

The risk management approach to macro-prudential policy*

*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

August 21, 2025

Abstract

Macro-prudential policymakers are tasked with assessing medium-term downside risks to the real economy stemming from financial imbalances and the associated threat of a future financial crisis. Before implementing policy measures to manage financial vulnerabilities, however, these policymakers must also weigh the policies' cost in terms of the economy's upside potential. This setting gives rise to a trade-off between limiting downside risks and preserving upside potential. We formalize the decision problem under uncertainty by stipulating an explicit loss function for macro-prudential policy, drawing inspiration from key policymakers' speeches. To evaluate policy, flexible and potentially asymmetric predictive distributions are obtained from a novel Bayesian structural quantile vector autoregressive model. We study the implied quantile impulse responses, assess the potential of macro-prudential policy for reducing growth volatility and negative skewness, infer the macro-prudential policy stance, and study the circumstances under which macro-prudential interventions are likely to be beneficial.

Keywords: Growth-at-risk, structural quantile vector autoregression, Bayesian inference, financial conditions, macro-prudential policy.

JEL classification: C32, C54.

*E-mail contacts: sulkhan.chavleishvili@econ.au.dk; rengle@stern.nyu.edu; stephan.fahr@ecb.europa.eu; manfred.kremer@ecb.europa.eu; felt@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.

1 Introduction

The Global Financial Crisis (GFC) of 2008–2009 underscored the critical need for robust macro-prudential policies; see e.g. the reports by [Brunnermeier et al. \(2009\)](#) and [de Larosiere \(2009\)](#). In particular, the severity of the GFC exposed a pronounced underestimation of medium-term downside risks to the economy resulting from a gradual accumulation of financial vulnerabilities over time. Since 2009, a plethora of theoretical and empirical frameworks have been developed to assess and manage downside risks, including, for example, [Brunnermeier and Sannikov \(2014\)](#), [Boissay et al. \(2016\)](#), [He and Krishnamurthy \(2019\)](#), [Adrian et al. \(2019\)](#), [Caldara et al. \(2021\)](#), [Adrian et al. \(2022\)](#), [Suarez \(2022\)](#), and [Carriero et al. \(2023\)](#).

Yet, effective macro-prudential policies cannot focus solely on mitigating downside risks, but must balance risk mitigation with growth considerations. As former UK Chancellor George [Osborne \(2012\)](#) put it, “*we do not want the financial stability of a graveyard.*” How to formalize this key trade-off as an economic decision problem under uncertainty, and how to study it empirically for an advanced market-based economy, however, is currently unclear. This is despite policymakers’ urgent need for analytical support and a formal decision framework’s theoretical appeal.

This paper develops a comprehensive framework to formalize the decision problem faced by the macro-prudential policymaker, and applies it to euro area data. To our knowledge, we are the first to extend the “risk management” approach of [Greenspan \(2003, p. 3\)](#), [Cecchetti \(2006\)](#), and [Kilian and Manganelli \(2008\)](#) to consider macro-prudential decisions. Thinking quantitatively about macro-prudential policy requires a loss (or utility) function, and, to our knowledge, we are the first to stipulate an explicit loss function for macro-prudential policy, drawing inspiration from key policymakers’ speeches. To evaluate policy, we propose a novel fully-tractable structural quantile vector autoregressive (SQVAR) model, and show how to estimate its parameters equation-by-equation using a novel four-step Gibbs sampler.¹ Although several “macro-at-risk” stylized

¹While SQVAR models are not new to the literature (e.g., [Chavleishvili and Manganelli, 2024](#)), the SQVAR model in this paper is novel in considering a more general setting (modeling more than two endogenous variables at more than one lag), allowing for exogenous variables and additional deterministic terms, and proposing a particularly tractable Bayesian inference methodology for its statistical analysis.

facts have by now been pointed out in existing work, our framework is unique in that virtually *all of them* can be captured *simultaneously* and studied in one coherent, single-step inferential framework. Finally, whenever possible, we relate our empirical results to models of financial frictions or other macroeconomic models.

Current measures of economic downside risk are borrowed from the financial risk management literature. For example, Growth-at-Risk (GaR) is Value-at-Risk (VaR) applied to the gross domestic product (GDP) instead of a bank's investment portfolio (e.g., [Prasad et al., 2019](#), [Adrian et al., 2022](#), and [Lang et al., 2023](#)). GaR is the tail quantile (usually 5%) of the variable of interest. The main advantage of GaR is its simplicity: by definition, GDP growth will not be below the estimated GaR with 95% probability. However, measuring downside risk by a single quantile has been criticized in the risk management literature on the grounds that it does not take the whole tail of the distribution into account, as well as for not being a coherent (e.g. sub-additive) measure of risk. In addition, it captures only one part (the downside) of the macro-prudential decision problem, and is silent on the upside potential for the economy.

We address the limitations of GaR by introducing *growth shortfall* (GS) as a measure of downside risk, and by complementing it with *growth longrise* (GL) to gauge the economy's upside potential. GS (GL) equals expected growth conditional on growth realizing a value below (above) a certain threshold, multiplied by the probability of the conditioning event. Zero serves as a natural threshold when applied to GDP growth, distinguishing economic contractions from expansions. By design, the sum of growth shortfall and longrise equals expected growth.² If the macro-prudential policymaker assigns equal importance to both components, its goal would be to maximize expected growth. However, a study of policymakers' speeches (see below) strongly suggests that their objective is, typically, to prevent severe and painful economic contractions without unnecessarily limiting growth potential. The decision framework proposed in this paper formalizes this trade-off by assigning greater weight to GS than to GL.

We illustrate our decision framework by applying it to the euro area economy. The empirical

²GS and GL do *not* coincide with expected shortfall and expected longrise as defined in [Adrian et al., 2019](#), p. 1277; see also Section 2.1 below.

implementation necessitates estimating the joint predictive distribution of all macroeconomic variables of interest.³ We recover flexible, and potentially asymmetric predictive distributions using a SQVAR model. SQVAR extends the concept of structural VAR to quantile modeling. Essentially, while VARs explain the *mean* dynamic interaction of the endogenous variables of a system, QVARs explain the dynamic interaction of any *quantile* of the variables of interest. Structural identification can be obtained, for example, through short-run timing restrictions. Once the model is estimated for all quantiles, estimates of all predictive distributions can be obtained by simulation; the main text and web appendix provide the necessary details.

To our knowledge, we are the first to estimate an SQVAR model’s parameters equation-by-equation using Bayesian methods. We rely predominantly on established methodology (Yu and Moyeed, 2001, Kozumi and Kobayashi, 2011, Khare and Hobert, 2012), but also make an econometric contribution. Leveraging the Bayesian approach, our estimates of the euro area parameters depend on an informative prior density to sharpen inference, which we obtain from a different but related sample (several decades of United States (U.S.) data). That is, we first obtain posterior estimates from U.S. data, which then serve as informative priors for the euro area parameters. The prior variance, which represents the “weight” assigned to the prior density, can be estimated from the euro area data following the Empirical Bayes approach of Giannone et al. (2015). Estimating the prior variance extends the three-step Gibbs sampler of Khare and Hobert (2012) to a novel four-step sampler. Overall, our approach to inference allows us to obtain precise posterior estimates of measures of downside risk and upside potential, among other nonlinear functions of the model’s deterministic parameters. In addition, we obtain appropriate finite-sample credible intervals, despite a relative paucity of euro area quarterly macro data and a considerable number of model parameters to be estimated.

³Because of the asymmetric emphasis on downside risk, expected growth is no longer a sufficient statistic to solve the macro-prudential risk management problem. In principle, the mean could still be a sufficient statistic if the random variables of interest were symmetric and with time-invariant higher moments. In such a case, estimates of the conditional means of the random variables of interest, augmented by suitable assumptions of the distribution of the residuals, would suffice. However, there is by now substantial empirical evidence that financial variables exert a time-varying and asymmetric impact on real variables (see, inter alia, Adrian et al., 2019, Adrian et al., 2022, Delle Monache et al., 2023, Lang et al., 2023, and Iseringhausen et al., 2024). This necessitates the use of suitable econometric models that account for these characteristics.

Our preferred SQVAR model includes a measure of the financial cycle, real GDP growth, HICP inflation, financial stress, and a short-term risk-free interest rate and uses four lags. Its parameters are estimated from euro area data between 1990Q1 and 2022Q4. A global commodity price index is included as an additional exogenous variable. The financial cycle is a measure of credit dynamics in the economy and represents an intermediate target variable on which macro-prudential policymakers can act by activating or adjusting macro-prudential instruments. Financial stress characterizes financial crises and can amplify other shocks (e.g., [He and Krishnamurthy, 2019](#)). Overall, our choice of variables reflects the idea that financial stability is of concern to policymakers if it is triggered by an impairment of the financial system and has real economic consequences, e.g. in terms of future economic activity.⁴ The monetary part of the system (with inflation and interest rate) is included to take the impact of the central bank’s actions into account.

We highlight four results. First, we argue that macro-prudential policymakers’ speeches are sufficiently clear about the ends and means of macro-prudential policy for us to derive an explicit loss (utility) function from them. An explicit loss function is a prerequisite to thinking quantitatively about optimal macro-prudential policy. For example, [Danthine \(2012\)](#), [Constancio \(2016\)](#), and [Carney \(2020\)](#) all agree that the main aim of macro-prudential policy is to make sure the financial system supports the economy; this requires that the financial system is strong enough to continue lending to households and businesses when economic shocks occur. All policymakers stress a counter-cyclical dimension of macro-prudential policy, aimed at smoothing (or “taming”) the financial cycle, and agree that applying macro-prudential instruments in a counter-cyclical way is crucial. The speech by [Carney \(2020\)](#) comes closest to proposing an explicit objective function. We build on his formulation and show that its arguments are related to GS and GL.

