Factor Stochastic Volatility: Identification, Order Invariance and Structural Analysis". Unpublished manuscript.

Chan, Joshua C. C., Gary Koop, and Xuewen Yu (2024). "Large Order-Invariant Bayesian VARs with Stochastic Volatility". In: *Journal of Business & Economic Statistics* 42.2, pp. 825–837.

Chavleishvili, Sulkhan and Simone Manganelli (2024). "Forecasting and stress testing with quantile vector autoregression". In: *Journal of Applied Econometrics* 39.1, pp. 66–85.

Chronopoulos, Ilias, Aristeidis Raftapostolos, and George Kapetanios (2024). "Forecasting value-at-risk using deep neural network quantile regression". In: *Journal of Financial Econometrics* 22.3, pp. 636–669.

Clark, Todd E., Florian Huber, Gary Koop, Massimiliano Marcellino, and Michael Pfarrhofer (2023). "Tail forecasting with multivariate Bayesian additive regression trees". In: *International Economic Review* 64.3, pp. 979–1022.

Cogley, Timothy and Thomas J. Sargent (2005). "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S". In: *Review of Economic Dynamics* 8.2, pp. 262–302.

De Santis, Roberto A. and Wouter Van der Veken (2022). *Deflationary financial shocks and inflationary uncertainty shocks: an SVAR Investigation*. Working Paper Series. European Central Bank.

Delle Monache, Davide, Andrea De Polis, and Ivan Petrella (2024). "Modeling and Forecasting Macroeconomic Downside Risk". In: *Journal of Business & Economic Statistics* 42.3, pp. 1010–1025.

Duprey, Thibaut and Alexander Ueberfeldt (2020). *Managing GDP Tail Risk*. Staff Working Papers 20-3. Bank of Canada.

Figueres, Juan Manuel and Marek Jarociński (2020). "Vulnerable growth in the euro area: Measuring the financial conditions". In: *Economics Letters* 191, p. 109126.

Forni, Mario, Luca Gambetti, Nicolò Maffei-Faccioli, and Luca Sala (2023). *The impact of financial shocks on the forecast distribution of output and inflation*. Working Paper. Norges Bank.

Forni, Mario, Luca Gambetti, and Luca Sala (2025). "Downside and upside uncertainty shocks". In: *Journal of the European Economic Association* 23.1, pp. 159–189.

Gertler, Mark (2020). "Comments and discussion on "When is growth at risk?" by Plagborg-Møller, M., Ricco, G., Reichlin, L., and Hasenzagl, T." In: *Brookings Papers on Economic Activity*.

Giglio, Stefano, Bryan Kelly, and Seth Pruitt (2016). "Systemic risk and the macroeconomy: An empirical evaluation". In: *Journal of Financial Economics* 119.3, pp. 457–471.

Goulet Coulombe, Philippe, Mikael Frenette, and Karin Klieber (2023). "From reactive to proactive volatility modeling with hemisphere neural networks". In: *Available at SSRN 4627773*.

Guerrieri, Luca and Matteo Iacoviello (2017). "Collateral constraints and macroeconomic asymmetries". In: *Journal of Monetary Economics* 90, pp. 28–49.

Guerrón-Quintana, Pablo, Alexey Khazanov, and Molin Zhong (2023). *Financial and Macroeconomic Data Through the Lens of a Nonlinear Dynamic Factor Model*. Finance and Economics Discussion Series 2023-027. Board of Governors of the Federal Reserve System (U.S.)

Huang, Yu-Fan, Wenting Liao, Sui Luo, and Jun Ma (2024). "Financial conditions, macroeconomic uncertainty, and macroeconomic tail risks". In: *Journal of Economic Dynamics and Control* 163, p. 104871.

Iseringhausen, Martin (2024). "A time-varying skewness model for Growth-at-Risk". In: *International Journal of Forecasting* 40.1, pp. 229–246.

Iseringhausen, Martin, Ivan Petrella, and Konstantinos Theodoridis (2023). "Aggregate skewness and the business cycle". In: *Review of Economics and Statistics*, pp. 1–37.

Koenker, Roger and Quanshui Zhao (1996). "Conditional Quantile Estimation and Inference for Arch Models". In: *Econometric Theory* 12.5, pp. 793–813.

