

Web Appendix to
Macro-prudential policy under asymmetric risks:
A Bayesian Structural Quantile VAR approach*

*Sulkhan Chavleishvili,^(a) Robert F. Engle,^(b) Stephan Fahr,^(c)
Manfred Kremer,^(c) Frederik Lund-Thomsen,^(d) Simone Manganelli,^(c)
Bernd Schwaab^(c)*

^(a) Aarhus University, ^(b) NYU Stern School of Business,

^(c) European Central Bank, ^(d) Danmarks Nationalbank

*E-mail contacts: sulkhan.chavleishvili@econ.au.dk; rengle@stern.nyu.edu; stephan.fahr@ecb.europa.eu; manfred.kremer@ecb.europa.eu; folt@nationalbanken.dk; simone.manganelli@ecb.europa.eu; bernd.schwaab@ecb.europa.eu (corresponding author). The views expressed in this paper are those of the authors and do not necessarily reflect the views or policies of the European Central Bank or Danmarks Nationalbank.

Contents

A	Additional details on macro-prudential policies	1
A.1	Macro-prudential measures' impact on the financial cycle	1
A.2	Continued discussion of the policymaker's loss function	2
B	SQVAR model	4
C	Bayesian estimation	6
C.1	Gibbs sampler	6
C.2	Posterior density for λ	8
C.3	Specification of prior densities	9
D	Quantile IRFs, multi-step-ahead prediction, and counterfactual scenarios	12
D.1	Quantile impulse response functions	12
D.2	The SQVAR's predictive distribution: tree structure and resolution	15
D.3	SQVAR-based counterfactual scenarios	17
E	Euro area data and time series plots	19
E.1	GDP and inflation data	19
E.2	Composite indicator of systemic stress	19
E.3	The financial cycle	20
E.4	Short-term interest rates	20
E.5	Time series plots	21
F	Model adequacy & robustness checks	23
F.1	In- and out-of-sample assessment of predictive quantiles	23
F.2	Robustness to different priors for the euro area SQVAR's parameters	26
F.3	Robustness to a different hyperprior parameterization	27
F.4	Robustness to modeling the data at a monthly frequency	31

F.5	Robustness to adopting different sample endpoints	32
F.6	Robustness to removing the transmission lag	32
G	Additional euro area results	38
G.1	Euro area posterior estimates	38
G.2	Selected predictive quantiles for RGDP growth rates	44
H	U.S. posterior estimates	46

A Additional details on macro-prudential policies

A.1 Macro-prudential measures' impact on the financial cycle

For the counterfactual policy scenarios in Sections 5.4 – 5.5 to be meaningful, policymakers must have access to macro-prudential instruments that are effective in managing the financial cycle. Empirical evidence on their effectiveness, however, has only recently emerged. This section provides a brief and highly selective review of the relevant literature.

Macro-prudential instruments are commonly categorized by the targets they operate on. Capital-based instruments, such as counter-cyclical capital buffers, aim to increase bank capital during booms to absorb losses during downturns and curb excessive credit growth. Liquidity-based instruments, including liquidity coverage and net stable funding ratios, address maturity mismatches and funding risks. And borrower-based instruments, such as loan-to-value and debt-to-income limits, directly constrain household and firm leverage, often in real estate markets.

An academic consensus has formed that macro-prudential instruments can influence their respective targets, particularly in curbing credit growth and house price inflation (Cerutti et al., 2017). A meta-analysis by Araujo et al. (2020) points to statistically and economically significant effects on credit, despite considerable instrument heterogeneity, and finds weaker and less precise effects on house prices. Importantly, their analysis also suggests that tightening actions (leaning against the wind) have a stronger impact than loosening ones, likely due to imposing binding constraints on economic agents, akin to the "pushing on a string" metaphor in monetary policy. Borrower-based measures consistently prove most effective in moderating household credit growth and house price dynamics (Ampudia et al., 2021, Richter et al., 2019). The literature also indicates that macro-prudential tightening and the associated credit slowdown adversely affect economic activity primarily in the short term. However, these short-term costs may be outweighed by long-term benefits: by reducing the probability and severity of financial crises, macro-prudential policy could be linked to longer economic expansions, shallower recessions, and thus stronger, less volatile economic growth (Ampudia et al., 2021).

Despite evidence of their direct influence on credit and asset markets, the ability of macro-prudential policies to fully “tame” the financial cycle, as well as their potential side effects, remains actively debated. [Forbes \(2019\)](#) highlights several critical knowledge gaps. First, understanding the implications of leakages and cross-border spillovers is crucial. For example, domestic regulation might shift credit to less regulated entities or across borders, potentially creating new vulnerabilities. Second, the optimal calibration and combination of instruments (and approaches to identifying emerging risks as the financial system evolves) are largely unknown. In addition, long decision and implementation lags can hamper the effectiveness of macro-prudential policies. Macro-prudential policies may also face significant political economy constraints, given that the tangible costs of tighter policy are easily felt, while the benefits (crises that never materialize) are hard to quantify.

To sum up, the available empirical literature indicates that macro-prudential instruments can be effective in targeting specific vulnerabilities, especially credit growth and house prices, with borrower-based instruments being particularly potent. Especially when used in combination, macro-prudential instruments are increasingly seen as a crucial first line of defense against financial instability, helping build resilience and mitigate the build-up of systemic risk.

A.2 Continued discussion of the policymaker’s loss function

Several further issues regarding the macro-prudential policymaker’s objective function (6) / (7) deserve comment. First, while the specific mathematical form of our loss function is based on the framework in [Carney \(2020\)](#), we emphasize that the underlying principles and objectives that motivate this form are widely shared among policymakers. For example, also [Danthine \(2012\)](#) emphasizes the system-wide perspective of macro-prudential policy, distinguishing it from micro-prudential supervision. He highlights the need to understand how individual actions collectively generate systemic risk and how risk can build over time. This aligns with our focus on GS, which inherently captures system-wide downside risk and the potential for collective amplification of adverse shocks. While not explicitly providing a loss function, his emphasis on preventing se-

vere and painful economic contractions (as stated in our introduction) underpins the asymmetric weighting we apply, giving greater importance to avoiding large negative outcomes. Similarly, [Constancio \(2016\)](#) emphasizes that “*The ultimate objective of macro-prudential policy is to prevent and mitigate systemic risk, which includes strengthening the resilience of the financial system and smoothing the financial cycle, in order to preserve the effective provision of financial services to the real economy.*” He further stresses that “*Our interventions to strengthen the system have to be expected to support the real economy more in the bad times than they hold it back in the good.*” This sentiment directly supports the concept of balancing risk mitigation with growth considerations and, crucially, the asymmetric weighting of downside (GS) versus upside (GL) that forms the core of the loss function (6). The idea of ‘*smoothing the financial cycle*’ as a means to achieve this objective also reinforces our inclusion of the financial cycle in the SQVAR model below and our counterfactual policy simulations.

Second, one may question the applicability of a loss function articulated by a former Bank of England governor to the euro area economy. While national specificities and institutional frameworks differ, the fundamental objectives and underlying philosophy of macro-prudential policy are remarkably consistent across major advanced economies, particularly those that experienced the GFC (cf. [Yellen, 2010](#), [Tarullo, 2011](#), and [Ingves, 2012](#)). This shared experience has led to a significant convergence in macro-prudential policy design internationally. We emphasize that we do not assume identical institutional mandates or specific quantitative targets, but rather that the qualitative nature of the macro-prudential objective and the inherent trade-offs are sufficiently similar to warrant using this framework as a basis for a quantitative analysis in the euro area.

Finally, the objective function (6) is reminiscent of the mean with downside risk model in asset allocation; see e.g. [Fishburn \(1977\)](#). [Kilian and Manganelli \(2008\)](#) show that many existing downside risk measures are special cases of the downside risk notion proposed by [Fishburn \(1977\)](#). From a normative rather than positive perspective, [Suarez \(2022\)](#) provides a micro-foundation of a related objective function based on a representative agent with a CARA utility function on GDP and assumptions regarding the structure of the economy.

B SQVAR model

SQVAR system

We can reformulate the SQVAR model (9) as

$$x_t = \omega(\gamma) + A_0(\gamma)x_t + \sum_{j=1}^p A_j(\gamma)x_{t-j} + B(\gamma)d_t + \sum_{j=0}^p C_j(\gamma)z_{t-j} + \varepsilon_t^\gamma, \quad (\text{B.1})$$

for a given vector of quantiles $\gamma \equiv [\gamma_1, \dots, \gamma_n]'$ and state its identifying assumption as $Q_\gamma(\varepsilon_t^\gamma | \Omega_t) = 0_{n \times 1}$. We emphasize that matrix $A_0(\gamma)$ is lower-triangular for any γ . The identifying assumption $Q_\gamma(\varepsilon_t^\gamma | \Omega_t) = 0_{n \times 1}$ is a stacked (vectorized) version of n univariate quantile regressions' identifying assumptions $Q_{\gamma_i}(\varepsilon_{it}^{\gamma_i} | \Omega_{it}) = 0$, $i = 1, \dots, n$. Importantly, the model does not assume a particular statistical distribution (such as Gaussian or Student's t) for each error term ε_{it} . Instead, it assumes merely that $Q_{\gamma_i}(\varepsilon_{it}^{\gamma_i} | \Omega_{it}) = 0$, or, equivalently, that $F_i(\varepsilon_{it} < 0 | \Omega_{it}) = \gamma_i$ for some cumulative distribution function $F_i(\cdot)$. For completeness, we note that each exogenous variable z_{vt} , $v = 1, \dots, r$, is modeled as a univariate quantile autoregressive process,

$$z_{vt} = \mathcal{W}_v(\gamma_v^*) + \sum_{j=1}^p \mathcal{A}_{j,v}(\gamma_v^*)z_{v,t-j} + \mathcal{B}_v(\gamma_v^*)d_t + \xi_{vt}^{\gamma_v^*}, \quad (\text{B.2})$$

for $\gamma_v^* \in (0, 1)$.