Second, our estimated SQVAR is characterized by substantial asymmetries. Our quantile impulse response function estimates in particular suggest that the system’s dynamic properties differ

⁴The ECB definition of financial stability refers to “*the risk that the provision of necessary financial products and services by the financial system will be impaired to a point where economic growth and welfare may be materially affected*,” see [ECB \(2019\)](#). Similarly, the Financial Stability Board, International Monetary Fund, and the Bank for International Settlements define systemic risk as a “*risk of disruption to financial services that is (i) caused by an impairment of all or parts of the financial system, and (ii) has the potential to have serious negative consequences for the real economy*,” see [FSB \(2009\)](#).

considerably across quantiles. For example, a shock to financial stress shifts the left tail of euro area future real GDP growth further to the left, while leaving its conditional median and right tail approximately unaffected. In addition, a shock to financial stress deflates the financial cycle in the medium term (approximately similarly across quantiles), and impacts the lower and upper quantiles of inflation with different signs, resulting in less inflation uncertainty. Vice versa, financial stress responds to a shock to the financial cycle by initially decreasing for six quarters and then increasing; this pattern is considerably more pronounced for stress’s upper tail (the 0.9 quantile) than its lower tail (the 0.1 quantile). In sum, the estimated macro-financial interactions imply that the upper quantiles of the predictive GDP growth distribution are less volatile than its lower quantiles (cf. [Adrian et al., 2022](#), [Lang et al., 2023](#)), and lend empirical support to macroeconomic models that allow for asymmetric impacts of financial variables on macroeconomic outcomes. This includes, for example, nonlinear macro-financial models with occasionally binding financing constraints such as [He and Krishnamurthy \(2019\)](#), [Van der Gucht \(2021\)](#), and [Mendicino et al. \(2024\)](#). Our estimates lend less support to macro DSGE models that are solved to first order (log-linearized) and subsequently equipped with Gaussian shocks.

Third, counterfactual simulations suggest that an active macro-prudential policy regime, targeted at smoothing the financial cycle, could considerably improve economic welfare by increasing economic trend (mean) real GDP growth while lowering growth volatility, negative growth skewness, and its excess kurtosis. This is achieved mostly by decreasing the probability of observing periods of high financial stress.⁵ Interestingly, managing the financial cycle also reduces inflation volatility, its positive skewness, and (mild) kurtosis. This, in turn, allows short-term (monetary policy) interest rates to be less volatile.

Fourth, our framework provides a metric to assess whether the macro-prudential stance is too tight or too loose conditional on available information. Welfare calculations for leaning against an exuberant financial cycle are based on the objective function as derived above. Flexible predictive

⁵To compare macro-economic outcomes across counterfactual settings we need to assume that the presence of a new policy regime does not change that SQVAR’s parameters to such an extent that it makes the thought experiment uninteresting. We discuss the Lucas’ critique in the main text.

distributions are obtained from our SQVAR model as needed. We find that the welfare gains from macro-prudential tightening can be positive or negative. Historically, the gains from macro-prudential policy tightening tended to be positive when the financial cycle takes high values (i.e. when credit growth and house price appreciation are high), short-term interest rates and inflation are low, the economy is growing, and financial stress is low. Our results thus lend support to the view that macro-prudential policymakers should act in a counter-cyclical fashion by easing policy in bad times (to increase loss-absorbing capacity) and tightening policy in good times (to lean against exuberance; see also [Van der Ghote, 2021](#)).

Our work is part of a rapidly growing body of research that studies downside risk in macroeconomic outcomes. Most of this work has focused on the risk of considerable declines in real GDP, brought about by a deterioration of financial conditions; see e.g. [Adrian et al. \(2019\)](#), [Prasad et al. \(2019\)](#), and [Caldara et al. \(2021\)](#). The International Monetary Fund (IMF), the European Central Bank (ECB), and the Federal Reserve Bank of New York now routinely publish GaR estimates for major world economies; see [IMF \(2017\)](#), [ECB \(2019\)](#), and [New York Fed \(2022\)](#). These developments have motivated a proliferation of modeling frameworks to assess the severity of extreme events associated with key economic variables, including single-equation QR models ([Adrian et al., 2019](#)), panel QR models ([Brandao-Marques et al., 2020](#), [Adrian et al., 2022](#), [Lang et al., 2023](#)), panel-GARCH models ([Brownlees and Souza, 2021](#)), fully non-parametric kernel regression models ([Adrian et al., 2021](#)), combined linear vector autoregressive (VAR) and single-equation QR models ([Duprey and Ueberfeldt, 2020](#), [Forni et al., 2021](#)), nonlinear Bayesian VAR models ([Caldara et al., 2021](#), [Carriero et al., 2023](#)), and quantile FAVAR models ([Korobilis and Schröder, 2024](#)). [Plagborg-Møller et al. \(2020\)](#) provide a critical review of this literature.

We proceed as follows. Section 2 presents our decision framework. Section 3 introduces the statistical model. Section 4 describes the euro area data. Section 5 discusses our key empirical results. Section 6 concludes. A web appendix provides further technical and empirical results.

2 The risk management decision framework

This section first defines measures of downside risk and of upside potential. Some of these measures are related to well-known concepts from the financial risk management literature. We then derive an explicit objective function from policymaker's speeches, and integrate these parts into an encompassing decision framework.

2.1 Measures of downside risk and upside potential

Growth-at-risk: Our first measure of adverse impact is *growth-at-risk* ($\text{GaR}_{t,t+h}^\gamma$) at confidence level $\gamma \in (0, 1)$, defined implicitly by the probability

$$\mathbb{P} [y_{t+h} \leq \text{GaR}_{t,t+h}^\gamma | \Omega_{yt}] = \gamma, \quad (1)$$

where y_t is the quarterly annualized real GDP growth rate between time $t - 1$ and t , $h = 1, \dots, H$, and H is a certain prediction horizon. The information set Ω_{yt} contains all data known at time t ; see Section 3 for details. In words, $\text{GaR}_{t,t+h}^\gamma$ is implicitly defined by the time t probability of quarterly annualized output growth at $t + h$ falling below $\text{GaR}_{t,t+h}^\gamma$, which by definition is set equal to γ (see, for example, McNeil et al., 2005, Ch. 2.2).

Growth shortfall: Our second measure of adverse real economic impact is *growth shortfall* (GS), defined as

$$\begin{aligned} \text{GS}_{t,t+h}^\tau &= \int_{-\infty}^{\tau} y_{t+h} dF_{t,t+h}(y_{t+h}) \\ &= \mathbb{E} [y_{t+h} | y_{t+h} < \tau, \Omega_{yt}] \times \mathbb{P} [y_{t+h} < \tau | \Omega_{yt}], \end{aligned} \quad (2)$$

where $F_{t,t+h}$ is a time- t conditional cdf, and $\mathbb{E} [\cdot | \Omega_{yt}]$ denotes a time- t conditional expectation. In principle, the threshold $\tau \in \mathbb{R}$ could be set to a low conditional quantile, say $\tau = \text{GaR}_{t,t+h}^\gamma$. If so, then the first factor in (2) would coincide with the familiar notion of expected shortfall; see e.g. McNeil et al. (2005, Ch. 2) and Adrian et al., 2019, p. 1277. Alternatively, τ could be set to a

certain unconditional quantile, or to zero. In this case, the first factor in (2) does *not* coincide with expected shortfall from the financial risk management literature.⁶ The second factor in (2) is the conditional (on Ω_{yt}) probability of experiencing $y_{t+h} < \tau$.

We set $\tau = 0$ in Section 5 below. If $\tau = 0$, GS can be factored into two intuitive terms: the expected loss conditional on a contraction, and the probability of experiencing a contraction.⁷ While both components could, in principle, be of interest in their own right, GS summarizes them tractably into one metric. When $\tau = 0$, GS corresponds to the economic question: what is the time t -expected contraction of the economy at time $t + h$.

All the above risk measures are economically intuitive and straightforward to communicate. GS (2), however, has theoretical and practical advantages over GaR (1). First, while both risk measures can account for the asymmetric impact of financial variables on the economy, only (2) takes the entire left tail into account. Second, (scaled) expected shortfall is a coherent risk measure, while any single conditional quantile in isolation is not (Artzner et al., 1999). For example, GS contributions are sub-additive, while GaR contributions are not. This feature is desirable if one, for instance, sought to study sector contributions to aggregate downside risk.

Growth longrise: When considering financial stability policies aimed at containing downside risk, the upper quantiles of future GDP growth, should, ideally, not be negatively affected. For setting up the decision framework in Section 2.2, we consider a measure of upside potential that complements GS. We define the *growth longrise* (GL) as

$$\begin{aligned} \text{GL}_{t,t+h}^\tau &= \int_\tau^\infty y_{t+h} dF_{t,t+h}(y_{t+h}) \\ &= \mathbb{E}[y_{t+h} | y_{t+h} > \tau, \Omega_{yt}] \times \mathbb{P}[y_{t+h} > \tau | \Omega_{yt}]. \end{aligned} \quad (3)$$

If $\tau = 0$, then (3) corresponds to the question: what is the time- t expected expansion of the economy between $t + h - 1$ and $t + h$? Similarly to GS, GL combines the expected growth

⁶GS also does not coincide with “downside entropy” in Adrian et al., 2019, p. 1276, which measures the extent to which time- t ’s predictive density differs from its unconditional counterpart below the predictive density’s median.

⁷To see this, note that $\mathbb{E}[y_{t+h} | y_{t+h} < \tau, \Omega_{yt}] \equiv \frac{\int_{-\infty}^\tau y_{t+h} \cdot 1\{y_{t+h} < \tau\} dF_{t,t+h}(y_{t+h})}{\int_{-\infty}^\tau 1\{y_{t+h} < \tau\} dF_{t,t+h}(y_{t+h})} = \frac{\int_{-\infty}^\tau y_{t+h} dF_{t,t+h}(y_{t+h})}{\mathbb{P}[y_{t+h} < \tau | \Omega_{yt}]}$.

given an expansion with the conditional probability of experiencing an expansion.⁸ Given the complementarity between GS and GL, their sum equals the expected growth rate of the economy between $t + h - 1$ and $t + h$,

$$\mathbb{E}[y_{t+h}|\Omega_{yt}] = \int_{-\infty}^{\tau} y_{t+h} dF_{t,t+h}(y_{t+h}) + \int_{\tau}^{\infty} y_{t+h} dF_{t,t+h}(y_{t+h}) = \text{GS}_{t,t+h}^{\tau} + \text{GL}_{t,t+h}^{\tau}. \quad (4)$$

2.2 Deriving an objective function from policymakers' speeches

This section derives an explicit loss (or utility) function from macro-prudential policymakers' speeches. An explicit loss function is a prerequisite to thinking quantitatively about optimal macro-prudential policy. We focus on three speeches by European policymakers: [Danthine \(2012\)](#), then Vice Chairman of the Governing Board of the Swiss National Bank, [Constancio \(2016\)](#), then Vice-President of the ECB, and [Carney \(2020\)](#), then Governor of the Bank of England.