Koenker, Roger W and Gilbert Bassett Jr. (1978). "Regression Quantiles". In: *Econometrica* 46.1, pp. 33–50.

Korobilis, Dimitris and Maximilian Schröder (2025). "Monitoring multi-country macroeconomic risk: A quantile factor-augmented vector autoregressive (QFAVAR) approach". In: *Journal of Econometrics* 249, p. 105730.

Kremer, Manfred (2016). "Macroeconomic effects of financial stress and the role of monetary policy: a VAR analysis for the euro area". In: *International Economics and Economic Policy* 13, pp. 105–138.

Kremer, Manfred, Marco Lo Duca, and Dániel Holló (2012). *CISS - a composite indicator of systemic stress in the financial system*. Working Paper Series. European Central Bank.

Krippner, Leo (2013). "Measuring the stance of monetary policy in zero lower bound environments". In: *Economics Letters* 118.1, pp. 135–138.

Lanne, Markku and Helmut Luetkepohl (2008). "Identifying monetary policy shocks via changes in volatility". In: *Journal of Money, Credit and Banking* 40.6, pp. 1131–1149.

Lanne, Markku and Pentti Saikkonen (2007). "A multivariate generalized orthogonal factor GARCH model". In: *Journal of Business & Economic Statistics* 25.1, pp. 61–75.

Lenza, Michele and Giorgio E. Primiceri (2022). "How to estimate a vector autoregression after March 2020". In: *Journal of Applied Econometrics* 37.4, pp. 688–699.

Lewis, Daniel J. (2021). "Identifying Shocks via Time-Varying Volatility". In: *The Review of Economic Studies* 88.6, pp. 3086–3124.

Loria, Francesca, Christian Matthes, and Donghai Zhang (2025). "Assessing macroeconomic tail risk". In: *The Economic Journal* 135.665, pp. 264–284.

Luetkepohl, Helmut and George Milunovich (2016). "Testing for identification in SVAR-GARCH models". In: *Journal of Economic Dynamics and Control* 73, pp. 241–258.

Montes-Galdón, Carlos and Eva Ortega (2022). *Skewed SVARs: tracking the structural sources of macroeconomic tail risks*. Working Papers. Banco de España.

Peersman, Gert (2005). "What caused the early millennium slowdown? Evidence based on vector autoregressions". In: *Journal of Applied Econometrics* 20.2, pp. 185–207.

Primiceri, Giorgio (2005). "Time varying structural vector autoregressions and monetary policy". In: *The Review of Economic Studies* 72.3, pp. 821–852.

Santos Silva, J. M. C. and F. N. Cardoso (2001). "The Chow-Lin method using dynamic models". In: *Economic Modelling* 18.2, pp. 269–280.

Uhlig, Harald (2005). "What are the effects of monetary policy on output? Results from an agnostic identification procedure". In: *Journal of Monetary Economics* 52.2, pp. 381–419.

Van der Weide, Roy (2002). "GO-GARCH: a multivariate generalized orthogonal GARCH model". In: *Journal of Applied Econometrics* 17.5, pp. 549–564.

Wolf, Elias (2021). "Estimating growth at risk with skewed stochastic volatility models". In: *Available at SSRN 4030094*.

A Additional Details on Structural Identification

This section complements section 2.3 by providing a more detailed explanation of the structural identification. The primary objective is to determine the impact matrix B_0, a procedure that consists of the following four steps.

I. Estimate the unconditional covariance matrix Similar to a homoscedastic VAR, structurally identifying the model requires connecting the covariance matrix of reduced-form errors to the structural parameters: the impact matrix B_0 and the time-varying variances of structural shocks Λ_t. To do so, we assume that the impact matrix B_0 is constant over the sample and, without loss of generality, that the structural error variances are normalized to one on average ($\mathbb{E}[\Lambda_t] = I$).[26] These assumptions allow us to relate the unconditional covariance matrix of reduced-form errors, $\Sigma_u^M = \mathbb{E}\left[\mathbf{u}_t\mathbf{u}_t'\right]$, directly to the impact matrix B_0 as follows:

$$\Sigma_u^M = \mathbb{E}[B_0\Lambda_t B_0'] = B_0\mathbb{E}[\Lambda_t]B_0' = B_0 I B_0' = B_0 B_0' \tag{A.1}$$

This approach is similar to methods used in the SVAR-GARCH literature, where using the unconditional covariance simplifies the estimation (see, e.g., Van der Weide, 2002; Lanne and Saikkonen, 2007; Luetkepohl and Milunovich, 2016).