Small SQVAR example

This section illustrates the general form of the SQVAR model (9) by providing a miniature model example. Specifically, we consider a bivariate model for the data vector $x_t = (y_t, s_t)'$, where y_t is the quarterly annualized real GDP growth and s_t is a coincident indicator of systemic financial stress. We consider one lag ($p = 1$), no exogenous variables, no deterministic terms beyond the constant, and only two quantiles: 0.1 for GDP, and 0.9 for financial stress. The system (9) can then

be written as

$$\begin{bmatrix} Q_{.1}(y_t|\Omega_{1t}) \\ Q_{.9}(s_t|\Omega_{2t}) \end{bmatrix} = \begin{bmatrix} \omega_1(.1) \\ \omega_2(.9) \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ a_{021}(.9) & 0 \end{bmatrix} \begin{bmatrix} y_t \\ s_t \end{bmatrix} + \begin{bmatrix} a_{11}(.1) & a_{12}(.1) \\ a_{21}(.9) & a_{22}(.9) \end{bmatrix} \begin{bmatrix} y_{t-1} \\ s_{t-1} \end{bmatrix}.$$

We note, in particular, the lower-triangular matrix A_0 , the contemporaneous term $x_t = (y_t, s_t)'$ on the right-hand side, and that different quantiles can be (are) considered for different variables. In this example, a unit-sized shock to y_t moves the 0.9 quantile of s_t by $a_{021}(.9)$, and a unit-sized shock to s_t moves the 0.1 quantile of y_{t+1} by $a_{12}(.1)$.

Comment on quantile crossing

Quantile-based statistical models have occasionally been criticized because their predictive quantiles can cross under certain circumstances. For example, in single-equation linear QR models with slope heterogeneity, the predictive quantiles will cross with certainty when the conditioning variables are sufficiently far from their center; see, e.g., [Koenker \(2005, Ch. 2.5\)](#). This does, of course, not imply that single-equation QR models are problematic: They can point to important asymmetries in-sample, see e.g. [Adrian et al. \(2019\)](#), and can do well in out-of-sample forecasting, see, e.g., [Korobilis \(2017\)](#). Quantile crossing is less of a concern in our SQVAR context because we will simulate data from the model, implicitly reordering them as in [Chernozhukov et al. \(2010\)](#), rather than focus on its nq one-step-ahead conditional quantile estimates. We note, in particular, that we do not need them as inputs to fit a skewed parametric density in a second step (cf. [Adrian et al., 2019](#)).

C Bayesian estimation

C.1 Gibbs sampler

We obtain posterior inference relying predominantly on established methods for Bayesian quantile regressions (see e.g. [Yu and Moyeed, 2001](#), [Kozumi and Kobayashi, 2011](#), [Khare and Hobert, 2012](#), and [Korobilis, 2017](#)). We estimate the SQVAR parameters in (9) equation by equation, which is possible given that $A_0(\gamma)$ is lower-triangular and as long as their prior density does not introduce cross-equation restrictions. We start by considering the endogenous variable x_i at quantile γ ,

$$x_{it} = w'_{it}\beta_i(\gamma) + \varepsilon_{it}^\gamma, \quad (\text{C.1})$$

where x_{it} is a scalar, $i = 1, \dots, n$, $t = 1, \dots, T$, w_{it} is a vector of regressors (in our case, contemporaneous values and lags of the endogenous and exogenous variables, a constant, as well as dummy variables), and $\beta_i(\gamma)$ is a vector of quantile-specific coefficients. The error term ε_{it}^γ is assumed to have an asymmetric Laplace distribution of the form

$$f(\varepsilon_{it}^\gamma) = \gamma(1 - \gamma) [\sigma_i(\gamma)]^{-1} e^{-\rho_\gamma(\varepsilon_{it}^\gamma)}, \quad (\text{C.2})$$

where $\sigma_i(\gamma)$ is a scale parameter, and $\rho_\gamma(\varepsilon) \equiv \varepsilon(\gamma - I(\varepsilon < 0))$ is the standard quantile regression check function; see also [Koenker and Bassett \(1978\)](#) and [Engle and Manganelli \(2004\)](#). It is clear from (C.2) that minimizing the usual quantile regression objective function, as e.g. defined in [Koenker and Bassett \(1978\)](#), is equivalent to maximizing a corresponding asymmetric Laplace log-likelihood; see e.g. [Yu and Moyeed \(2001\)](#).

An asymmetric Laplace random variable can be represented as a mixture of a standard normal and an exponential random variable ([Yu and Moyeed, 2001](#), [Kozumi and Kobayashi, 2011](#)). As a result, (C.1) can be restated as

$$x_{it} = w'_{it}\beta_i(\gamma) + \theta(\gamma)\nu_{it}(\gamma) + \tilde{\tau}(\gamma)\sqrt{\sigma_i(\gamma)\nu_{it}(\gamma)}u_{it}^\gamma, \quad (\text{C.3})$$

where $\tilde{\tau}^2(\gamma) = \frac{2}{\gamma(1-\gamma)}$, $\theta(\gamma) = \frac{1-2\gamma}{\gamma(1-\gamma)}$, $u_{it}^\gamma \sim N(0, 1)$, $\nu_{it}(\gamma) \sim \mathcal{E}(\sigma_i(\gamma))$, and $\mathcal{E}(\tilde{e})$ denotes an exponential distribution with mean \tilde{e} .

Since this section exclusively considers univariate regressions at a single quantile γ , we drop the quantile notation for convenience and leave the dependence implied. The mixture representation in (C.3) allows us to draw from the conditional posterior distributions using an appropriate set of prior distributions. We follow [Kozumi and Kobayashi \(2011\)](#) and choose the prior distributions

$$\beta_i \sim N(\underline{\mu}_{\beta,i}, \lambda_i \underline{\Sigma}_{\beta,i}), \quad \sigma_i \sim IG(\underline{\alpha}_{\sigma,i}, \underline{\zeta}_{\sigma,i}), \quad (\text{C.4})$$

where $IG(\alpha, \zeta)$ denotes an inverse-gamma distribution with shape and scale parameters α and ζ .

The addition of the scalar $\lambda_i \geq 0$ to the prior variance of β_i in (C.4) lets us control the looseness/tightness of the respective prior. A lower value of λ_i implies a tight, informative prior, while a high value of λ_i implies a loose, uninformative prior. This extra degree of latitude becomes particularly useful when considering regression quantiles in the tails of the distribution, where much weight is assigned to only a few observations. Similarly, a different degree of latitude may be appropriate for different macroeconomic and/or financial endogenous variables in the SQVAR. We give λ_i its own prior density and let $\lambda_i \sim IG(\underline{\alpha}_\lambda, \underline{\zeta}_\lambda)$.

We acknowledge that other prior choices for $\beta_i(\gamma)$ and $\sigma_i(\gamma)$ are possible. As one example, a global-local shrinkage prior, such as a horseshoe prior centered either on zero or the prior information, is one alternative for $\beta_i(\gamma)$. Alternatively, a Minnesota or sum-of-coefficients prior could be appropriate for $\beta_i(\gamma)$. Similarly, a tight ‘‘global’’ prior on $\lambda_i(\gamma)$, without an i subscript, or without a dependence on γ , could be an alternative hyperprior for the shrinkage parameter (but probably not a good one, as it leans against different degrees of prior variance that may be appropriate across quantiles and/or economic variables; see the discussion just above, and also [Web Appendix F.2](#) on prior robustness checks and [Web Appendix F.3](#) on a hyperprior robustness check).

With the above priors in place, and continuing to rely on [Kozumi and Kobayashi \(2011\)](#) and [Khare and Hobert \(2012\)](#), we end up with the following four-step Gibbs sampler for all nq equa-

tions, dropping subscript i for convenience. Steps 1 – 3 are identical to those in [Khare and Hobert \(2012\)](#). Step 4 is new but straightforward; Section [C.2](#) provides a derivation.