All three policymakers stress that the *raison d'être* of macro-prudential policy is to make sure the financial system supports the economy (e.g., [Carney, 2020](#), p. 7). This requires the financial system to be strong enough to continue lending to households and businesses when economic shocks occur. In particular, macroeconomic downturns should not be amplified by unsustainable debt burdens or funding mismatches. To accomplish these goals, macro-prudential policymakers concentrate on systemic risks, i.e., those large enough to materially impact economic growth.

In addition to addressing structural (non-cyclical) systemic risks, the three policymakers stress a counter-cyclical dimension of macro-prudential policy, aimed at smoothing (“influencing”; or “taming” in [Danthine, 2012](#)) the financial cycle. Macro-prudential policy should be “preemptive and ... counter-cyclical” ([Constancio, 2016](#)), while being cognizant of the trade-offs between supporting economic growth and mitigating downside risks.

The speech by [Carney \(2020, p. 11\)](#) comes closest to proposing an explicit objective function,

⁸To our knowledge, the term “longrise” was coined by [Adrian et al. \(2019\)](#) as the antonym to “shortfall.” GL does not coincide with their “expected longrise” nor “upside entropy.”

stated as

$$\min_{m_t} \mathcal{L} \equiv \sum_{h=0}^H \delta^h \left[\mathbb{E}_t[f(\text{GaR}_{t+h})] - \tilde{\lambda} \mathbb{E}_t[y_{t+h}] \right], \quad (5)$$

where m_t is a “policy setting” at time t , H is the policymaker’s time horizon (which should be thought of as representing “the medium-to-long run”), $\delta \leq 1$ is a discount factor, \mathbb{E}_t is a conditional (on current information) expectation, f is a possibly-convex but otherwise unspecified function, GaR_{t+h} is a measure of downside risk, indicating how much GDP could fall as a consequence of current financial vulnerabilities, $\tilde{\lambda} > 0$ is a preference parameter, and $\mathbb{E}_t[y_{t+h}]$ is the expected growth rate of the economy.

While (5) goes a long way towards putting forward an explicit objective function, it is not entirely explicit about the policymaker’s horizon H , which particular quantile is indicated by GaR_{t+h} , and the precise shape of f . Starting from (5), we obtain a closely-related objective function in two steps. First, we identify its first term, $\mathbb{E}_t[f(\text{GaR}_{t+h})]$, with our expectation-based measure of downside risk, $\text{GS}_{t+h} = \mathbb{E}_t[(y_{t+h} | y_{t+h} < \tau) \times \mathbb{1}\{y_{t+h} < \tau\}]$; see (2). This allows us consider a continuum of lower quantiles rather than tracking a particular one while ignoring others, and take the entire left tail into account. Both factors comprising GS implicitly depend on GaR .⁹ Second, we reformulate (5) as a utility (=negative loss) maximization problem and move the penalty term to linearly scale downside risk. This yields

$$\max_{m_t} \mathcal{U} \equiv \sum_{h=1}^H \delta^h \left[(\lambda^p - 1) \text{GS}_{t,t+h}(y_{t+h}(c_t(m_t), \dots)) + \mathbb{E}_t(y_{t+h}(c_t(m_t), \dots)) \right], \quad (6)$$

where $\lambda^p > 1$ is a weight determining the policymaker’s aversion to negative realizations of output growth, and where the dependence of y_{t+h} on the financial cycle c_t (via m_t) and other contemporaneous and future variables is made explicit.

Since $\text{GS}_{t,t+h}^\tau$ and $\text{GL}_{t,t+h}^\tau$ add to expected growth, see (4), the objective function (6) can be

⁹Recall that $\text{GS}_{t+h} \leq 0$ is bounded from above (assuming $\tau \leq 0$). GS_{t+h} increases non-linearly with the conditional mean of y_{t+h} , owing to its two factors.

rewritten in terms of GS and GL. The utility maximization problem then becomes

$$\max_{m_t} \mathcal{U} \equiv \sum_{h=1}^H \delta^h \left[GL_{t,t+h}(y_{t+h}(c_t(m_t), \dots)) + \lambda^p GS_{t,t+h}(y_{t+h}(c_t(m_t), \dots)) \right], \quad (7)$$

where the penalty parameter $\lambda^p > 1$ is the same as in (6). The expression (7) makes it obvious that, although weighted unequally, *all* predictive quantiles of y_{t+h} are taken into account when evaluating the likely impact of macro-prudential interventions.

Several points regarding the policymaker’s objective function (6) / (7) deserve comment. First, the objective function (6) trades off supporting expected growth (a good) against mitigating downside risks (also a good). While linear-quadratic specifications are commonly assumed for monetary policy objective functions (Carney, 2020), they are less natural for macro-prudential policy because there are no obvious “target” values to aim at. The objective function (6) therefore does not take this familiar form. (This said, the closer GS is to zero, the better.)

Second, when proposing (6), we assumed that the macro-prudential policymaker has a macro-prudential instrument m_t , or vector of instruments, that she can use to mitigate medium-term downside risks to the economy by influencing the financial cycle $c_t(m_t)$. The influence of the financial cycle on the economy’s predictive growth distribution can be direct ($c_t(m_t) \rightarrow y_t$), or indirect via its impact on other contemporaneous and future variables ($c_t(m_t) \rightarrow c_{t+1}, \dots, c_{t+h}, y_t, \dots, y_{t+h-1}, s_t, \dots, s_{t+h-1}, \dots \rightarrow y_{t+h}$), including financial conditions (stress) and short-term interest rates. Our empirical model allows us to capture both direct and indirect forms of transmission; see Section 3 below. Web Appendix A.1 provides a brief and highly selective review of the literature on how capital-, liquidity-, and borrower-based macro-prudential instruments m_t can influence the financial cycle $c_t(m_t)$.

Finally, a number of further points regarding (6) and (7) deserve comment. For space considerations, Web Appendix A.2 continues the present discussion. We emphasize that, while the specific mathematical form of our loss function builds most closely on Carney (2020), the underlying principles and objectives motivating this form are widely shared across policymakers.

3 Structural quantile vector autoregression

3.1 The statistical model

This section describes the SQVAR model used to obtain flexible and potentially asymmetric predictive distributions of all endogenous variables.

We consider a series of random variables $\{x_t : t = 1, \dots, T\}$, where $x_t \in \mathbb{R}^n$ is an n -vector with i th element denoted by x_{it} for $i = 1, \dots, n$ and $n \in \mathbb{N}$. Before discussing the model, we need to define a sequence of information sets. This allows us to work with the stratified modeling strategy as suggested by [Wei \(2008\)](#) and adapted in [Chavleishvili and Manganelli \(2024\)](#). The information sets are defined recursively as

$$\begin{aligned}\Omega_{1t} &\equiv \{x_{t-1}, x_{t-2}, \dots\} \\ \Omega_{it} &\equiv \{x_{i-1,t}, \Omega_{i-1,t}\} \quad i = 2, \dots, n,\end{aligned}$$

adding observations a scalar at a time, nT times, starting from $i = t = 1$. As an example, Ω_{3t} , say, contains all lagged values x_{t-1}, \dots, x_1 as well as the contemporaneous x_{1t} and x_{2t} .

Vector x_t follows a SQVAR(1) process if the γ_i quantile of x_{it} can be written as

$$\begin{aligned}Q_{\gamma_1}(x_{1t}|\Omega_{1t}) &= \omega_1(\gamma_1) + a_{11}(\gamma_1)x_{1,t-1} + a_{12}(\gamma_1)x_{2,t-1} + \dots + a_{1n}(\gamma_1)x_{n,t-1} \\ Q_{\gamma_2}(x_{2t}|\Omega_{2t}) &= \omega_2(\gamma_2) + a_{021}(\gamma_2)x_{1t} + \\ &\quad + a_{21}(\gamma_2)x_{1,t-1} + a_{22}(\gamma_2)x_{2,t-1} + \dots + a_{2n}(\gamma_2)x_{n,t-1} \\ &\quad \vdots \\ Q_{\gamma_n}(x_{nt}|\Omega_{nt}) &= \omega_n(\gamma_n) + a_{0n1}(\gamma_n)x_{1t} + \dots + a_{0n,n-1}(\gamma_n)x_{n-1,t} + \\ &\quad + a_{n1}(\gamma_n)x_{1,t-1} + a_{n2}(\gamma_n)x_{2,t-1} + \dots + a_{nn}(\gamma_n)x_{n,t-1},\end{aligned}\tag{8}$$

for any $\gamma_i \in (0, 1)$, $i \in \{1, \dots, n\}$. When $n = 1$, this simplifies to the quantile autoregressive

process of [Koenker and Xiao \(2006\)](#). In compact notation, we write an SQVAR(p) process as

$$Q_\gamma(x_t|\Omega_t) = \omega(\gamma) + A_0(\gamma)x_t + \sum_{j=1}^p A_j(\gamma)x_{t-j},$$

where $\gamma \equiv [\gamma_1, \dots, \gamma_n]'$, $Q_\gamma(x_t|\Omega_t)$ is shorthand notation for the left hand-side of system (8) that associates each element of x_t with its corresponding information set, and $\omega(\gamma)$ and A_j , $j = 0, \dots, p$ stack their respective terms. $A_0(\gamma)$ is a lower-triangular $n \times n$ matrix, with zeros on the main diagonal. In the context of a standard linear VAR, this representation is equivalent to identifying the system by assuming a Cholesky decomposition for the covariance matrix of the residuals. In the present, more general context, the lower-triangular $A_0(\gamma)$ allows for orthogonal structural shocks that have distinct, quantile-specific impacts; see also 3.2 below. The remaining $n \times n$ coefficient matrices $A_j(\gamma)$ for $j = 1, \dots, p$ can be restricted or unrestricted.