Estimating Σ_u^M is therefore the first key step in the structural identification. Building on the assumption that the reduced-form errors are normally distributed with a time-varying standard deviation, we can define each element of Σ_u^M as:

$$\Sigma_{u,i,j}^M = \mathbb{E}\left[u_{i,t}u_{j,t}\right] = \mathbb{E}\left[\sigma_{i,t}\sigma_{j,t}u_{i,t}^*u_{j,t}^*\right] = \mathbb{E}\left[\sigma_{i,t}\sigma_{j,t}\right]\underbrace{\mathbb{E}\left[u_{i,t}^*u_{j,t}^*\right]}_{\rho_{i,j}^*} + \mathbb{C}\text{ov}[u_{i,t}^*u_{j,t}^*,\, \sigma_{i,t}\sigma_{j,t}], \tag{A.2}$$

where $u_{i,t}^*$ is defined as the reduced-form error normalized (i.e., divided) by its time-varying standard deviation $\sigma_{i,t}$, and $\rho_{i,j}^*$ is the time-invariant cross-correlation between the normalized reduced-form residuals.[27] Moreover, for the diagonal elements, this simplifies further:[28]

$$\Sigma_{u,i,i}^M = \mathbb{E}\left[u_{i,t}^2\right] = \mathbb{E}\left[\sigma_{i,t}^2\right] \tag{A.3}$$

Together, eq. (A.2) and eq. (A.3) allow us to estimate the unconditional covariance matrix Σ_u^M. In our application, these parameters are informed by the variance estimates from the

[26] This normalization is standard in the heteroscedastic VAR literature, as the variances and impact matrix are only identified up to scale. Note that this still allows for time-varying cross-correlations of the reduced-form errors, as these depend on the variances of the structural shocks.

[27] Since normalized errors have unit variance by construction, $\mathbb{E}\left[u_{i,t}^*u_{j,t}^*\right] = \mathbb{C}\text{ov}[u_{i,t}^*, u_{j,t}^*] = \rho_{i,j}^*$.

[28] Since $\rho_{i,i}^* = 1$ and $\mathbb{C}\text{ov}[u_{i,t}^*u_{i,t}^*, \sigma_{i,t}\sigma_{i,t}] = \mathbb{C}\text{ov}[(u_{i,t}^*)^2, \sigma_{i,t}^2] = \mathbb{E}\left[\frac{(u_{i,t})^2}{\sigma_{i,t}^2}\sigma_{i,t}^2\right] - \mathbb{E}\left[(u_{i,t}^*)^2\right]\mathbb{E}\left[\sigma_{i,t}^2\right] = \mathbb{E}\left[u_{i,t}^2\right] - \mathbb{E}\left[\sigma_{i,t}^2\right] = 0$.

quantile regressions (for the σ terms as described in section 2.3, and by the residuals u of the VAR for the covariance and ρ terms.

II. Identify a set of valid impact matrices While other standard VAR identification methods are also feasible in our setting, in our empirical application we use sign and narrative restrictions, leading us to identify a set of impact matrices $\{B_0(m)\}_{m=1}^{M}$ that are consistent with the decomposition of the unconditional variance-covariance matrix expressed in eq. (A.1). We describe in more detail our identification approach in sections 2.3 and 3.2.

III. Compute the implied scaling factors in sample As is standard in the structural VAR literature, reduced-form errors are linear combinations of the structural shocks. This fact, combined with the assumption that structural shocks are independent, implies that the variances of the reduced-form shocks are also linear combinations of the structural variances. This relationship forms a linear system that can be solved for the unknown structural variances, $\text{diag}(\Lambda_t)$:

$$\text{diag}(\Lambda_t) = (B_0 \odot B_0)^{-1}\text{diag}(\Sigma_{u,t}) \tag{A.4}$$

where $\odot$ is the Hadamard (element-wise) product.[29]

This step also acts as an identification restriction. The time-varying nature of the model requires that a selected B_0 yields admissible (i.e., non-negative) structural variances for the entire sample. Therefore, we restrict the set of valid impact matrices to ensure that $\text{diag}(\Lambda_t)$ is positive at all times t.[30]