1. Draw $\sigma|x, w, \dots \sim IG(\bar{\alpha}_\sigma, \bar{\zeta}_\sigma)$,
 where $\bar{\alpha}_\sigma = \underline{\alpha}_\sigma + \frac{3}{2}T$, and $\bar{\zeta}_\sigma = \underline{\zeta}_\sigma + \sum_{t=1}^T \frac{(x_t - w_t' \beta - \theta \nu_t)^2}{2\bar{\tau}^2 \nu_t} + \sum_{t=1}^T \nu_t$.
2. Draw $\beta|x, w, \dots \sim N(\bar{\mu}_\beta, \bar{\Sigma}_\beta)$,
 where $\bar{\Sigma}_\beta^{-1} = \sum_{t=1}^T \frac{w_t w_t'}{\bar{\tau}^2 \sigma \nu_t} + \lambda^{-1} \underline{\Sigma}_\beta^{-1}$, and $\bar{\mu}_\beta = \bar{\Sigma}_\beta \left[\sum_{t=1}^T \frac{w_t (x_t - \theta \nu_t)}{\bar{\tau}^2 \sigma \nu_t} + \lambda^{-1} \underline{\Sigma}_\beta^{-1} \underline{\mu}_\beta \right]$.
3. Draw $\nu_t^{-1}|x_t, w_t, \dots \sim \mathcal{IGN}(\bar{\kappa}_1, \bar{\kappa}_2)$, where \mathcal{IGN} denotes the inverse Gaussian distribution, with $\bar{\kappa}_1 = \frac{\sqrt{\theta^2 + 2\bar{\tau}^2}}{|x_t - w_t' \beta|}$, and $\bar{\kappa}_2 = \frac{\theta^2 + 2\bar{\tau}^2}{\sigma \bar{\tau}^2}$.
4. Draw $\lambda|\beta, \dots \sim IG(\bar{\alpha}_\lambda, \bar{\zeta}_\lambda)$,
 where $\bar{\alpha}_\lambda = \underline{\alpha}_\lambda + \frac{\tilde{k}}{2}$, $\bar{\zeta}_\lambda = \underline{\zeta}_\lambda + \frac{1}{2} (\beta - \underline{\mu}_\beta)' \underline{\Sigma}_\beta^{-1} (\beta - \underline{\mu}_\beta)$, and \tilde{k} is the dimension of β (the number of regressors).

To conclude, we note that the Gibbs sampler of [Kozumi and Kobayashi \(2011\)](#) and [Khare and Hobert \(2012\)](#), steps 1–3 above, comes with theoretical guarantees: draws from the sampler converge to the intractable true posterior, and do so at a geometric rate. To prove this result, [Khare and Hobert \(2012\)](#) make no assumptions regarding the dimensions of the data (so that the result continues to hold even if \tilde{k} were large relative to T).

C.2 Posterior density for λ

This section derives the fourth step of the Gibbs sampler in Section [C.1](#). This step estimates a tightness parameter that determines the weight given to the U.S. parameter estimates in forming the priors for estimating the euro area model. Allowing the data to help decide on the tightness of the prior density across quantiles and variables appears particularly appropriate in our case, as we have no reliable prior knowledge of how informative the Minnesota prior is for the parameters of the U.S. model at any given quantile, nor how informative the U.S. posterior density is for the parameters of the euro area model. This issue is particularly acute when modeling tail quantiles

far from the center of the predictive distribution, and/or when modeling dissimilar variables in the SQVAR. Bayesian updating of the tightness parameter through a hyper-prior is a solution adopted, for example, in the literature on shrinkage in Bayesian regression models (see e.g. [Huber and Feldkircher, 2019](#) and [Korobilis and Pettenuzzo, 2019](#)), and in the analysis of sequential medical trials (see, e.g., [Ibrahim et al., 2015](#) and [Ibrahim and Chen, 2000](#)).

To arrive at the conditional posterior distribution for λ , again dropping subscripts γ and i for clarity, we start by writing out the kernel of the joint posterior distribution, $\mathcal{P}(\cdot)$, with the terms relating to λ ,

$$\mathcal{P}(\lambda|\cdot) \propto |\lambda \underline{\Sigma}_\beta|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\beta - \underline{\mu}_\beta)' \lambda^{-1} \underline{\Sigma}_\beta^{-1} (\beta - \underline{\mu}_\beta) \right\} \times \lambda^{-\alpha_\lambda - 1} \exp \left\{ -\underline{\zeta}_\lambda \lambda^{-1} \right\},$$

where the first term comes from the normally distributed prior for β and the second term from the inverse-gamma prior distribution for λ . Rewriting the above expression, we obtain

$$\begin{aligned} \mathcal{P}(\lambda|\cdot) &\propto \lambda^{-\alpha_\lambda - \frac{\tilde{k}}{2} - 1} \exp \left\{ -\lambda^{-1} \left(\frac{1}{2} (\beta - \underline{\mu}_\beta)' \underline{\Sigma}_\beta^{-1} (\beta - \underline{\mu}_\beta) + \underline{\zeta}_\lambda \right) \right\} \\ &\propto IG(\bar{\alpha}_\lambda, \bar{\zeta}_\lambda), \end{aligned}$$

where $\bar{\alpha}_\lambda = \alpha_\lambda + \frac{\tilde{k}}{2}$ and $\bar{\zeta}_\lambda = \underline{\zeta}_\lambda + \frac{1}{2} (\beta - \underline{\mu}_\beta)' \underline{\Sigma}_\beta^{-1} (\beta - \underline{\mu}_\beta)$, and \tilde{k} denotes the number of regressors. As a result, we can immediately draw λ from its posterior distribution conditional on a posterior draw of β .

C.3 Specification of prior densities

As explained in the main text, we first estimate the SQVAR model's parameters for U.S. data. We do so using the same model specification (i.e., variables, deterministic terms, number of lags and zero restrictions). Next, we let these estimates inform the euro area parameters' priors. We are thus taking advantage of decades of data available for the U.S. between 1973Q1 and 2022Q4.

Priors for U.S. data

We need to specify the prior parameters $\underline{\mu}_{\beta,i}^{US}$, $\underline{\alpha}_{\sigma,i}^{US}$, $\underline{\zeta}_{\sigma,i}^{US}$, $\underline{\alpha}_{\lambda,i}^{US}$, and $\underline{\zeta}_{\lambda,i}^{US}$. For U.S. data, we keep the prior parameters homogeneous across quantiles. We have therefore dropped a potential dependence on γ .

We employ a Minnesota prior for the vector of coefficients β_i^{US} (see, e.g., [Litterman, 1986](#), [Lütkepohl, 2005](#), p. 225, [Giannone et al., 2015](#)). This means that the prior density is Gaussian, and pointing to a persistent autoregressive process of order one. Specifically, we set the coefficient in β_i^{US} referring to variable i 's own first lag equal to either 0.9 or 1, depending on the variable. For the CISS, the financial cycle, and the real GDP growth rate, the own-lag coefficient is set to 0.9; for CPI inflation and the Federal Funds Rate the coefficient is set to one. All other elements of β_i^{US} are given a prior mean of zero. We further specify the Minnesota prior's covariance matrix $\underline{\Sigma}_{\beta,i}^{US}$ as diagonal, with $\underline{\sigma}_{\beta,ij,l}^{US}$ the element corresponding to the l th lag of the j th endogenous variable in equation i . The diagonal elements are given by

$$\begin{aligned}\underline{\sigma}_{\beta,ij,l}^{US} &= \phi_0^2 && \text{if } i = j, l = 0 \\ \underline{\sigma}_{\beta,ij,l}^{US} &= \left(\frac{\phi_0}{l\phi_3}\right)^2 && \text{if } i = j, l > 0 \\ \underline{\sigma}_{\beta,ij,l}^{US} &= \left(\phi_0\phi_1\frac{\hat{\sigma}_i}{\hat{\sigma}_j}\right)^2 && \text{if } i \neq j, l = 0 \\ \underline{\sigma}_{\beta,ij,l}^{US} &= \left(\frac{\phi_0\phi_1}{l\phi_3}\frac{\hat{\sigma}_i}{\hat{\sigma}_j}\right)^2 && \text{if } i \neq j, l > 0 \\ \underline{\sigma}_{\beta}^{US} &= (\phi_0\phi_2)^2 && \text{otherwise,}\end{aligned}$$

where ϕ_0 is a general tightness parameter, ϕ_1 a tightness parameter on endogenous variables other than variable i , ϕ_2 a tightness parameter on exogenous variables and deterministic terms, and ϕ_3 a tightness parameter controlling the importance of lags of endogenous variables. We choose $\phi_0 = 0.2$, $\phi_1 = 0.5$, $\phi_2 = 10^5$, and $\phi_3 = 1$. These are common choices in the literature (see, e.g., [Litterman, 1986](#)). Parameter $\hat{\sigma}_i$ denotes the standard error of the residuals from a univariate quantile autoregression on endogenous variable i at the median. Finally, we set $\underline{\alpha}_{\sigma,i}^{US} = \underline{\zeta}_{\sigma,i}^{US} = 0.01$, corresponding to a non-informative prior for the scale parameter σ . We also set $\underline{\alpha}_{\lambda,i}^{US} = 3$ and $\underline{\zeta}_{\lambda,i}^{US} = 6$, implying a prior mean and a prior standard deviation of 3 for λ_i .