In addition to endogenous variables x_t , empirically-appropriate SQVAR model specifications can also include further deterministic terms d_t and exogenous variables z_t . This yields the model's general form as

$$Q_\gamma(x_t|\Omega_t) = \omega(\gamma) + A_0(\gamma)x_t + \sum_{j=1}^p A_j(\gamma)x_{t-j} + B(\gamma)d_t + \sum_{j=0}^p C_j(\gamma)z_{t-j}. \quad (9)$$

where $B(\gamma)$ and $C_j(\gamma)$ are conformable matrices. Web Appendix B provides additional detail on (9), discusses when model-implied predictive quantiles can cross (and when not), and presents a simple bivariate example.

3.2 Shock identification

Structural identification in SQVAR models is invariably linked to the definition of a “multivariate quantile” and “structural shock” in a non-linear, state-dependent setting. In a standard linear VAR, the contemporaneous impact matrix A_0 captures the instantaneous causal relationships among variables, and a Cholesky decomposition implies a specific recursive ordering for identifying orthogo-

nal structural shocks. In an SQVAR, these contemporaneous relationships are themselves allowed to be quantile-dependent. To disentangle truly orthogonal structural shocks that have distinct, quantile-specific impacts, some form of recursive (lower-triangular) specification of $A_0(\gamma)$ is often a natural and necessary starting point. This ensures that the contemporaneous structural shocks, which drive movements in specific conditional quantiles, are orthogonal to each other within that quantile. In the absence of a universally accepted definition for a “multivariate quantile” that inherently yields a non-triangular, fully identified contemporaneous impact matrix for all quantiles, the recursive approach offers a clear and implementable method for defining such structural shocks.

Matrix $A_0(\gamma)$ in (9) is lower-triangular with zeros on the main diagonal. This is sufficient to exactly identify the model. Identification can be sharpened by considering additional short-run, long-run, and/or sign restrictions. We do so below when we impose two additional short-run restrictions on $A_1(\gamma)$; see Sections 5.1 and 5.2 for details.

While long-run and sign restrictions are a powerful tool in linear VARs (e.g., [Antolín-Díaz and Rubio-Ramírez, 2024](#)), applying them to SQVARs is considerably more intricate. In an SQVAR, the impulse responses, and thus the signs of the effects, depend on the specific quantile being considered. Take a positive “supply shock,” for instance: it might boost median output but also increase overall volatility. This could translate to a negative impact on the left tail of the output distribution (worsening extreme downturns) while still positively affecting the right tail (boosting extreme upturns). This quantile-dependent behavior means that a simple set of sign restrictions — typically imposed on the median or mean response — might not hold true across all relevant quantiles. Instead, researchers would need a richer, quantile-specific set of sign restrictions.¹⁰ Such an approach would be challenging to justify in a macro-prudential setting, and potentially difficult to implement consistently across the entire distribution.

Finally, if an appropriate instrument for a specific shock of interest were available, then that

¹⁰One possibility suggested by [Chavleishvili and Manganelli \(2024\)](#) aims to mitigate some of these complexities. Their approach involves applying common VAR identification strategies, including sign and long-run restrictions, to the *conditional expectations* implied by the QVAR model, which can be recovered e.g. by simulation. This approach implicitly assumes that the chosen identification scheme, when applied to the mean dynamics, remains valid for the entire conditional distribution.

instrument could potentially be used to identify the shock’s impact on all variables across their conditional quantiles. Such an instrument could, for example, be added to a recursively-identified SQVAR model as a first variable in the system (as an “internal instrument”; see [Plagborg-Møller and Wolf, 2021](#), p. 957). To our knowledge, however, the formal development and robust implementation of such an approach is still an area of active research, and theoretical guarantees do not yet exist. We do not use instrumental variables below for these reasons.

3.3 Parameter estimation

We estimate the SQVAR parameters using Bayesian methods. Building upon already-established methodology (particularly [Kozumi and Kobayashi, 2011](#) and [Khare and Hobert, 2012](#)) is possible since the lower-triangular structure on A_0 allows us to estimate the model’s parameters equation-by-equation over a discretized set of $q = 19$ quantiles $\mathcal{Q} = \{0.05, 0.10, \dots, 0.90, 0.95\}$ using $n \times q$ univariate quantile regressions.

We use informative priors below to sharpen inference. Specifically, we first obtain parameter estimates from decades of U.S. data. The U.S. parameters’ posterior is then used to formulate an informative prior for the corresponding euro area model parameters. The weight put on U.S. prior information is estimated from euro area data, extending the three-step Gibbs sampler of [Khare and Hobert \(2012\)](#) to a novel four-step sampler. The difference between the U.S. and euro area estimates reflects the informational content of euro area data.

Web Appendix [C.1](#) presents our novel four-step Gibbs sampler used for posterior inference. The additional step estimates a prior precision parameter by sampling from a known inverse-Gamma distribution. Web Appendix [C.3](#) discusses prior specifications. We take particular care that the prior does not introduce cross-quantile or cross-variable restrictions, as this would invalidate an equation-by-equation estimation approach. Web Appendix [C.2](#) derives the fourth step of the Gibbs sampler.

3.4 Density forecasting and counterfactual scenarios

This section explains how density forecasts and counterfactual scenarios can be obtained from the SQVAR model (9). To simplify the exposition, we here consider a special case of (9), with $p = 1$ and $B(\gamma) = C_0(\gamma) = C_1(\gamma) = 0$.

Consider an n -vector $u_1^* \equiv [u_{11}^*, \dots, u_{n1}^*]'$ whose elements are draws from the i.i.d. uniform distribution with support on $(0, 1)$. Then a draw from the one-step ahead forecast distribution of x_{T+1} can be obtained as

$$x_{T+1}^* = (I_n - A_0(u_1^*))^{-1}(\omega(u_1^*) + A_1(u_1^*)x_T),$$

where I_n is the n -dimensional identity matrix, and where we have replaced γ with the draw u_1^* . The parameters associated with the selected quantiles u_1^* can be estimated (or were estimated previously and then stored). Conditional on x_{T+1}^* , a draw from the two-step ahead forecast distribution of x_{T+2} is $x_{T+2}^* = (I_n - A_0(u_2^*))^{-1}(\omega(u_2^*) + A_1(u_2^*)x_{T+1}^*)$, where u_2^* is another n -vector with i.i.d. draws from the standard uniform distribution. Iterating this process forward, we can obtain a sample path $(x_{T+1}^*, x_{T+2}^*, \dots, x_{T+H}^*)$ of any desired length H .

With n variables, q quantiles, and H steps ahead, there are $(q^n)^H$ possible paths at any forecast origin. Even for moderate values (e.g., $q = 19$ quantiles, $n = 5$ variables, and a forecasting horizon of $H = 16$ 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.” To explore the “tree” of all potential future paths at any time $t = 1, \dots, T$, we simulate S future paths for x_{t+h} , $h = 1, \dots, H$ quarters ahead. Once predictive densities are available, the estimation of GS and GL is straightforward. Web Appendix D.2 provides additional detail on the SQVAR’s multi-step ahead predictive distribution.

Rather than moving through the “tree” of potential future values of x_{t+h} completely at random, we can also focus on a subset of potential paths, or, in the extreme, consider only one path in isolation. Such a subset of potential paths can be thought of as a counterfactual ‘scenario,’ or

model-based thought experiment, that conditions on a fixed sequence of future conditional quantile realizations for a subset of variables in x_t . We use such scenarios when considering the potential of active macro-prudential policy for reducing economic growth volatility and negative skewness in Section 5.4. Web Appendix D.3 provides additional detail.

3.5 Quantile impulse response functions

We obtain quantile impulse response function estimates by simulation. Intuitively, given a draw from the parameters' posterior distribution, a structural shock to any variable can be propagated forward using (9). Doing this many times allows us to obtain predictive quantiles of all variables in the system at any horizon $h = 1, \dots, H$. These predictive quantiles are compared to their no-shock counterparts, which are also obtained by simulation. Obtaining QIRF estimates in this way requires a loop within a loop: an outside loop drawing from the posterior density of the parameters, and an inside loop propagating forward both the shock and no-shock scenarios. Web Appendix D discusses implementation details.

4 Euro area data

Our baseline statistical model for the endogenous variables x_t contains five endogenous variables and one exogenous variable (a global commodity price index). The five endogenous variables consist of a financial cycle indicator, annualized quarter-on-quarter HICP inflation, annualized quarter-on-quarter real GDP growth, a measure of financial stress, and the three-month OIS interest rate. Web Appendix E provides a detailed description of our data.

The financial cycle indicator used in our empirical analysis below is that of [Lang et al. \(2019\)](#). It is available at the quarterly frequency and 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 measure of system-wide financial stress, we use quarterly averages of the ECB’s revised daily composite indicator of systemic stress (CISS) as discussed in [Chavleishvili and Kremer \(2025\)](#). 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 *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.

To construct a consistent time series of short-term euro area interest rates, we splice together three time series: quarterly averages of the daily German Frankfurt Interbank Offered Rate (FI-BOR, 1990Q1 to 1993Q4), the three-month Euro Interbank Offered Rate (EURIBOR, 1994Q1 to 1998Q4), and the three-months-ahead euro Overnight Index Swap rate (1999Q1 to 2022Q4); see Web Appendix [E.4](#) for details.

5 The risk management approach in practice

This section first discusses model specification, model adequacy, and quantile impulse response functions. It then establishes that macro-prudential policy can contribute to decreasing an economy’s growth rate volatility, negative skewness, and excess kurtosis, mainly by stabilizing the overall macro-financial environment. Finally, we offer a metric of the macro-prudential policy stance, and study when a tightening of macro-prudential policy is most likely to be beneficial.