IV. Select a unique impact matrix Since our identification approach relied on sign and narrative restrictions, we arrive at this step with a *set* of valid impact matrices, $\{B_0(n)\}_{n\in\mathcal{N}}$, rather than a unique B_0 matrix. We select a unique matrix from this set via maximum likelihood, inspired by Lanne and Luetkepohl (2008):

$$\begin{aligned}\max_{B_0\in\{B_0(n)\}_{n\in\mathcal{N}}} \log\mathcal{L} = &-\frac{nT}{2}\log 2\pi - \frac{1}{2}\sum_{t=1}^{T}\log\left(\det\left(B_0\Lambda_t B_0'\right)\right)\\ &-\frac{1}{2}\sum_{t=1}^{T}(\mathbf{y}_t - A\mathbf{y}_{t-1})'(B_0\Lambda_t B_0')^{-1}(\mathbf{y}_t - A\mathbf{y}_{t-1})\end{aligned}$$

where T is the number of periods and n is the number of endogenous variables. As mentioned in section 2.3, we exploit the model's heteroscedasticity to identify a unique impact matrix B_0 from the set of suitable candidates. As eq. (A.4) shows, alternative impact matrices B_0 lead to

[29] From eq. (2.5) in the main text, $\Sigma_{u,t} = B_0\Lambda_t B_0' = B_0\text{diag}(\Lambda_t)B_0'$. This implies $\text{diag}(\Sigma_{u,t}) = \text{diag}(B_0\text{diag}(\Lambda_t)B_0') = (B_0 \odot B_0)\text{diag}(\Lambda_t)$.

[30] For simplicity in the simulation exercises, any implied 'negative variance' was constrained to zero.

different in-sample structural variances Λ_t. Unlike in homoscedastic VARs, this time variation in variances makes the likelihood function sensitive to the parameters in B_0, allowing us to use maximum likelihood to pin down B_0 uniquely.

B Data

Variable	**Transformations**	**Source**
GDP	log of interpolated monthly real GDP, multiplied by 100	Quarterly Real GDP Euro Area (ECB), interpolated to monthly frequency with the method of Santos Silva and Cardoso (2001) using data on Industrial Production excluding the construction sector (ECB).
HICP	log of HICP multiplied by 100	Headline HICP (ECB)
STR	Short Term Rate, percent per annum, demeaned	We use the monthly average Euro Area Shadow Short Term rate Jan-1995:Sep-2022 from Krippner (2013) and linearly extrapolate with the 3-month Euribor.
CISS	square root as suggested by Kremer (2016), demeaned	Composite Indicator of Systemic Stress Jan-1987:Oct-2023 (ECB), extrapolated with VIX (ECB)

Table B1: Database

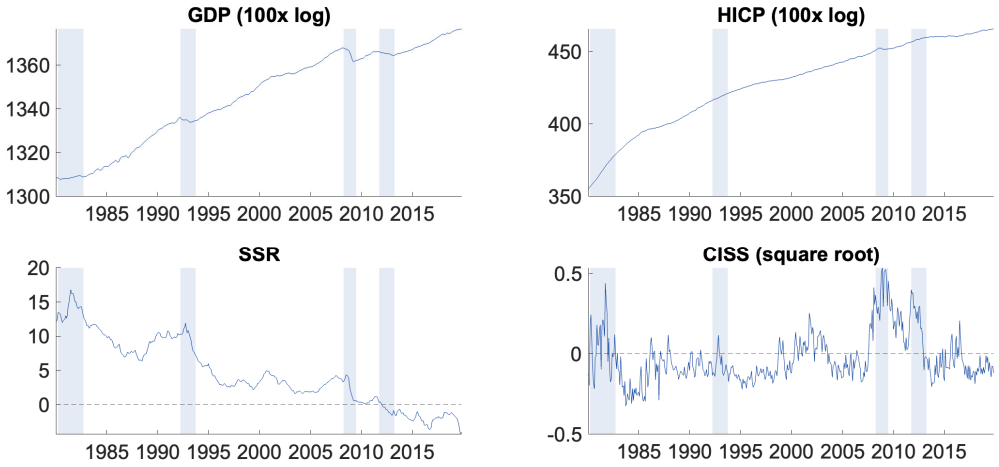

Figure B1: Database for the Euro Area VAR-QR Model

Note: Shaded areas depict euro area recessions dated by CEPR-EABCN. Series: log of euro area real GDP, log of euro area HICP, Shadow Short Term Rate, and the square root of CISS (demeaned).