Priors for euro area data

We set the euro area prior parameters for $\beta_i^{EA}(\gamma)$ to match the corresponding posterior moments obtained from U.S. data. Specifically,

$$\begin{aligned}\underline{\mu}_{\beta,i}^{EA}(\gamma) &= \frac{1}{N^S} \sum_{s=1}^{N^S} \hat{\beta}_{i,s}^{US}(\gamma) \\ \underline{\Sigma}_{\beta,i}^{EA}(\gamma) &= \frac{1}{N^S} \sum_{s=1}^{N^S} \left(\hat{\beta}_{i,s}^{US}(\gamma) - \underline{\mu}_{\beta,i}^{EA}(\gamma) \right)^2,\end{aligned}$$

where $\hat{\beta}_{i,s}^{US}$ are the N^S posterior draws of β_i^{US} . All other euro area hyperparameters are given values identical to their U.S. counterparts. Specifically, $\underline{\alpha}_{\sigma,i}^{EA} = \zeta_{\sigma,i}^{EA} = 0.01$, $\underline{\alpha}_{\lambda,i}^{EA} = 3$, and $\zeta_{\lambda,i}^{EA} = 6$.

D Quantile IRFs, multi-step-ahead prediction, and counterfactual scenarios

D.1 Quantile impulse response functions

This section presents the simulation algorithm used to obtain quantile impulse response functions (QIRFs). For a given set of coefficients, we require two sets of conditional distributions to compare: First, we require the conditional forecast distribution of the data in a baseline scenario without an initial shock. Second, we require the conditional forecast distribution of the data in a counterfactual scenario in which a single shock arrives in the first forecasting period. The QIRFs are obtained as the difference between the latter and the former distribution.

To fix the notation, we let H denote the forecast horizon, o the forecast origin, M the number of posterior draws required for posterior inference of the QIRFs, and S the number of forward simulations for each posterior draw. We assume that $\mathcal{Q} = \{0.05, 0.10, \dots, 0.90, 0.95\}$ is a sufficiently large set of $q = 19$ quantiles distributed symmetrically around the median. We obtain $N^S = 2,500$ draws of each $\beta_i(\gamma)$, after discarding a burn-in sample of $N^B = 2,500$. The algorithm proceeds as follows.

Algorithm:

1. **Obtain posterior draws.** Obtain and store N^S posterior draws for all the SQVAR's parameters in (B.1) and (B.2) at all quantiles, see Appendix C.1. We store them in a four-dimensional array of dimensions $[n + r] \times [p(n + r) + r + 1 + k] \times q \times N^S$.
2. **Choose forecast origin.** Fix the initial conditions for the endogenous variables $x_{o-(p-1):o}$ and exogenous variables $z_{o-(p-1):o}$. We use the unconditional median over the estimation sample for all variables.
3. Set $m = 1$.

3.1. **Draw parameters.** Let $\hat{\beta}^{(m)}$ be a random draw from the set of N^S posterior draws, and $\hat{\beta}_i^{(m)}(\tilde{\gamma})$ the subset of parameters corresponding to the posterior draw of β_i for some quantile $\tilde{\gamma}$.

3.2. Let $\hat{\epsilon}_{o+1}$ be an $n \times 1$ vector of zeros.

3.3. Set $s = 1$.

3.3.1. Set $h = 1$.

3.3.1.1. **Draw quantiles at random.** Obtain n random draws from the uniform distribution $U(0, 1)$ and map them to the corresponding quantiles in \mathcal{Q} based on proximity. Stack the mapped quantiles in the $n \times 1$ vector ϱ . Define the $r \times 1$ vector ϱ^* in a similar fashion.

3.3.1.2. **Set up SQVAR system matrices.** Let $f_{A_j}(\hat{\beta}_i^{(m)}(\tilde{\gamma}))$ be a mapping from posterior draw $\hat{\beta}_i^{(m)}(\tilde{\gamma})$ to the i th row of A_j in (B.1). Define similar mappings for the remaining matrix coefficients. Using these mappings, stack the variable-specific posterior draws into matrices of quantile coefficients, such that

$$\tilde{A}_j^{o+h} = \begin{bmatrix} f_{A_j}(\hat{\beta}_1^{(m)}(\varrho_1)) \\ \vdots \\ f_{A_j}(\hat{\beta}_i^{(m)}(\varrho_i)) \\ \vdots \\ f_{A_j}(\hat{\beta}_n^{(m)}(\varrho_n)) \end{bmatrix} \quad \tilde{\omega}^{o+h} = \begin{bmatrix} f_{\omega}(\hat{\beta}_1^{(m)}(\varrho_1)) \\ \vdots \\ f_{\omega}(\hat{\beta}_i^{(m)}(\varrho_i)) \\ \vdots \\ f_{\omega}(\hat{\beta}_n^{(m)}(\varrho_n)) \end{bmatrix}$$

$$\tilde{C}_j^{o+h} = \begin{bmatrix} f_{C_j}(\hat{\beta}_1^{(m)}(\varrho_1)) \\ \vdots \\ f_{C_j}(\hat{\beta}_i^{(m)}(\varrho_i)) \\ \vdots \\ f_{C_j}(\hat{\beta}_n^{(m)}(\varrho_n)) \end{bmatrix} \quad \tilde{B}^{o+h} = \begin{bmatrix} f_B(\hat{\beta}_1^{(m)}(\varrho_1)) \\ \vdots \\ f_B(\hat{\beta}_i^{(m)}(\varrho_i)) \\ \vdots \\ f_B(\hat{\beta}_n^{(m)}(\varrho_n)) \end{bmatrix}$$

for $j = 0, \dots, p$. Define the corresponding matrices for the exogenous vari-

ables analogously using ϱ^* , resulting in matrices $\tilde{\mathcal{W}}_v^{o+h}$, $\tilde{\mathcal{A}}_{j,v}^{o+h}$ and $\tilde{\mathcal{B}}_v^{o+h}$.

3.3.1.3. Iterate exogenous variables forward. Compute the conditional forecast of each of the r exogenous variables, $z_{o+h}^{(s)}$, using (B.2) and the relevant quantile coefficients determined in the previous step as

$$z_{v,o+h}^{(s)} = \tilde{\mathcal{W}}_v^{o+h} + \sum_{j=1}^p \tilde{\mathcal{A}}_{j,v}^{o+h} z_{v,o+h-j}^{(s)} + \tilde{\mathcal{B}}_v^{o+h} d_{t+h}.$$

3.3.1.4. Iterate endogenous variables forward. Compute the conditional forecast of $x_{o+h}^{(s)}$ using (B.1) and the relevant quantile coefficients determined in step 3.3.1.2 as

$$x_{o+h}^{(s)} = \left(I - \tilde{A}_0^{o+h} \right)^{-1} \left[\begin{array}{c} \tilde{\omega}^{o+h} + \sum_{j=1}^p \tilde{A}_j^{o+h} x_{o+h-j}^{(s)} \\ + \tilde{B}^{o+h} d_{t+h} + \sum_{j=0}^p \tilde{C}_j^{o+h} z_{o+h-j}^{(s)} + \hat{\epsilon}_{o+h} \end{array} \right]$$

3.3.1.5. If $h < H$, set $h = h + 1$ and return to step 3.3.1.1.

3.3.2. If $s < S$, set $s = s + 1$ and return to step 3.3.1.

3.4. Obtain predicted quantiles. Let $\check{\mathbf{x}}_{o+h} = \{x_{o+h}^{(s)}\}_{s=1}^S$ be the set of S simulated forecasts of x_{o+h} . Compute the γ th conditional quantile forecast of $x_{i,o+h}$ as

$$Q_\gamma(x_{i,o+h} | \Omega_o) = Q_\gamma(\check{\mathbf{x}}_{i,o+h}),$$

where $Q_\gamma(\cdot | \Omega_o)$ is the γ th quantile conditional on the information set Ω_o , and letting $Q_\gamma(\cdot)$ denote the empirical quantile function. This concludes the baseline (no shock) forward simulations.

3.5. Choose the shock of interest. Let \mathcal{S} be an $n \times 1$ selection vector picking the endogenous variable i to be shocked. (That is, the i th element of \mathcal{S} is 1 and 0 otherwise.) Redefine $\hat{\epsilon}_{o+1} = \psi \mathcal{S}$, where ψ is a scalar. A common choice of magnitude ψ is the estimated standard deviation of quantile shocks obtained at the median. Repeat steps

3.3 - 3.4 to obtain the quantile projection conditional on the shock, $Q_\gamma(x_{i,o+h}|\Omega_o, \psi S)$.

3.6. **Compute the γ th quantile impulse response function as**

$$l_{\gamma,i,o+h}^{(m)} = Q_\gamma(x_{i,o+h}|\Omega_o, \psi S) - Q_\gamma(x_{i,o+h}|\Omega_o)$$

3.7. If $m < M$, set $m = m + 1$ and return to step 3.2.

Posterior inference of the QIRF of the conditional quantile γ of variable i in period $o+h$ is obtained from the sequence $\left\{ l_{\gamma,i,o+h}^{(m)} \right\}_{m=1}^M$.

D.2 The SQVAR's predictive distribution: tree structure and resolution

Deriving the SQVAR (9)'s multi-step-ahead predictive distribution requires a sequential process that rapidly increases in complexity. A “tree structure” arises because, at each forecast step, we are interested in multiple quantiles $(\gamma_1, \dots, \gamma_q)$ for each endogenous variable. To build the predictive distribution at horizon H , we recursively propagate the quantile forecasts through the model; see Section 3.4. For a model with n endogenous variables and q estimated quantiles, if we consider all possible combinations of quantiles at each step, the number of potential “paths” through the tree to a given forecast horizon H grows exponentially. Specifically, the number of possible trajectories of the entire vector of endogenous variables is $(q^n)^H$. Even for moderate values (e.g., $q = 19$ quantiles, $n = 5$ variables, and a forecasting horizon of $H = 4$ quarters), the number of possible paths becomes astronomically large, making it computationally infeasible to explicitly enumerate and store every single terminal node of the “tree.”