5.1 Model specification

We select a $n = 5$ -variable SQVAR (9) with $p = 4$ lags as our preferred model.¹¹ The vector of endogenous variables x_t contains a financial cycle indicator c_t , annualized quarter-on-quarter HICP inflation π_t , annualized quarter-on-quarter real GDP growth y_t , the financial stress measure CISS s_t , and the three-month OIS interest rate r_t , in this order, as $x_t = (c_t, \pi_t, y_t, s_t, r_t)'$.¹² Slow-moving variables, such as the financial cycle (which is constructed from multi-year growth rates), inflation, and real GDP growth, are ordered first, while fast-moving financial variables, such as the CISS and the OIS rate, are ordered last. The ordering of variables matters given the recursive information set defined in Section 3.1. Timing restrictions are standard in the empirical macro-financial VAR literature; see e.g. [Christiano et al. \(1999\)](#), [Kilian \(2009\)](#), and [Gilchrist and Zakrajsek \(2012\)](#), among many others.

We include four dummy variables d_t corresponding to four Covid-related observations between 2020Q1 and 2020Q4. This inclusion neutralizes the effect of these influential but unusual observations on the posterior parameter estimates.

Finally, we prevent the occurrence of a “price puzzle” (i.e., the counter-intuitive finding that a contractionary monetary policy shock causes an increase in inflation, rather than the expected decrease) in two complementary ways. First, we include global commodity prices as an exogenous variable z_t (and also consider $p = 4$ lags there). Including commodity prices has been shown to prevent or mitigate the price puzzle for U.S. data; see e.g. [Sims \(1992\)](#) and [Hanson \(2004\)](#).¹³ Second, we follow [Estrella \(2015\)](#) and impose a transmission lag on the structural (direct) impact of short-term interest rates on the real economy. Specifically, two elements of $A_1(\gamma)$ are restricted to zero (one element in the row for inflation and one element in the row for real GDP). As a result,

¹¹Information criteria, such as the Deviance Information Criterion DIC, AIC, and BIC, point to different numbers of lags p , depending on the considered criterion, variable, and quantile. Weighing all evidence, $p = 4$ is at the conservative upper end of the indicated values for our quarterly data.

¹²The ECB working paper version of this paper undertakes an extensive variable selection exercise, indicating that short-term interbank rates can be a useful variable to consider in an SQVAR. In addition, short-term interest rates have been linked to the accumulation of financial vulnerabilities over time, see e.g. [Boissay et al. \(2016\)](#), and to increasing leverage and reach-for-yield behavior in financial markets, see e.g. [Brunnermeier and Sannikov \(2014\)](#) and [Akinci et al. \(2020\)](#). Short-term interest rates, in turn, respond to GDP and inflation through systematic monetary policy.

¹³The parameters in $C(\gamma)$ are restricted such that global commodity prices can only directly impact the inflation process. Doing so prevents a positive response of GDP to a global negative supply shock.

short-term interest rates (i.e., monetary policy) can still have a first-lag impact on the real economy, but, given the triangular structure of (9), only through the financial cycle (ordered first in x_t).

5.2 Model adequacy, prior robustness, and additional robustness checks

Web Appendix F presents the outcomes of an extensive battery of model adequacy, prior robustness, and additional robustness checks.

Web Appendix F.1 provides an out-of-sample assessment of the SQVAR model’s predictive density for real GDP growth, the key variable of interest in Sections 5.3 – 5.5. We find that the estimated model’s out-of-sample predictive quantiles for real GDP growth are appropriately calibrated, meaning that the nominal and empirical confidence levels are close. This is an important assessment, as maximizing the asymmetric Laplace log-likelihood function (or minimizing the Quantile Regression “check function”) guarantees the correct calibration of the one-step-ahead conditional quantiles *in-sample*, but not out-of-sample. We evaluate the predictive quantiles using Engle and Manganelli (2004)’s Dynamic Quantile test, and find that the empirical model captures the dynamic dependencies in the data appropriately.¹⁴

Web Appendix F.2 studies the robustness of our parameter estimates to different prior choices for the euro area SQVAR’s parameters. We perform two additional estimation exercises, deliberately moving away from the U.S.-data-based informative prior. We consider a standard Minnesota prior to estimate the euro area SQVAR’s parameters (shrinking coefficients towards a random walk, favoring parsimony) and an uninformative flat prior (allowing the euro area data alone to dominate the estimation). The QIRFs resulting from these two alternative prior specifications are approximately similar, but less precisely estimated in the second case.

Web Appendix F.3 studies the robustness of our parameter estimates to adopting a different hyperprior parameterization that implies more global shrinkage towards the U.S. estimates. The QIRFs remain approximately similar when this alternative hyperprior for the shrinkage parameter

¹⁴A recent study by Surprenant (2025) undertakes a comprehensive out-of-sample forecasting horse race exercise involving different QVAR, linear VAR, and VAR-SV model specifications for macro-financial data quite similar to ours. The QVAR models frequently improve upon the benchmarks and seldom perform worse.

$\lambda_i(\gamma)$ is adopted. We strongly prefer a loose hyperprior, however, corresponding to a prior view that $\lambda_i(\gamma)$ may need to be different in the center of the distribution and the tail, and that they may need to be different for financial and macro variables.

Web Appendix F.4 studies the effects of modeling the macro-financial data at a monthly instead of a quarterly frequency. The effects are modest. This is reassuring as our shock identification approach relies predominantly on timing restrictions, which could potentially be affected by the data sampling frequency.

Web Appendix F.5 studies the effect of ending the estimation samples in 2008Q2 (excluding the euro sovereign debt crisis, Covid-19 period, and 2022 energy crisis) and in 2019Q4 (excluding the Covid-19 period and 2022 energy crisis). Removing informative episodes from the estimation sample changes the estimates to some extent, as is intuitive.

Finally, Web Appendix F.6 studies the effect of removing Estrella (2015)’s transition lag as described in Section 5.1, allowing for an immediate impact of interest rates on inflation and real GDP. Without the transmission lag, a price puzzle occurs, with inflation showing a counter-intuitive initial increase in response to a monetary policy tightening. The other QIRFs remain similar.

Weighing the above, we conclude that the empirical model as specified in Section 5.1 and studied in Sections 5.3 – 5.5 combines model tractability with the ability to explore a rich set of relevant questions empirically given the data at hand.

5.3 Parameter estimates and quantile impulse response functions

Figures G.1 to G.8 in Web Appendix G.1 report our posterior mean estimates of all SQVAR model parameters. Posterior mean estimates are shown with corresponding 95% credible intervals. Least squares parameter estimates are visible as horizontal lines and are provided as a point of comparison. The parameter posterior estimates can differ substantially across quantiles, and from their respective ordinary least squares estimates. Parameter heterogeneity across quantiles is particularly visible for the parameter matrices A_1 to A_4 .

Figure 1 plots quantile impulse response functions (QIRF) for all endogenous variables as

implied by the parameter estimates. For readability, the figure plots the upper (90th percentile) and lower (10th percentile) tails as well as the median (50th percentile). The QIRFs summarize all contemporaneous and lagged interactions between all endogenous variables. The generally rather low contemporaneous relationships between the model variables (captured via A_0 ; see Figure G.2) imply that the QIRF are qualitatively but also quantitatively not strongly affected by the specific ordering of the variables in x_t .

In some cases, the QIRFs show pronounced asymmetries, while in other cases the differences between the quantiles are less pronounced and within the credible intervals.