C Additional Figures

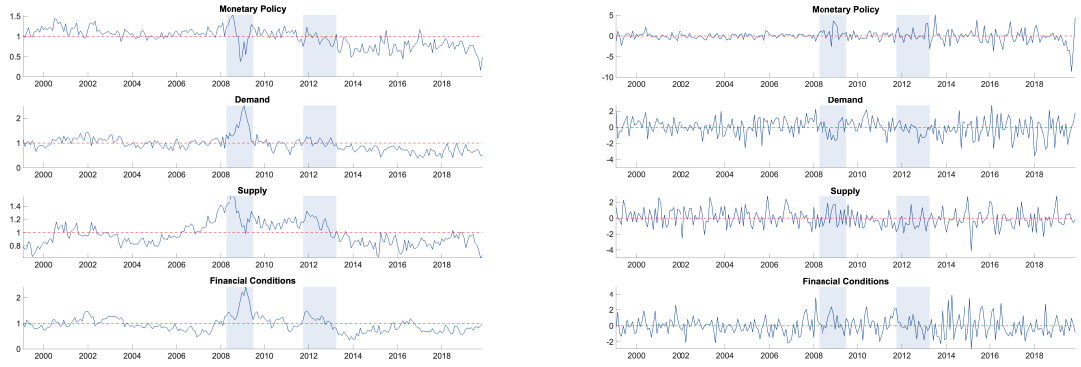

(a) Standard Deviations of Structural Shocks (b) Identified Structural Shocks

Figure C2: Standard deviations the structural shocks (panel a) and realisations of (unscaled) structural shocks (panel b).

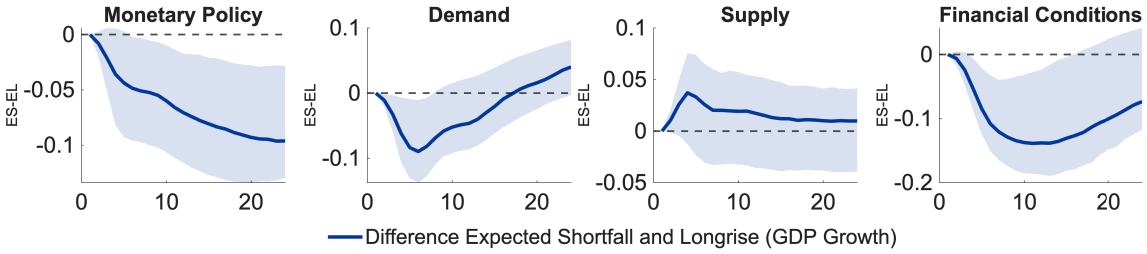

Figure C3: GIRFs to Uncertainty of GDP Growth

Note: The figure plot the response over time of the difference between expected shortfall and longrise in cumulative GDP growth following various structural shocks. The expected shortfall and longrise are calculated using the 5^{th} and 95^{th} percentiles, respectively, and the analysis considers one standard deviation structural shocks simulated as of June 2005. The blue-shaded area indicates the 90% confidence interval, constructed using the bootstrap method by Bruder (2018).