Monte Carlo sampling becomes necessary to characterize the predictive distribution as a result. Instead of attempting to enumerate every discrete point in the high-dimensional, expanding tree, we generate a large number of random draws (paths) through the tree structure. Each simulated path represents a possible future realization of the system. By aggregating the simulated paths, we can then efficiently approximate the full predictive distribution and extract any desired predictive quantiles (e.g., the 5th, 50th, 95th percentiles of the forecasted variable of interest for a given

quarter). This is the standard approach for obtaining predictive distributions from non-linear, state-dependent models where analytical solutions are intractable. Although the underlying structure is indeed discrete, due to the finite number of estimated quantiles, the simulation process effectively provides a refined and smooth approximation of this complex distribution.

The SQVAR framework, by estimating a finite number of conditional quantiles, implies a discrete (albeit highly resolved) representation of the predictive distribution. The “resolution” of the predictive density is determined by the number of estimated quantiles, the number of endogenous variables, and the forecast horizon. In our application, we use $q = 19$ quantiles per variable. For a single variable at a single step, this gives 19 discrete points. However, as we combine these conditional quantiles across $n = 5$ endogenous variables and propagate them over horizons up to $H = 20$ quarters, the number of potential paths – and thus potential “unique values” or points in the discretized joint predictive distribution – quickly becomes immense. For instance, considering only the 19 quantiles for a single variable, at $H = 20$, there are $19^{20} \approx 3.76 \times 10^{25}$ possible paths, which is an extraordinarily large number.

Despite the underlying discrete nature, when a sufficiently large number of simulations are performed (e.g., tens of thousands as in our study), the resulting predictive distributions for individual variables appear continuous and smooth. As illustrated in the selected predictive distributions plotted and discussed in Web Appendix G.2, these distributions typically exhibit asymmetry, which is a key advantage of our quantile framework over mean-based models. They are often skewed, reflecting higher downside risks or upside potentials depending on the state of the economy. Importantly, we do not observe significant “clumping” or “bimodality” in the GDP growth predictive distributions in our empirical application. The distributions tend to be unimodal, albeit asymmetric. This suggests that the chosen number of quantiles (19) and the simulation methodology provide a sufficiently rich and smooth representation of the underlying dynamics without artificial discontinuities.

D.3 SQVAR-based counterfactual scenarios

Rather than moving through the tree of potential future values of x_{t+h} completely at random, we may sometimes wish to consider only a subset of paths, or even just one path, in isolation. Such paths can also be thought of as a ‘counterfactual scenario,’ or model-based thought experiment, that conditions on an arbitrary but fixed sequence of future quantile realizations.


To this effect, let Γ^* denote a design matrix of size $H \times (n + r)$. Each element of Γ^* is either empty or in the interval $(0, 1)$. Any non-empty h th element of the i th column designates the quantile path to be traveled by variable i in period h in all S simulations. An empty element means that the path is chosen at random. The matrix Γ^* can then be readily applied in Step 3.3.1.1 in the simulation algorithm of Section D.1 to obtain density forecasts in a counterfactual scenario.

An illustration of this principle is given in Figure D.1. There, a single variable, y , is projected forward for $H = 3$ periods using two estimated quantiles, $\gamma \in \{u, d\}$. An unrestricted projection as in Figure D.1a yields a total of $2^3 = 8$ unique paths for y to travel along. Consider now the counterfactual, in which we require y to initially increase, then decrease, and then finally increase again. This maps to the path satisfying $\Gamma^* = \{u, d, u\}'$, highlighted by the red dashed lines in Figure D.1b.

Instead of a fully restricted tree, as illustrated above, one may also consider a partially restricted tree. In a partially restricted tree, the quantile paths are only fixed for some of the periods or variables. In the example above, one may, for instance, require that the chosen quantile from period 1 to 2 is always $\{d\}$, while leaving the remaining branches unrestricted. In this case, the total number of paths that can be traveled equals four $(Q^{u,d,u}, Q^{u,d,d}, Q^{d,d,u}, Q^{d,d,d})$.

Figure D.1: Illustration of counterfactual scenario analysis through quantile restrictions

Left panel: Solid blue lines indicate unrestricted paths. Right panel: The red dashed line indicates a restricted path, while light blue lines indicate paths excluded by the imposed restrictions.



E Euro area data and time series plots

E.1 GDP and inflation data

Sections 2 and 5 aim to provide analytical support for macro-prudential policies that safeguard the stability of the entire financial system to support the real economy's overall welfare. Real Gross Domestic Product (GDP) is the most comprehensive and policy-relevant measure of aggregate economic activity, encompassing all sectors of the economy, including the crucial services sector. Macro-prudential authorities, such as the European Central Bank, primarily monitor GDP as their key indicator of real activity when assessing systemic risks and their potential impact.

We use pre-1999 macro-financial time series data for the euro area when available. Such pre-1999 data were urgently needed for monetary policy analysis during the ECB's early years. As a result, counterfactual data were constructed "as if" the euro area had already existed earlier. Pre-1999 euro area data is publicly available;¹ see e.g. [Fagan et al. \(2001\)](#). We obtain quarterly real GDP and monthly consumer price index data between 1990Q1 and 1998Q4 from this source and splice it with official Eurostat data on real GDP and the Harmonised Index of Consumer Prices (HICP) between 1999Q1 and 2022Q4. Quarterly consumer prices are then computed as averages of monthly index levels. This results in $T = 132$.

E.2 Composite indicator of systemic stress

As a measure of system-wide financial distress, we use the revised daily version of the ECB's composite indicator of systemic stress (CISS) as introduced and described in [Chavleishvili and Kremer \(2025\)](#). This version of the CISS includes 15 individual market-based financial stress indicators that cover the main segments of a typical modern financial system: financial intermediaries, money markets, equity markets, bond markets, and foreign exchange markets. The 15 indicators are aggregated into a single statistic in a way that takes their time-varying cross-correlations into account. As a result, the CISS takes higher values when stress prevails in most market segments

¹<https://eabcn.org/page/area-wide-model>.

at the same time, capturing the idea that financial stress is more systemic, and more dangerous for the economy as a whole, whenever financial instability spreads widely across different parts of the financial system. The CISS for the euro area (as well as that for the U.S. and other countries) is updated daily and publicly available via the ECB’s Statistical Data Warehouse.²

E.3 The financial cycle

The financial cycle indicator used in the empirical analysis is based on [Lang et al. \(2019\)](#). It is designed to capture risks stemming from domestic (real) credit volumes, real estate markets, asset prices, and external imbalances. [Lang et al. \(2019\)](#) demonstrate that the indicator increases, on average, three to four years before the onset of systemic financial crises and the ensuing economic recession, and that its early warning properties for euro area countries are superior to those of the total credit-to-GDP gap, a popular alternative referred to in Basel-III regulation. As a result, the financial cycle measurement offers useful information about both the probability and the likely cost of systemic financial crises several years in advance. In our model, a systemic financial crisis entails a sharp increase in the CISS and a subsequent large drop in real GDP.

Figure [E.1](#) below provides a time series plot of the financial cycle indicator. It takes high values during the dot-com boom years between 1997 and 2000 and during the credit boom years preceding the 2008–2009 global financial crisis. The financial cycle takes particularly low values in 2009 and 2011, at times associated with crisis-induced fire sales and financial system deleveraging.

E.4 Short-term interest rates

To construct a consistent time series of short-term euro area interest rates, we splice together three time series, as follows. From 1999Q1 to 2022Q4, we use quarterly averages of the three-months-ahead euro Overnight Index Swap (OIS) rate, a risk-free interest rate. Between 1994Q1 and 1998Q4, we use the three-month Euro Interbank Offered Rate (EURIBOR), corrected for the average EURIBOR-OIS spread between 2000Q1 and 2007Q2. Finally, between 1990Q1 and


²<https://sdw.ecb.europa.eu/>

1993Q4, we use the German Frankfurt Interbank Offered Rate (FIBOR), additionally adjusted for the average FIBOR-EURIBOR spread between 1994Q1 and 1998Q4. All interest rate data are taken from London Stock Exchange Group. Web Appendix [E.5](#) provides a time series plot.

E.5 Time series plots

Figure E.1: Euro area and U.S. time series

Time series data plots. Annualized inflation and growth rates are calculated as 400 times the log difference between quarterly averages. Sources: European Central Bank, U.S. Bureau of Labor Statistics, U.S. Bureau of Economic Analysis, Board of Governors of the Federal Reserve System, LSEG, Haver Analytics, and authors' calculations.



F Model adequacy & robustness checks

F.1 In- and out-of-sample assessment of predictive quantiles

This section provides a selective in- and out-of-sample assessment of the SQVAR model’s predictive density.