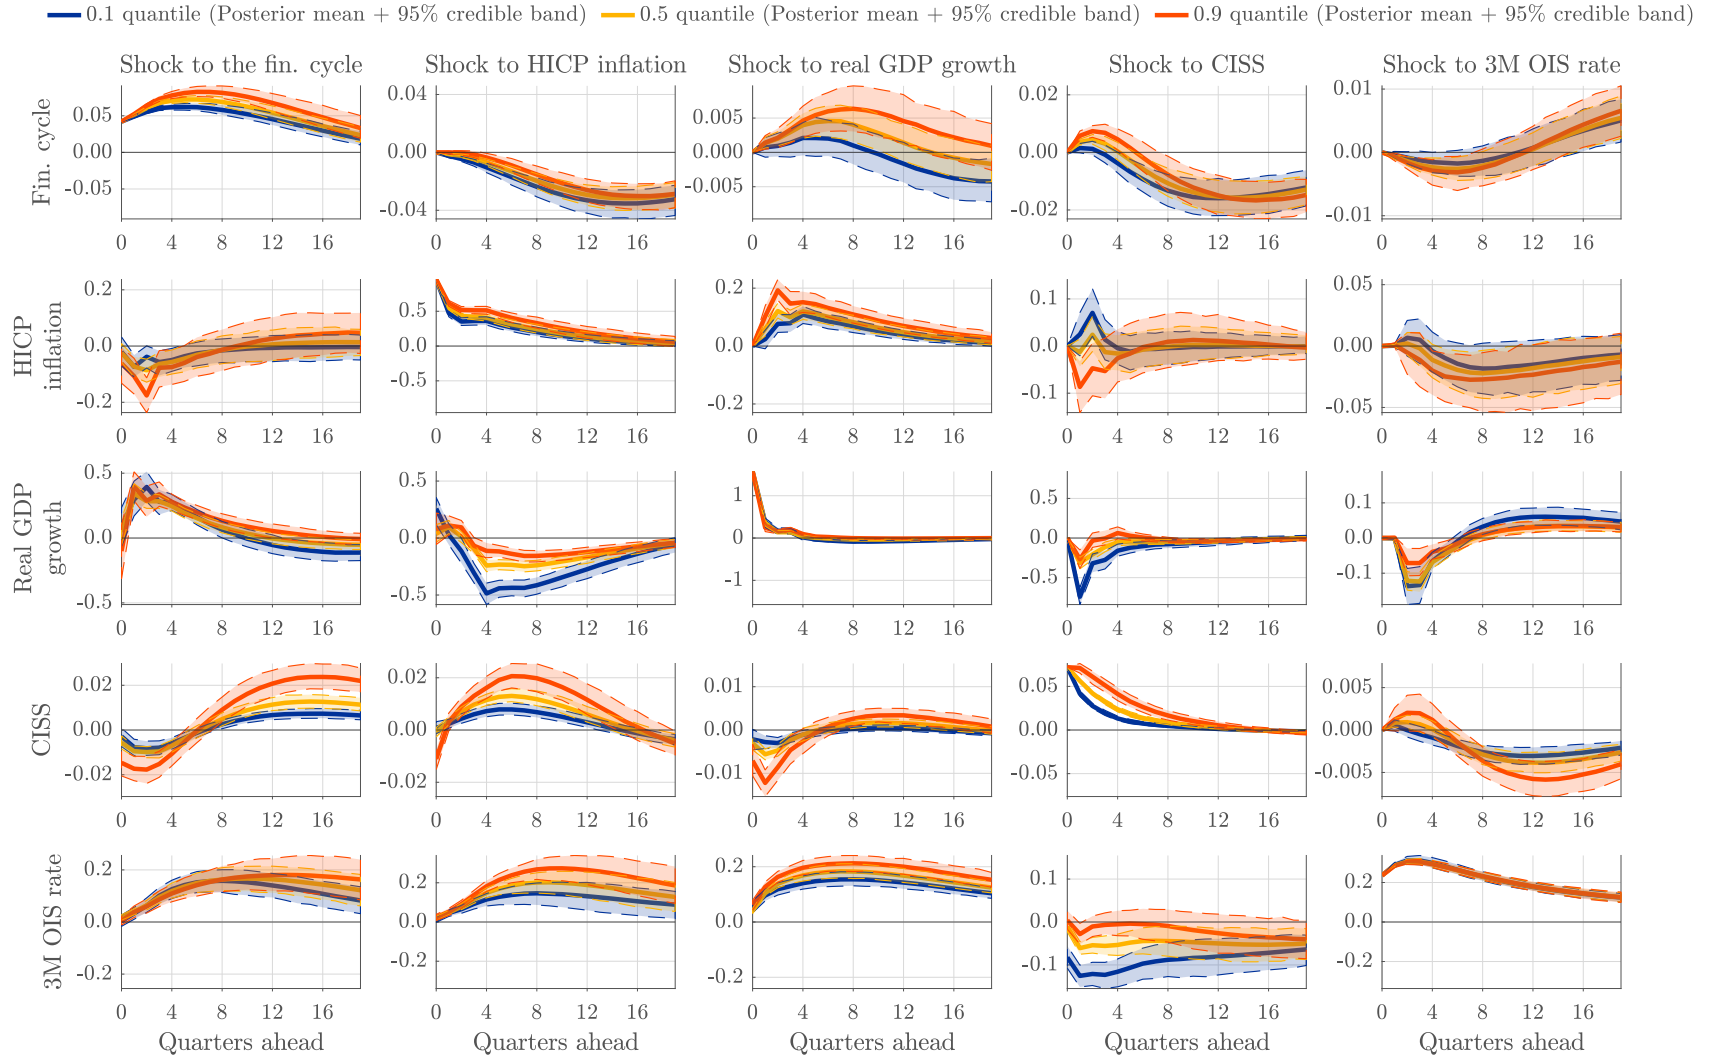
Consistent with the vulnerable growth literature (e.g., [Adrian et al., 2019](#), [Caldara et al., 2021](#), [Korobilis and Schröder, 2024](#)), the response of real GDP growth to a shock in the CISS (Figure 1's panel [3,4]) is much stronger in its left tail (the 0.1 quantile) than in its center and right tail (the 0.9 quantile). Thus, a given increase in financial stress depresses economic growth much more strongly in the lower part of the distribution than in the upper part.

Figure 1's panel [2,4] reports the impact of financial stress on euro area inflation. Both tails are sensitive to financial stress, and respond with opposite signs. As a result, a shock to stress (i.e., a tightening of financial conditions) leads to less inflation uncertainty in the future (cf., for example, [López-Salido and Loria, 2024](#)).

The response of the CISS to a shock to the financial cycle (Figure 1's panel [4,1]) indicates an initial dampening of stress followed by an increase after approximately two years. This pattern is particularly pronounced for the CISS's upper tail (the 0.9 quantile), which displays a marked negative response in the short term that reverses after approximately eight quarters. Vice versa, positive shocks to the CISS deflate the financial cycle in the medium term (see Figure 1's panel [1,4]). These findings are in line with work on the “term structure” of GaR (see, for example, [Adrian et al., 2021](#) and [Lang et al., 2023](#)).

Figure 1: Quantile impulse response function estimates

Quantile impulse response functions implied by the posterior estimates reported in Figures G.1-G.8. 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), HICP inflation (second row), real GDP growth (third row), financial stress (CISS, fourth row), and the three-month OIS rate (fifth row). Credible intervals are at a 95% level. The estimation sample is 1990Q1 to 2022Q4.



The numerous “macro at risk” stylized facts discussed above are not necessarily new to the literature. Instead, they are, overall, in line with cited previous work. We stress, however, that, to our knowledge, we captured virtually *all of them*, simultaneously, for euro area data, within a *single integrated dynamic framework*. The reported credible intervals around the QIRFs are obtained from appropriate single-step inference and did not require first modeling a set of conditional quantiles and afterwards fitting (skewed) parametric densities.

Figure G.9 in Web Appendix G.1 presents the posterior mean estimates of the hyperprior parameter ($\lambda_i(\gamma)$) that scales the (U.S.) prior’s variance. The estimates $\hat{\lambda}_i(\gamma)$ take an interesting U-shape. The U-shape suggests that the euro area economy resembles the U.S. economy most closely in the center of the predictive distributions, and less so in the lower and upper tails. This is intuitive, as the euro area has been in trouble, or doing very well, mostly for euro area-specific reasons (such as, for example, the 1992 recession following the German reunification, and the 2010-2012 euro area sovereign debt crisis and subsequent recovery).¹⁵

Figure G.11 in Web Appendix G.2 plots the 5%, 10%, 50%, 90%, and 95% predictive quantile of real GDP growth, at any time t , as implied by our statistical model. As a result of the pronounced macro-financial interactions depicted in Figure 1, the lower quantiles for future real GDP growth are considerably more volatile than its upper quantiles. For example, for the one-quarter-ahead forecast, the predictive quantiles’ standard deviations decrease (almost) monotonically: 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. We also note that the asymmetry in the predictive density decreases with the forecasting horizon.

Web Appendix H reports SQVAR parameter and credible interval estimates for U.S. data. The estimates are broadly in line with those for the euro area. For example, growing financial vulnerabilities shift the right tail of the U.S. CISS towards more positive values. A shock to the U.S. CISS shifts the left tail of the predictive GDP growth distribution towards more negative values while

¹⁵Figure H.9 in Web Appendix H presents the posterior mean estimates $\hat{\lambda}_i(\gamma)$ from U.S. data. These estimates do not follow a U-shape, suggesting that it is not the down-weighting of an increasing share of the sample observations (via the QR’s log-likelihood function as one moves further into the tail) that necessitates a less informative prior *per se*, but economic factors.

leaving the right tail less affected.

Overall, the pronounced asymmetries in Figures 1, G.11, and H.10 lend support to nonlinear macroeconomic models with a financial sector that feature occasionally binding financing constraints and/or a “volatility paradox.” Examples of such macro models include Brunnermeier and Sannikov (2014), Boissay et al. (2016), He and Krishnamurthy (2019), and Mendicino et al. (2024).

5.4 The potential of macro-prudential policy

This section investigates the difference an active macro-prudential policy regime could have made on the evolution of key macroeconomic variables had such an active regime existed between 1991Q1 and 2022Q4. Specifically, we consider three ‘worlds’: one actually-observed, and predominately characterized by a passive macro-prudential policy; one characterized by a counterfactual “aggressive” macro-prudential policy regime that forcefully smooths the financial cycle; and, finally, one characterized by a “moderate” macro-prudential policy regime that strikes a middle ground between the former two.¹⁶


To compare macro-economic outcomes across different counterfactual scenarios we need to assume that the presence of a new policy regime does not change that SQVAR’s parameters to such an extent that it makes the thought experiment uninteresting. This is known as the Lucas (1976) critique: changes in economic policy can influence agents’ expectations and thus behavior, potentially altering a VAR model’s reduced-form parameters. We here assume that the SQVAR parameters remain unaffected across scenarios, and note that this is more believable for the “moderate” than for the “aggressive” policy regime (cf., for example, Leeper and Zha, 2003).

To implement the counterfactual scenarios, we first determine the conditional quantiles that were realized historically using the estimated SQVAR. This is straightforward: At each $t = 1, \dots, T$, with SQVAR parameters fixed at, for example, their full sample posterior estimates, we forecast the endogenous variables one period ahead, and then note the conditional quantiles

¹⁶For our counterfactual policy scenarios to be meaningful, macro-prudential policy must employ instruments that are effective in managing the financial cycle. We recall that Web Appendix A.1 summarizes the nascent empirical literature on the efficacy of capital-, liquidity-, and borrower-based macro-prudential instruments.

Figure 2: The financial cycle in two different experiments

The historical financial cycle (“Realized”) and two estimates of the flattened financial cycle that would have been obtained had the “aggressive” or the “moderate” counterfactual active policy regime been adopted. The credible bands are at a 68% confidence level, based on 10,000 posterior draws from the full-sample posterior density, using 1991Q1 as the forecast origin. For each draw, we identify the historically-realized conditional quantiles of all variables and then perform the policy experiment as described in the main text. Gray-shaded areas indicate CEPR euro area recession periods.



that are closest to the realized data. Put differently, we determine the conditional quantiles (from among a grid of $q = 19$ quantiles; see Section 3.3) that minimize the forecast error. This procedure yields a series of conditional quantile realizations. By construction, starting from $t = 1$, the historical conditional quantile realizations replicate the observed data. This process is similar to a historical decomposition in a linear SVAR model; see e.g. Kilian and Lütkepohl (2017, Ch. 4).

The active policies are implemented by allowing the policymaker to adjust the financial cycle by influencing its conditional quantile realizations. The conditional quantiles of all other variables in the system remain as previously identified.

Under the “aggressive” policy regime, the financial cycle is set to its conditional median at each time. Here, a (potentially unrealistically) powerful macro-prudential policymaker does whatever it takes to ensure that the financial cycle evolves as “expected” conditional on the past. By contrast, the “moderate” policymaker is less aggressive and merely adjusts the financial cycle’s quantile realizations by half relative to the median. For example, this implies that a historical 0.9

conditional quantile realization is now adjusted downward to the 0.7 conditional quantile (halfway to the median), and a historical 0.3 conditional quantile realization is now adjusted upwards to the 0.4 conditional quantile. This strikes a middle ground between where the financial cycle has historically been and where it goes to under the aggressive scenario.