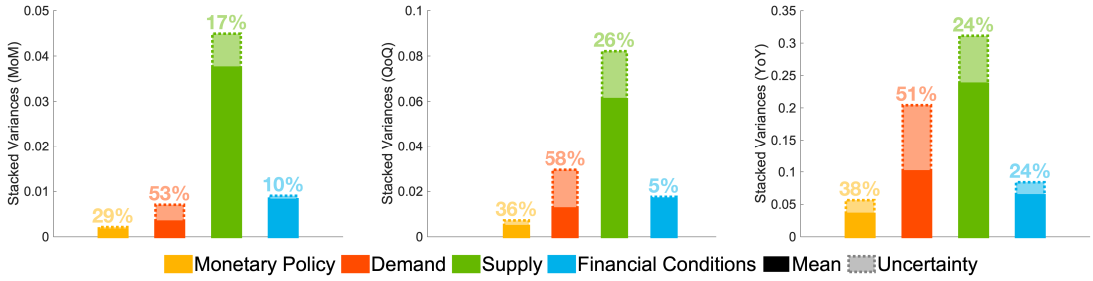

Figure C4: Mean and Uncertainty Channel: Expected Longrise of Inflation

Note: The figure illustrates the variance decomposition of the Expected Shortfall (ES) for inflation by shock and forecast horizon, highlighting the contribution from both the mean and uncertainty channels. The mean component is defined as $Mean_{t+h} = (\mathbb{E}_t[\Delta HICP_{t+h}] - \Delta HICP^{Trend}_{t+h})$, while the $Uncertainty_{t+h} = (\mathbb{E}_t[EL^{HICP}_{t+h}] - EL^{HICP,Trend}_{t+h} - Mean_{t+h})$, capturing the residual variation in the tail net of the mean component. Solid bars represent $= \mathbb{V}\text{ar}(Mean_{t+h})$, whereas the more transparent bars with a dashed outline represent $= \mathbb{V}\text{ar}(Uncertainty_{t+h}) + 2\mathbb{C}\text{ov}(Mean_{t+h}, Uncertainty_{t+h}))$ The percentages indicate the share of the uncertainty component to the total variance of the ES. The decomposition is conducted using the CFF method.

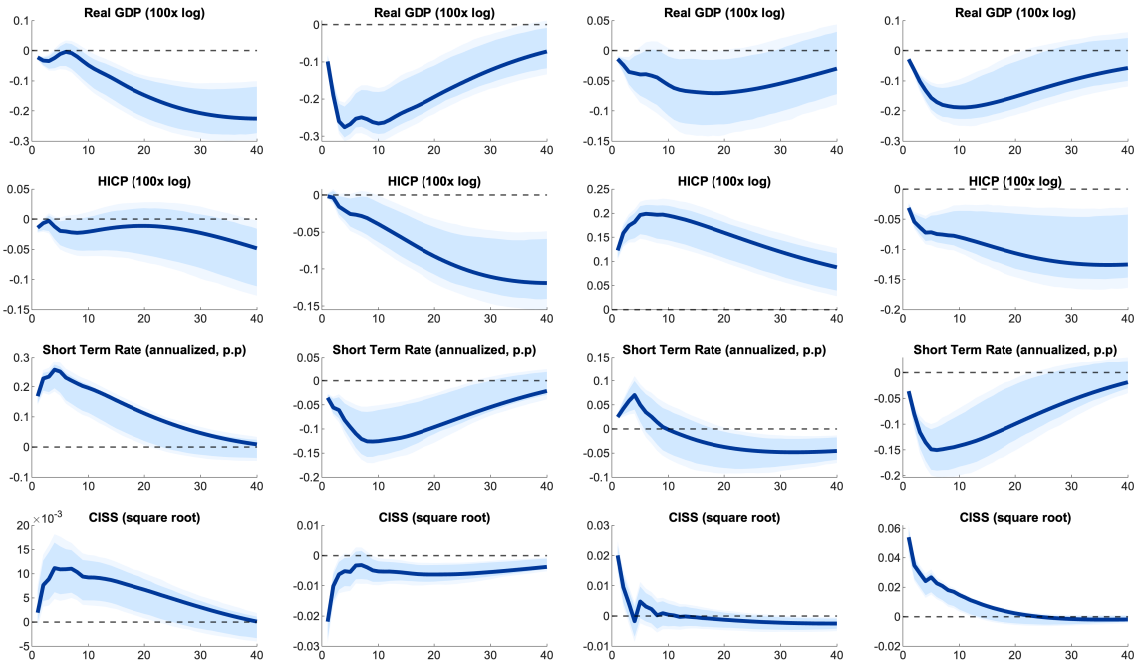

(a) Monetary Policy (b) Demand (c) Supply (d) Financial Conditions

Figure C5: Impulse Response Functions (Mean): GIRF in June 2005 (Calm Period).

Note: The dark blue solid lines are the GIRFs of the mean responses to a one standard deviation tightening structural shock as of June 2005. The blue-shaded areas are the 80% and 90% confidence intervals, capturing statistical uncertainty (based on the bootstrap method proposed by Bruder (2018)).