In-sample assessment: Maximizing the asymmetric Laplace likelihood function (C.2) (or, equivalently, minimizing the Quantile Regression “check function;” see the discussion below (C.2)) virtually guarantees the correct calibration of the time-varying one-step-ahead conditional quantiles for all endogenous variables in-sample. Put differently, getting the unconditional coverage (of the model’s one-step-ahead time-varying quantiles) right is what the QR objective function is designed to deliver.

In a Bayesian setting, the correct calibration of all one-step-ahead in-sample conditional quantiles will continue to hold provided the model parameters’ priors are not particularly dogmatic/ill-chosen. For the full-sample posterior mean estimates used in Sections 5.3 – 5.5, we report that 90,4% of the $T = 132$ in-sample observations for euro area real GDP growth lie above the model’s estimated time-varying .1 quantile for that variable, 75,0% lie above the model’s time-varying .25 quantile, 51,9% lie above the model’s time-varying median, 26,9% lie above its time-varying .75 quantile, and 9,6% lie above its 0.9 quantile. As a result, the SQVAR’s one-step-ahead predictive quantiles for euro area real GDP growth are virtually perfectly calibrated in-sample.

Out-of-sample assessment: In principle, many variables, forecasting horizons, alternative models for comparison, and different prior specifications for each of these models’ parameters could be considered in an all-encompassing out-of-sample assessment exercise; see e.g. [Surprenant \(2025\)](#) for a welcome and important first step in this direction. To keep the discussion on point, we continue to focus on the above five quantiles of the SQVAR model’s one-step-ahead predictive density for euro area real GDP growth, and study whether nominal and out-of-sample empirical confidence levels are approximately aligned. Real GDP growth is the key variable of interest in the empirical study presented in Section 5.

To assess the calibration of our quantile forecasts, we conduct an out-of-sample evaluation using the Dynamic Quantile (DQ) test proposed by [Engle and Manganelli \(2004\)](#), which in turn builds on insights by [Kupiec \(2000\)](#) and [Christoffersen \(1998\)](#). When discussing testing outcomes, this section implicitly (and exceptionally) adopts a frequentist perspective.

The core intuition behind the DQ test is that, for a quantile forecast to be well-calibrated, two conditions must hold: unconditional coverage and independence. Unconditional coverage implies that the actual proportion of observations falling below the γ -quantile forecast should be equal to γ . Independence implies that whether an observation falls below the quantile forecast should be independent of all information available at the time the forecast is made.

The DQ test is constructed as follows. The “exceedance” variable $Hit_{\gamma,t+1}$ is defined as

$$Hit_{\gamma,t+1} \equiv I(x_{t+1} < Q_{\gamma}(x_{t+1} | \Omega_t)) - \gamma,$$

where $I(\cdot)$ is the indicator function, taking value 1 if the condition in parentheses is true and zero otherwise. In this context, $Q_{\gamma}(x_{t+1} | \Omega_t)$ represents the γ quantile forecast based on the information set Ω_t at time t . If the quantile forecast is correctly calibrated, the sequence $Hit_{\gamma,t+1}$ should be a martingale difference sequence with mean zero. To test this, we regress $Hit_{\gamma,t+1}$ on a constant, one lag of the exceedance variable, and the quantile forecast itself. The regression is limited to a single lag of the exceedance variable due to the relatively small out-of-sample period. Since the quantile forecast is derived from variables in the information set, the quantile itself should not provide any predictive power for the exceedance in the subsequent period. If the quantile model is correctly calibrated, none of the regression coefficients should be significantly different from zero. This can be tested using a standard F-test.

We applied the DQ test to the one-step-ahead quantile forecasts for euro area real GDP. The model was estimated using data up to 2008Q4, and the DQ test was conducted on an out-of-sample basis for the subsequent period, from 2009Q1 to 2022Q4, yielding a total of 56 out-of-sample observations. [Table F.1](#) reports the p-values from the F-test for three specifications of the DQ test.

Table F.1: **Dynamic quantile tests**

P-values of [Engle and Manganelli \(2004\)](#)'s dynamic quantile tests for five quantiles of the one-step-ahead predictive real GDP growth distribution. Estimation sample: 1991Q1 – 2008Q4. Evaluation period: 2009Q1 – 2022Q4.

Model	Q10	Q25	Q50	Q75	Q90
1. Constant	0.07	0.35	0.35	0.03	0.00
2. Constant + hit lag	0.43	0.97	1.00	0.95	0.57
3. Constant + hit lag + quantile forecast	0.10	0.18	0.14	0.14	0.07

The first specification (“1. Constant”) tests only for unconditional coverage and is equivalent to [Kupiec \(2000\)](#)'s test. The second specification (“2. Constant + hit lag”) incorporates a test for the independence of hits (similar to [Christoffersen \(1998\)](#)'s test). Finally, the third specification (“3. Constant + hit lag + quantile forecast”) constitutes the full DQ test, where the quantile forecast itself is included as an explanatory variable. Since the quantile forecast is derived from variables in the information set, the quantile forecast should not provide any additional predictive power for the exceedance in the subsequent period if the model is well-calibrated.

Table F.1 indicates the following: For unconditional coverage (first row), the upper tail (Q75, Q90) quantiles' p-values are below 5%, suggesting a departure from unconditional coverage in the right tail over the out-of-sample evaluation period. A closer inspection suggests that the predictive density's dispersion may be on the low side following the 2020 Covid-19 recession, which in turn may be related to our decision to use Covid-dummies to neutralize their effect on the full-sample posterior estimates, coupled with the energy crisis / Russian war of aggression against Ukraine that follows shortly afterwards in 2022. When we test for independence of hits (second row), the p-values are high across all quantiles, ranging from 0.43 to 1.00. This result indicates that the quantile exceedances are largely independent and that the model captures the dynamic dependencies well, without a systematic pattern of under- or over-predicting the tails. For the full DQ test (third row), the p-values for all five quantiles are above a conventional 5% significance level, indicating a good calibration over the entire support of the predictive distribution.

We emphasize that the model is not re-estimated during the out-of-sample period. The fixed

estimation window (up to 2008Q4) means that the model does not adapt to structural changes or new information that emerged post-2008. Given the significant shocks in the evaluation period, we interpret these results as pointing to an adequate out-of-sample performance. In particular, the model successfully captures the dynamic dependencies (as shown by the high p-values for lagged hits).

F.2 Robustness to different priors for the euro area SQVAR's parameters

This section demonstrates the insensitivity of our main results to alternative prior choices, supporting the credibility of the empirical findings as presented in Section 5. To explicitly demonstrate the robustness of our results, we perform two additional estimation exercises, deliberately moving away from the U.S.-data-based informative prior. First, we estimate the euro area SQVAR model's parameters using a standard **Minnesota prior**. The Minnesota prior is a commonly used empirical prior for VAR models that shrinks coefficients towards a random walk, favoring parsimony. Second, we estimate the euro area SQVAR model's parameters using an **uninformative (flat) prior**. This choice reflects maximal prior ignorance, allowing the euro area data alone to dominate the estimation.

The quantile impulse-response functions (QIRFs) resulting from these two alternative prior specifications are presented in Figures F.1 and F.2. For direct comparison, the QIRFs from our primary analysis – using U.S.-data-based informative priors – are included as dotted lines in these figures.

The key findings from our prior robustness checks can be summarized as follows:

First, we observe an overall qualitative consistency, particularly for GDP. As expected, the QIRF estimates show some variation across prior specifications. However, the qualitative patterns remain largely consistent, particularly for real GDP growth (the central variable for assessing macro-prudential downside risks). The direction, persistence, and especially the asymmetric nature of responses (e.g., larger downside shifts in the left tail than in the median or right tail) for GDP to shocks (e.g., to financial stress or the financial cycle) are preserved across all prior set-

tings. This consistency for GDP is paramount, as our macro-prudential analysis critically relies on accurately characterizing the downside risks to economic activity.

Second, we do observe some interesting differences in the response of HICP inflation. When using the Minnesota prior, both the lower and upper tail quantiles of inflation tend to rise in response to a financial stress (CISS) shock. In contrast, our baseline model showed these quantiles moving in opposite directions, suggesting a reduction in inflation uncertainty. Additionally, the response of inflation to an interest rate shock becomes positive under the Minnesota prior. With the uninformative prior, the response of inflation to an interest rate shock is notably more negative than in our baseline model. The differences in inflation responses highlight the sensitivity of certain macroeconomic relationships, particularly inflation dynamics, to prior assumptions when data are limited or when specific relationships are less strongly identified by the likelihood alone. The informative U.S. prior, by bringing in information from a larger sample, helps to pin down these relationships more precisely, reflecting the common inflation dynamics observed in major advanced economies.

Despite the noted differences in inflation responses, we conclude that our primary findings related to macro-prudential policy and downside risks to GDP are robust to the selection of priors. The critical aspect for our analysis, the behavior of GDP growth, especially its left tail and its asymmetric response to financial shocks, remains qualitatively stable across all considered prior specifications. This robustness provides strong confidence in the policy implications derived from our framework.