Figure 2 plots the historical financial cycle (black line) against the financial cycle that obtains in the two counterfactual ‘worlds’ (blue and orange solid lines). By construction, each active policy flattens the financial cycle compared to its realized path (up to a discretization error), particularly in the first two-thirds of the sample. We observe that the financial booms during the late 1990s and mid-2000s, before the GFC in 2008–2009, are particularly affected, suggesting that these booms were driven primarily by the financial cycle’s high quantile realizations rather than those of the other endogenous variables in the system. The figure also indicates that the post-2008 financial cycle contractions would have been considerably less severe under the two active policies.

Table 1 reports the first four unconditional moments of key euro area macroeconomic variables across the three scenarios.¹⁷ Strikingly, both types of active macro-prudential policy are associated with higher trend (mean) real GDP growth, lower growth volatility, and lower unwelcome skewness and kurtosis. Both types of active policy are thus associated with clear welfare gains, using any sensible evaluation criterion that values growth and penalizes downside risk.¹⁸

The clear beneficial impact on GDP can be traced back to lower financial stress: the CISS’s mean shifts down, indicating looser financial conditions on average, and the CISS’s volatility decreases along with its positive skewness and kurtosis. Interestingly, managing the financial cycle also reduced inflation volatility, along with its positive skewness and (mild) kurtosis. This, in turn, allows short-term (monetary policy) interest rates to be less volatile.

¹⁷Web Appendix G.1’s Figure G.10 reports posterior densities of the first three unconditional moments.

¹⁸We can approximate the welfare gain by applying an unconditional version of the objective function (6), that is, using the unconditional mean growth rate and unconditional growth shortfall as arguments. If the policy maker cares twice as much about mean growth than expected growth shortfall, i.e. $\lambda^p = 1.5$, then leaning against the financial cycle leads to utility gains between 3% and 6%.

Table 1: Moments of euro area macro variables across policy experiments

Unconditional moments (in rows) of key euro area macroeconomic variables (in four panels) across three scenarios (“Realized,” “Aggressive,” “Moderate,” in columns). The “Realized” columns refer to historical data and are provided as points of comparison. The “Aggressive” and “Moderate” columns refer to two policy experiments and report posterior means and 95% credible intervals in square brackets. The policy experiments are described in the main text. The year 2020 is excluded when computing moments owing to Covid-19.

HICP inflation				Real GDP growth		
	Realized	“Aggressive” policy	“Moderate” policy	Realized	“Aggressive” policy	“Moderate” policy
Mean	2.304	2.327 [2.082, 2.645]	2.303 [2.162, 2.453]	1.698	1.708 [1.301, 2.064]	1.728 [1.604, 1.849]
Std. dev.	1.946	1.857 [1.681, 2.048]	1.852 [1.710, 2.011]	2.395	2.052 [1.823, 2.285]	2.109 [1.961, 2.268]
Skew	1.528	1.141 [0.671, 1.616]	1.231 [0.826, 1.636]	-1.742	-0.982 [-1.713, -0.268]	-1.295 [-1.881, -0.680]
Excess kurtosis	4.425	2.924 [1.312, 4.693]	3.259 [1.1823, 4.845]	9.967	6.521 [3.504, 9.834]	7.957 [5.455, 10.743]
CISS				Three-month OIS rate		
	Realized	“Aggressive” policy	“Moderate” policy	Realized	“Aggressive” policy	“Moderate” policy
Mean	0.148	0.127 [0.073, 0.170]	0.135 [0.123, 0.148]	3.074	3.072 [2.010, 4.108]	3.108 [2.795, 3.438]
Std. dev.	0.169	0.124 [0.095, 0.168]	0.132 [0.118, 0.147]	3.330	3.261 [2.867, 3.737]	3.211 [3.026, 3.397]
Skew	1.893	1.292 [0.929, 1.841]	1.487 [1.176, 1.812]	0.907	1.102 [0.586, 1.606]	1.064 [0.867, 1.261]
Excess kurtosis	4.400	1.492 [0.102, 3.963]	2.371 [0.993, 3.991]	-0.092	0.215 [-0.744, 1.420]	0.150 [-0.241, 0.603]

We conclude that an active macro-prudential policy regime can, in our experiments, considerably improve macroeconomic outcomes by increasing economic growth while decreasing its volatility, negative skewness, and excess kurtosis. This is achieved by decreasing the probability of observing periods of high financial stress.

5.5 Macro-prudential policy stance

While counterfactual experiments, such as those of Section 5.4, can encourage the use of macro-prudential instruments, policymakers tasked with policy implementation need to be forward-looking and condition their actions on available data. An active debate in policy circles revolves around the question of how to measure the macro-prudential policy stance. In other words, should macro-prudential policy be tightened or loosened today conditional on currently available information? We use the decision framework in Section 2.2 to address this question.

We assume that macro-prudential instruments are available and can be used to influence the financial cycle, and ask when it is beneficial to do so. The thought experiment of this section is the following: How does the macro-prudential policymaker's objective function value (7) change if the financial cycle is marginally lowered now, to be increased later on? If the change is positive, we conclude that the macro-prudential stance is too loose (as it would benefit from a lower, less buoyant financial cycle). If, on the other hand, the answer is negative, then macro-prudential policy is too tight.

We distinguish between a passive and an active policy scenario. The passive policy scenario consists of a time- t conditional forecast of the economy. The active policy scenario entails reducing the financial cycle in quarter $h = 1$ by a certain amount, and raising it again in quarter $H/2 + 1$ by the same amount when the economy is closer to its unconditional distribution. The financial cycle is shocked by one half standard deviation of the residuals in the SQVAR's financial cycle equation at the median. This amount is approximately equal to the impact of an average capital-based macro-prudential intervention on credit growth and house prices in the empirical study by Ampudia et al. (2021, Sec. 4.1), and is thus of an empirically-relevant and interpretable magnitude.

Each policy scenario is evaluated in welfare terms using the objective function (7). At any time t , future economic growth rates and growth shortfalls are estimated from 50,000 simulations of potential future values of y_{t+h} . To obtain credible intervals, each forward simulation is repeated 200 times, each time using a new draw from the parameters' posterior density. To fully operationalize (6), we choose $H = 16$ quarters, no discounting of the near future ($\delta = 1$ for $h = 1, \dots, 16$ and $\delta = 0$ thereafter), the penalty parameter $\lambda^p \in \{1.5, 3\}$, and $\tau = 0$. The choice $\lambda^p = 1.5$ implies that the policymaker cares twice as much about future trend growth than she cares about reducing downside risk. The choice $\lambda^p = 3$ implies that the policymaker cares twice as much about reducing downside risk than about fostering future growth. We evaluate the policymaker's utility function twice, once for the active scenario and once for the passive scenario, and study the difference between the two values at any time t .


Figure 3 presents the objective function difference Δu_t associated with adopting the active macro-prudential policy along with 68% and 95% credible intervals. We observe that adopting the active policy is not equally beneficial at all times. The benefits from leaning against the financial cycle are maximal during the late 1990s, before the bust of the dot-com boom in 2000, and during the mid-2000s before the onset of the GFC in 2007. These findings are intuitive, as the financial system was particularly sanguine during these times, arguably seeding the respective busts later on (e.g., [Boissay et al., 2016](#)).

The benefits from leaning against the financial cycle are estimated to be negative during the early 1990s, following the end of the dot-com boom in early 2000, during the GFC in 2008, and during the euro area sovereign debt crisis in 2011–2012. This is again intuitive, as the financial system was already deleveraging during these times, and requiring more would add insult to injury.

Importantly, Figure 3 can suggest an approximate rule-of-thumb to set the Basel-III counter-cyclical capital buffer in practice: increase the buffer by 0.25 percentage points every year, within the given regulatory range between zero and 2.5%, except when it's obvious that the economy is in, or about to enter, a contraction. In this way, the buffer can be brought from zero to full 2.5% capacity in about ten years of expansion, to lean against exuberance during the expansion and

Figure 3: Macro-prudential policy stance

The time- t benefit $\widehat{\Delta u}_t = [\hat{u}_t(\text{active}) - \hat{u}_t(\text{passive})]$ from marginally lowering the financial cycle now to increase it later on (active policy) vis-à-vis a time- t conditional unrestricted forecast (passive policy). If $\Delta u_t > 0$, then the economy would benefit from a lower financial cycle and the macro-prudential policy stance is too loose. If $\Delta u_t < 0$, then the policy stance is too tight. Credible bands and posterior mean estimates are based on 200 draws from the posterior distribution and $2 \times 50,000$ simulations of all conditional (on time- t information) distributions per draw from the posterior. Parameters are chosen as $\tau = 0$, $\delta = 1$ for $h = 1, \dots, 16$ and $\delta = 0$ thereafter, and $\lambda^p \in \{1.5, 3\}$; see the main text for details. Credible intervals for the $\lambda^p = 3$ -case are reported at a 68% and 95% confidence level. Parameter estimates are kept unchanged over time at their full sample (1990Q1 to 2022Q4) posterior estimates. Gray-shaded areas indicate CEPR-dated euro area recessions.



increase the loss absorbing capacity of the financial system once a financial crisis materializes.

We conclude this section by studying the circumstances under which macro-prudential interventions are most likely to be beneficial. To address this question we relate our estimate of macro-prudential stance, $\widehat{\Delta u}_t = \hat{u}_t(\text{active}) - \hat{u}_t(\text{passive})$, to the current macro-financial environment $(x'_t, z'_t)'$. Table 2 reports the respective least squares regression parameter estimates. We used $\lambda^p = 3$ to obtain the left hand-side variable, but obtain similar results for $\lambda^p = 1.5$.