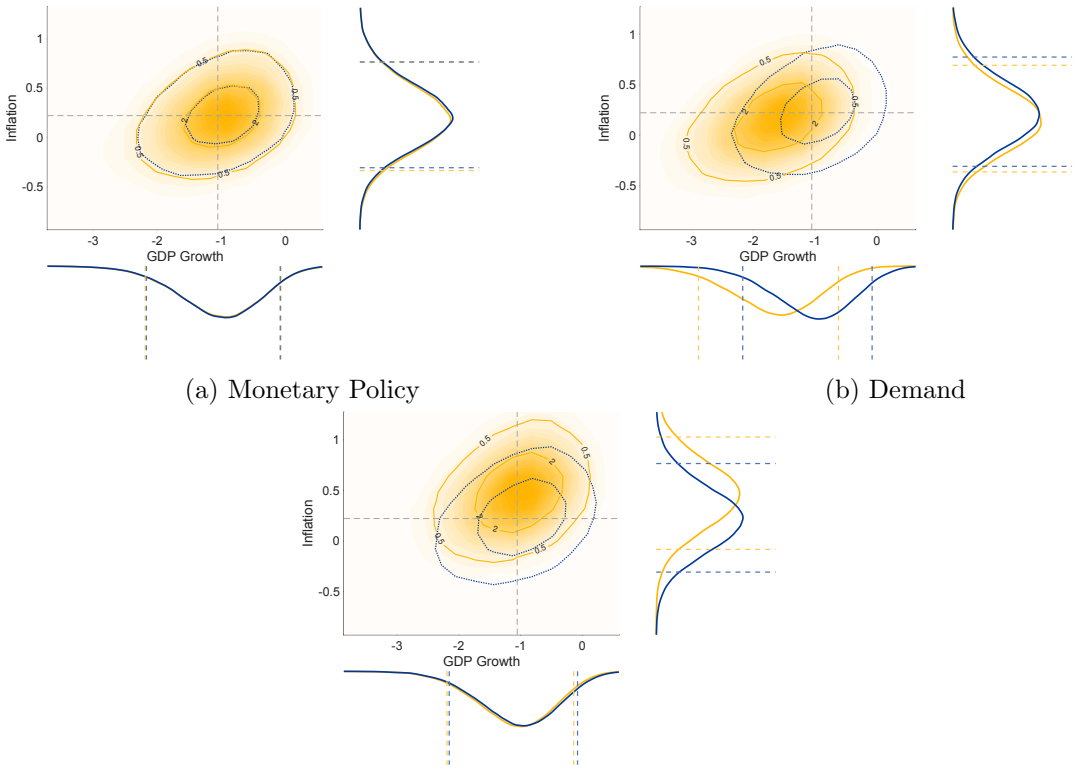

(a) Monetary Policy

(b) Demand

(c) Supply

Figure C6: Response of Joint Distribution of GDP Growth and Inflation

Note: The figure displays the simulated joint and marginal distributions of average quarter-on-quarter real GDP growth and inflation two quarters ahead, starting from January 2009. The blue lines refer to the baseline distribution, obtained by simulating the system under the combined impact of all shocks. The yellow distribution is derived by incorporating an additional two-standard-deviation structural shock. Dashed coloured lines indicate the 5^{th} and 95^{th} percentiles of the marginal distributions shown along the axes. The dashed grey lines indicate the conditional mean of the baseline distribution.

Acknowledgements

We thank Andrea Carriero, Luca Gambetti, Benedikt Kolb, Wolfgang Lemke, Massimiliano Marcellino, Rogier Quaedvlieg, internal ECB seminar participants, and Universität Leipzig seminar participants for helpful comments and suggestions.
The opinions in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank or the Eurosystem.

Giacomo Carboni
European Central Bank, Frankfurt am Main, Germany; email: giacomo.carboni@ecb.europa.eu

Luís Fonseca
European Central Bank, Frankfurt am Main, Germany; email: luis.fonseca@ecb.europa.eu

Fabio Fornari
European Central Bank, Frankfurt am Main, Germany; email: fabio.fornari@ecb.europa.eu

Leonardo Urrutia
Universität Leipzig, Leipzig, Germany; email: leon.urrutia@uni-leipzig.de

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899-7617-6 ISSN 1725-2806 doi:10.2866/9972273 QB-01-26-005-EN-N