F.3 Robustness to a different hyperprior parameterization

Our empirical results are robust to adopting plausible alternative priors for the SQVAR model's parameters and hyperparameters. Figure F.3 presents QIRF estimates using a tighter hyperprior on the shrinkage parameter $\lambda_i(\gamma)$. The prior on the 'looseness' parameter, $\lambda_i(\gamma)$, is set to $\underline{\alpha}_\lambda = 3$ and $\zeta_\lambda = 0.5$. The estimation sample is 1990Q1 to 2022Q4. Dotted lines denote the baseline QIRFs.

Figure F.1: Quantile impulse response function estimates with a Minnesota prior and no U.S. shrinkage

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. Credible intervals are at a 95% level. The estimation sample is 1990Q1 to 2022Q4. Dotted lines denote the baseline QIRFs.




Figure F.2: Quantile impulse response function estimates with an uninformative prior and no U.S. shrinkage

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. The estimation sample is 1990Q1 to 2022Q4. Dotted lines denote the baseline QIRFs.

— 0.1 quantile (Posterior mean + 95% credible band + baseline posterior mean) — 0.5 quantile (Posterior mean + 95% credible band + baseline posterior mean)
 — 0.9 quantile (Posterior mean + 95% credible band + baseline posterior mean)





Figure F.3: Quantile impulse response function estimates using a tighter prior on λ

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. The prior on the 'looseness' parameter, λ , are set to $\alpha_\lambda = 3$ and $\zeta_\lambda = 0.5$. The estimation sample is 1990Q1 to 2022Q4. Dotted lines denote the baseline QIRFs.



F.4 Robustness to modeling the data at a monthly frequency

Our empirical results are robust to modeling the macro-financial data at a monthly instead of a quarterly frequency. This is reassuring as our shock identification approach relies predominantly on timing restrictions, which could potentially be affected by the data sampling frequency.

Figure F.4 presents quantile impulse response functions for the monthly euro area B-SQVAR. The QIRFs are quantitatively similar. For this exercise, real GDP growth and the financial cycle indicator (SRI) for the U.S. and the euro area are interpolated from quarterly data, using industrial production to obtain monthly real GDP data. The monthly SQVAR model includes four lags, dummies for each quarter in 2020, and the same identifying timing restrictions as the quarterly model. Growth and inflation are annualized month-on-month log differences. The estimation samples are 1973M2 – 2022M12 for the U.S. and 1997M2 – 2022M12 for the euro area.

The interpolation details are as follows: The quarterly real GDP data – both for the U.S. and the euro area – are interpolated to the monthly frequency using state space methods and monthly industrial production as an interpolator variable; see [Litterman \(1983\)](#). The disaggregated monthly real GDP data add up to the quarterly values of real GDP by construction. Estimation is implemented using the procedure DISAGGREGATE.SRC in WinRATS (see [Doan \(2016\)](#)). The procedure is similar to the approach advocated by [Stock and Watson \(2010\)](#) and recently applied in [Jarocinski and Karadi \(2020\)](#). The quarterly data of the Systemic Risk Indicator (SRI) is interpolated to the monthly frequency using the procedure by [Chow and Lin \(1971\)](#). This procedure regresses the SRI on a constant, and the error is assumed to follow a stationary AR(1) process. Estimation is again based on the procedure by [Doan \(2016\)](#).

While the monthly QIRFs remain quantitatively similar to their quarterly counterparts, the reaction of monthly real GDP to a shock to the three-month OIS rate may deserve a brief comment. In the monthly model, an increase in short-term interest rates appears to have a positive impact on the left tail and a negative on the right tail of the real GDP distribution. This would be consistent with the shock reducing the overall volatility of real GDP. While this finding may be of interest for future work, it does not alter our main conclusions regarding the asymmetric impact of financial

stress or the effectiveness of counter-cyclical macro-prudential policy in mitigating downside risks.

F.5 Robustness to adopting different sample endpoints

Figures [F.5](#) and [F.6](#) present QIRFs based on subsample parameter estimates. Specifically, Figure [F.5](#) is based on an estimation sample from 1990Q1 to 2008Q2, excluding the euro sovereign debt crisis, Covid-19 period, and 2022 energy crisis. Figure [F.6](#) is based on an estimation sample from 1990Q1 to 2019Q4, excluding the Covid-19 period and 2022 energy crisis. The QIRFs remain approximately similar to those based on the full sample, particularly the response of the real GDP growth rate to shocks.

F.6 Robustness to removing the transmission lag

This section studies the effect of removing the transition lag mentioned in Section [5.1](#). Specifically, we impose a restriction on the immediate (first-lag) structural impact of short-term interest rates on HICP inflation and real GDP growth within the $A_1(\gamma)$ matrix. I.e., the direct coefficients from interest rates to inflation and to GDP in the first lag are restricted to zero. This implies that while short-term interest rates (monetary policy) can still affect inflation and real GDP at the first lag, they must do so indirectly via the financial cycle variable, given its ordering as the first endogenous variable in our recursive (triangular) identification scheme for $A_0(\gamma)$.

Our primary motivation for imposing this transmission lag is to prevent the occurrence of the “price puzzle” for the euro area data. The price puzzle, a common empirical anomaly in VAR models of monetary policy, manifests as a counter-intuitive positive response of inflation to a monetary policy tightening (i.e., an increase in short-term interest rates). As highlighted by [Estrella \(2015\)](#), imposing a transmission lag can be an effective and economically sensible way to address this issue. Indeed, our preliminary estimations confirmed that this simple restriction resolves the price puzzle in our dataset for the euro area.

To study the impact of this timing restriction, we can temporarily remove the transmission lag, allowing for an immediate direct impact of interest rates on inflation and real GDP. Figure

F.7 plots the QIRFs from this alternative specification. For ease of comparison, the QIRFs from our baseline model with the transmission lag are shown as dotted lines. The results for most variables and relationships are numerically and qualitatively very similar to our baseline model, underscoring the general robustness of our macro-financial dynamics. Unsurprisingly, the primary exception is the reaction of inflation to an interest rate shock. Without the transmission lag, the price puzzle reappears, with inflation showing a counter-intuitive initial increase in response to a monetary policy tightening. This confirms that the imposition of the transmission lag effectively serves its intended purpose of resolving this specific anomaly.

Figure F.4: Quantile impulse response functions for the monthly euro area QVAR

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. Real GDP growth and the financial cycle indicator for the U.S. and euro area are interpolated from quarterly data, using industrial production for the former. The monthly model includes four lags, dummies for each quarter in 2020 and the same zero restrictions as the quarterly model. Growth and inflation are annualized month-on-month log-differences. The estimation samples are 1973M2 – 2022M12 for the U.S. and 1997M2 – 2022M12 for the euro area. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. Credible intervals are at a 95% level.




Figure F.5: Quantile impulse response function estimates using data until 2008Q2

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. The estimation sample is 1990Q1 to 2008Q2. Dotted lines denote the baseline QIRFs.




Figure F.6: Quantile impulse response function estimates using data until 2019Q4

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. Credible intervals are at a 95% level. The estimation sample is 1990Q1 to 2019Q4. Dotted lines denote the baseline QIRFs.





Figure F.7: Quantile impulse response function estimates with no transmission lag

QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. The estimation sample is 1990Q1 to 2022Q4. Dotted lines denote the baseline QIRFs.

— 0.1 quantile (Posterior mean + 95% credible band + baseline posterior mean) — 0.5 quantile (Posterior mean + 95% credible band + baseline posterior mean)
 — 0.9 quantile (Posterior mean + 95% credible band + baseline posterior mean)



G Additional euro area results

G.1 Euro area posterior estimates

Figure G.1: Posterior inference for the euro area, $\omega(\gamma)$

Posterior mean and 95% credible intervals obtained from 2,500 posterior draws. Least squares estimates (in red) for the conditional mean are provided for comparison.




Figure G.2: Posterior inference for the euro area, $A_0(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.3: Posterior inference for the euro area, $A_1(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.4: Posterior inference for the euro area, $A_2(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.5: Posterior inference for the euro area, $A_3(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.6: Posterior inference for the euro area, $A_4(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.7: Posterior inference for the euro area exogenous variables, $C(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.8: Posterior inference for the euro area dummy variables, $B(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure G.9: Posterior means for the euro area tightness parameter, $\lambda_i(\gamma)$

Posterior mean estimates of $\lambda_i(\gamma)$ are obtained from 2,500 posterior draws.





Figure G.10: Posterior distributions of moments in counterfactual policy experiments

Kernel-smoothed posterior densities based on 10,000 draws from the full sample estimates and using 1991Q1 as the forecast origin. For each draw, we identify all historically-realized conditional quantiles and then perform two counterfactual experiments. One solid and two dashed vertical lines indicate historically-observed moments and counterfactual posterior means, respectively.



G.2 Selected predictive quantiles for RGDP growth rates


Figure G.11 plots two estimates of GaR, the 5% and 10% predictive quantile of real GDP growth, at any time t . Three other quantiles (the median, 90%, and 95%) are reported to indicate upside potential and for comparison. All estimates are based on the same (full-sample) parameter estimates but are otherwise conditional on variables observed up to time t only. The predictive densities correspond to four different forecast horizons: one quarter ahead, one year ahead, one year ahead one year out, and one year ahead four years out.

We highlight two findings. First, as a result of macro-financial interactions, the SQVAR's lower quantiles for future real GDP growth are considerably more volatile than its upper quantiles. For the one-quarter-ahead forecast, the predictive quantiles' standard deviations decrease almost monotonically as the quantiles increase: from 2.66 and 2.59 for the 5% and 10% quantile, to 1.85 for the median quantile, to 1.44 and 1.53 for the 90% and 95% quantile.

Second, the asymmetry in the predictive density decreases with the forecasting horizon. For the one-year-ahead forecast, the predictive quantiles' standard deviations continue to decrease as the quantiles increase: from 1.73 and 1.61 for the 5% and 10% quantile, to 1.27 for the median, to 1.04 and 1.02 for the 90% and 95% quantile. For the one-year-ahead forecast four years out, the predictive quantiles' standard deviations still decrease as the quantiles increase, although the pattern is now more muted: from 0.49 and 0.43 for the 5% and 10% quantile, to 0.31 for the median, to 0.28 and 0.29 for the 90% and 95% quantile. We conclude that the asymmetry in the predictive density is considerably reduced but still noticeable after five years.

Figure G.11: Predictive quantiles for euro area real GDP growth

Five time- t predictive quantiles for euro area real GDP growth over different forecast horizons. The top left and top right panels refer to annualized one quarter and one year-ahead growth, respectively. The bottom left and bottom right panels refer to the one year-ahead growth rate one year out, and one year-ahead growth four years out. All forecasts are obtained by simulation. SQVAR parameters are held constant at their full sample posterior mean estimated for $q = 19$ quantiles between 0.05 and 0.95.



H U.S. posterior estimates

Figure H.1: Posterior inference for the U.S., $\omega(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates (in red) for the conditional mean are provided for comparison.




Figure H.2: Posterior inference for the U.S., $A_0(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates for the conditional mean are provided for comparison.




Figure H.3: Posterior inference for the U.S., $A_1(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are for comparison.




Figure H.4: Posterior inference for the U.S., $A_2(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are for comparison.




Figure H.5: Posterior inference for the U.S., $A_3(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure H.6: Posterior inference for the U.S., $A_4(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure H.7: Posterior inference for the U.S. exogenous variables, $C(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure H.8: Posterior inference for the U.S. dummy variables, $B(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.




Figure H.9: Posterior mean estimates for U.S. prior tightness parameter $\lambda_i(\gamma)$

Posterior mean estimates of $\lambda_i(\gamma)$ for 19 quantiles γ between 0.05 and 0.95. Based on 2,500 draws from the posterior.





Figure H.10: U.S. quantile impulse response function estimates

Quantile impulse response functions implied by the posterior estimates reported in Figures H.1-H.9. QIRF estimates are based on 400 draws from the posterior distribution, and $2 \times 20,000$ simulations per posterior draw to obtain shocked and baseline quantiles of all variables. The shock size is equal to one standard deviation of the shocked variable's median regression residuals. Variables are ordered financial cycle (first row), CPI inflation (second row), real GDP growth (third row), financial stress (US CISS, fourth row), and the Federal funds rate (fifth row). Credible intervals are at a 95% level. The estimation sample is 1973Q1 to 2022Q4.



References

- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review* 109(4), 1263–89.
- Ampudia, M., M. L. Duca, M. Farkas, G. Perez-Quiros, M. Pirovano, G. Rünstler, and E. Tereanu (2021). On the effectiveness of macroprudential policy. *ECB Discussion Paper 2559*, 1–42.
- Araujo, J., M. Patnam, A. Popescu, F. Valencia, and W. Yao (2020). Effects of macroprudential policy: Evidence from over 6,000 estimates. *IMF Working Paper WP/20/67*.
- Carney, M. (2020). The grand unifying theory (and practice) of macroprudential policy. *Speech given at University College London on 5 March 2020*, 1 – 22.
- Cerutti, E., S. Claessens, and L. Laeven (2017). The use and effectiveness of macroprudential policies: New evidence. *Journal of Financial Stability* 28, 203—224.
- Chavleishvili, S. and M. Kremer (2025). CISS of death: Measuring financial crises in real time. *Review of Finance* 29(3), 685 – 710.
- Chernozhukov, V., I. Fernandez-Val, and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica* 78, 1093–1125.
- Chow, G. C. and A.-I. Lin (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *Review of Economics and Statistics* 53(4), 372–275.
- Christoffersen, P. F. (1998). Evaluating Interval Forecasts. *International Economic Review* 39, 841–862.
- Constancio, V. (2016). Principles of macroprudential policy. *Speech at the ECB-IMF conference on Macroprudential Policy, Frankfurt am Main, on 26 April 2016*.
- Danthine, J.-P. (2012). Taming the financial cycle. *Speech at the 30th SUERF Colloquium, Zurich, on 5 September 2012*.
- Doan, T. (2016). Disaggregate: Rats procedure to implement general disaggregation. Statistical Software Components, Boston College.

- Engle, R. F. and S. Manganelli (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics* 22(4), 367–381.
- Estrella, A. (2015). The price puzzle and VAR identification. *Macroeconomic Dynamics* 19(8), 1880–1887.
- Fagan, G., J. Henry, and R. Mestre (2001). An area-wide model (AWM) for the euro area. *ECB Working Paper* 42.
- Fishburn, P. (1977). Mean-risk analysis with risk associated with below-target returns. *American Economic Review* 67, 116–126.
- Forbes, K. J. (2019). Macroprudential policy: what we’ve learned, don’t know, and need to do. *AEA Papers and Proceedings* 109, 470–475.
- Giannone, D., M. Lenza, and G. E. Primiceri (2015). Prior selection for vector autoregression. *The Review of Economics and Statistics* 97(2), 436–451.
- Huber, F. and M. Feldkircher (2019). Adaptive shrinkage in bayesian vector autoregressive models. *Journal of Business & Economic Statistics* 37(1), 27–39.
- Ibrahim, J. G. and M.-H. Chen (2000). Power prior distributions for regression models. *Statistical Science* 15(1), 46–60.
- Ibrahim, J. G., M.-H. Chen, Y. Gwon, and F. Chen (2015). The power prior: Theory and applications. *Statistics in Medicine* 34(28), 3724–3749.
- Ingves, S. (2012). Financial crises and financial regulation – thoughts after five turbulent years. *Speech to the Swedish Economics Association, Stockholm, 13 June 2012*, 1–11.
- Jarocinski, M. and P. Karadi (2020). Deconstructing monetary policy surprises — the role of information shocks. *American Economic Journal: Macroeconomics* 12, 1–43.
- Khare, K. and J. P. Hobert (2012). Geometric ergodicity of the Gibbs sampler for Bayesian quantile regression. *Journal of Multivariate Analysis* 112, 108–116.

- Kilian, L. and S. Manganelli (2008). The central banker as a risk manager: Estimating the Federal Reserve's preferences under Greenspan. *Journal of Money, Credit, and Banking* 40, 1103–1129.
- Koenker, R. (2005). *Quantile Regression*. Cambridge: Cambridge University Press.
- Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica* 46(1), 33–50.
- Korobilis, D. (2017). Quantile regression forecasts of inflation under model uncertainty. *International Journal of Forecasting* 23, 11 – 20.
- Korobilis, D. and D. Pettenuzzo (2019). Adaptive hierarchical priors for high-dimensional vector autoregressions. *Journal of Econometrics* 212(1), 241–271.
- Kozumi, H. and G. Kobayashi (2011). Gibbs sampling methods for Bayesian quantile regression. *Journal of Statistical Computation and Simulation* 81(11), 1565–1578.
- Kupiec, P. H. (2000). Stress tests and risk capital. *Journal of Risk* 2, 27–40.
- Lang, J. H., C. Izzo, S. Fahr, and J. Ruzicka (2019). Anticipating the bust: a new cyclical systemic risk indicator to assess the likelihood and severity of financial crises. *ECB Occasional Paper No. 219*, 1 – 77.
- Litterman, D. (1983). A random walk, markov model for the distribution of time series. *Journal of Business and Economic Statistics* 1, 169–173.
- Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions - five years of experience. *Journal of Business & Economic Statistics* 4, 25–38.
- Lütkepohl, H. (2005). *New introduction to multiple time series analysis*. Springer, Germany.
- Richter, B., M. Schularick, and I. Shim (2019). The costs of macroprudential policy. *Journal of International Economics* 118, 263–282.
- Stock, J. H. and M. W. Watson (2010). Research memorandum: Distribution of quarterly values of GDP/GDI across months within the quarter. *Mimeo*.
- Suarez, J. (2022). Growth-at-risk and macroprudential policy design. *Journal of Financial Stability* 60, 101008.

- Surprenant, S. (2025). Quantile VARs and macroeconomic risk forecasting. *Bank of Canada Working Paper 2025-4*, 1–44.
- Tarullo, D. K. (2011). Regulating systemic risk. *Speech at the 2011 Credit Markets Symposium, Charlotte, North Carolina, on 31 March 2011*.
- Yellen, J. L. (2010). Macroprudential supervision and monetary policy in the post-crisis world. *Speech at the Annual Meeting of the National Association for Business Economics, Denver, 11 October 2010*, 1–11.
- Yu, K. and R. A. Moyeed (2001). Bayesian quantile regression. *Statistics & Probability Letters 54*, 437–447.