We highlight five findings. First, a tighter macro-prudential policy stance is called for when the financial cycle takes high values. This is intuitive and in line with the purpose of the indicator; see [Lang et al. \(2019\)](#). Second, low short-term interest rates (3mOIS) are associated with benefits from

Table 2: Macro-prudential stance and the macroeconomy

Least squares estimates of the posterior mean of $100 \times \Delta u_t$ (i.e., $100 \times$ the solid blue line in Figure 3) on the macro variables $(x'_t, z'_t)'$ from the SQVAR model. The t-statistics are based on Newey-West HAC standard errors. The regression constant is not reported. Estimation sample: 1991Q1 to 2022Q4, $N = 128$.

Variable	val	t-stat	Variable	val	t-stat
Fin. cycle _t	2.17	7.66	CISS _t	-5.13	-10.56
3mOIS _t	-0.16	-8.07	HICP inflation _t	-0.19	-5.00
RGDP growth _t	0.06	2.65	Commodity idx _t	0.00	0.82
R ²	0.80				

macro-prudential tightening. This is in line with studies explaining how low-for-long monetary policy interest rates can encourage a “reach-for-yield” behavior and low lending standards; see e.g. [Rajan \(2006\)](#) and [Dell’Ariccia et al. \(2014\)](#). Low short-term interest rates require low and stable inflation to materialize. Third, macro-prudential tightening is relatively more beneficial in an economic expansion (or boom) than a contraction. This is again intuitive and in line with e.g. [Boissay et al. \(2016\)](#), where an economic boom seeds optimism and increases the probability of a bust later on. Fourth, tightening macro-prudential policy tends to be beneficial when financial stress (CISS) is low. This is in line with an important literature on the “financial stability-” or “volatility paradox,” according to which the financial system can look strongest precisely when it is most vulnerable; see e.g. [Brunnermeier and Sannikov \(2014\)](#) and [Mendicino et al. \(2024\)](#). In these models, what looks ostensibly like low risk is, in fact, a sign of aggressive risk taking. Finally, global commodity prices do not appear to be important for explaining variation in $\widehat{\Delta u}_t$.

6 Conclusion

We formalized the macro-prudential policymaker’s decision problem under uncertainty and studied it for the euro area. Flexible and potentially asymmetric predictive distributions are obtained from a novel Bayesian structural quantile vector autoregressive model that allows us to study the nonlinear relationship between economic growth and various other factors such as financial stress,

the financial cycle, short-term interest rates, and inflation. We documented substantial asymmetries in the predictive distribution of real GDP growth and the responses of macroeconomic variables to, for example, financial shocks. Considering counterfactual scenarios allowed us to study the potential of macro-prudential policy, quantify the macro-prudential policy stance, and study when macro-prudential interventions are likely to be beneficial.

This paper entails several possible routes for future research. First, to our knowledge, methodological work on shock identification in multivariate quantile models, e.g. through instrumental variables or sign restrictions, is currently an open field (but see [Schüler \(2020\)](#), [Chavleishvili and Manganelli \(2024\)](#), and [Iacopini et al. \(2024\)](#)). Second, panel-SQVAR approaches could be considered, possibly incorporating hierarchical specifications of random coefficients as in, e.g., [Jarocinski \(2010\)](#). Finally, the relative forecasting performance of recent macro-at-risk models, and combinations thereof, is only now starting to be comprehensively assessed ([Surprenant, 2025](#)).

References

- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review* 109(4), 1263–89.
- Adrian, T., N. Boyarchenko, and D. Giannone (2021). Multimodality in macro-financial dynamics. *International Economic Review* 62(2), 861–86.
- Adrian, T., F. Grinberg, N. Liang, and S. Malik (2022). The term structure of growth-at-risk. *American Economic Journal: Macroeconomics* 14(3), 283–323.
- Akinci, O., G. Benigno, M. Del Negro, and Q. Albert (2020). The financial (in)stability real interest rate, r^{**} . *Federal Reserve Bank of New York Staff Reports* 946.
- 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.
- Antolín-Díaz, J. and J. F. Rubio-Ramírez (2024). Narrative Sign Restrictions for SVARs. *American Economic Review* 108(10), 2802–29.
- Artzner, P., F. Delbaen, J. M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance* 9(3), 203–228.

- Boissay, F., F. Collard, and F. Smets (2016). Booms and banking crises. *Journal of Political Economy* 124(2), 489–538.
- Brandao-Marques, L., G. Gelos, M. Narita, and E. Nier (2020). Leaning against the wind: An empirical cost-benefit analysis. *IMF Working Paper*, 1–54.
- Brownlees, C. and A. B. M. Souza (2021). Backtesting global growth-at-risk. *Journal of Monetary Economics* 118, 312–330.
- Brunnermeier, M., A. Crocket, C. Goodhart, A. Persaud, and H. Shin (2009). The fundamental principles of financial regulation. *Geneva Report on the World Economy* 11.
- Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *The American Economic Review* 104(2), 379–421.
- Caldara, D., C. Scotti, and M. Zhong (2021). Macroeconomic and financial risks: A tale of volatility. *Federal Reserve Board International Finance discussion papers*.
- Carney, M. (2020). The grand unifying theory (and practice) of macroprudential policy. *Speech given at University College London on 5 March 2020*, 1 – 22.
- Carriero, A., T. E. Clark, and M. Marcellino (2023). Capturing macroeconomic tail risks with Bayesian vector autoregressions. *Journal of Money, Credit and Banking*, forthcoming 2023.
- Cecchetti, S. G. (2006). Measuring the macroeconomic risks posed by asset price booms. *NBER Working Paper No. 12542*, 1–31.
- Chavleishvili, S. and M. Kremer (2025). CISS of death: Measuring financial crises in real time. *Review of Finance* 29(3), 685 – 710.
- Chavleishvili, S. and S. Manganelli (2024). Forecasting and stress testing with quantile vector autoregression. *Journal of Applied Econometrics* 39(1), 66–85.
- Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary policy shocks: What have we learned and to what end? In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics* (1 ed.), pp. 1–84. New York: Elsevier.
- 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*.

- de Larosiere, J. (2009). The high level group on financial supervision in the E.U. *Report for the European Commission*.
- Delle Monache, D., A. De Polis, and I. Petrella (2023). Modeling and forecasting macroeconomic downside risk. *Journal of Business and Economic Statistics*, *forthcoming*.
- Dell’Ariccia, G., L. Laeven, and R. Marquez (2014). Real interest rates, leverage, and bank risk-taking. *Journal of Economic Theory* 149, 65–99.
- Duprey, T. and A. Ueberfeldt (2020). Managing GDP tail risk. *Bank of Canada working paper 2020–03*, 1–63.
- ECB (2019). European Central Bank Financial Stability Review, May 2019. .
- 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.
- Forni, M., L. Gambetti, and L. Sala (2021). Downside and Upside Uncertainty Shocks. *mimeo*, 1–63.
- FSB (2009). Report to G20 Finance ministers and governors: Guidance to assess the systemic importance of financial institutions, markets and instruments – initial considerations. www.fsb.org.
- Giannone, D., M. Lenza, and G. E. Primiceri (2015). Prior selection for vector autoregression. *The Review of Economics and Statistics* 97(2), 436–451.
- Gilchrist, S. and E. Zakrajsek (2012). Credit spreads and business cycle fluctuations. *American Economic Review* 102(4), 1692–1720.
- Greenspan, A. (2003). Opening Remarks to Monetary policy and uncertainty: Adapting to a changing economy. In *Proceedings of the Federal Reserve Bank of Kansas City Symposium*, pp. 1–12. Carnegie-Rochester conference series on public policy.
- Hanson, M. S. (2004). The “price puzzle” reconsidered. *Journal of Monetary Economics* 51(7), 1385–1413.
- He, Z. and A. Krishnamurthy (2019). A macroeconomic framework for quantifying systemic risk. *American Economic Journal: Macroeconomics* 11(4), 1–37.
- Iacopini, M., F. Ravazzolo, and L. Rossini (2024). Bayesian multivariate quantile regression with alternative time-varying volatility specifications. *mimeo*.

- IMF (2017). Global financial stability report: Is growth at risk? *IMF Global Financial Stability Report* 3, 91–118.
- Iseringhausen, M., I. Petrella, and K. Theodoridis (2024). Aggregate skewness and the business cycle. *Review of Economics and Statistics*, forthcoming.
- Jarocinski, M. (2010). Responses to monetary policy shocks in the east and the west of Europe: a comparison. *Journal of Applied Econometrics* 25(5), 833–868.
- 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. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053–1069.
- Kilian, L. and H. Lütkepohl (2017). *Structural vector autoregressive analysis*. Cambridge University Press.
- 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. and Z. Xiao (2006). Quantile autoregression. *Journal of the American Statistical Association* 101(475), 980–990.
- Korobilis, D. and M. Schröder (2024). Monitoring multicountry macroeconomic risk. *Journal of Econometrics*, forthcoming.
- Kozumi, H. and G. Kobayashi (2011). Gibbs sampling methods for Bayesian quantile regression. *Journal of Statistical Computation and Simulation* 81(11), 1565–1578.
- 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.
- Lang, J. H., M. Rusnak, and M. Greiwe (2023). Medium-term growth-at-risk in the euro area. *ECB Working Paper Nr. 2808*, 1–41.
- Leeper, E. M. and T. Zha (2003). Modest policy interventions. *Journal of Monetary Economics* 50, 1673–1700.
- Lucas, R. (1976). Econometric policy evaluation: A critique. In K. Brunner and A. Meltzer (Eds.), *The Phillips Curve and labor markets* (1 ed.), pp. 19–46. New York: Carnegie-Rochester conference series on public policy, American Elsevier.

- López-Salido, D. and F. Loria (2024). Inflation at risk. *Journal of Monetary Economics*, 103570.
- McNeil, A., R. Frey, and P. Embrechts (2005). *Quantitative Risk Management*. Princeton: Princeton University Press.
- Mendicino, C., K. Nikolov, J. Rubio-Ramirez, J. Suarez, and D. Supera (2024). Twin defaults and capital requirements. *Journal of Finance*, forthcoming.
- New York Fed (2022). Outlook-at-Risk: real GDP Growth, unemployment, and inflation. <https://www.newyorkfed.org/research/policy/outlook-at-risk..>
- Osborne, G. (2012). Speech by the Chancellor of the Exchequer, Rt Hon George Osborne MP at the Lord Mayor’s dinner for bankers and merchants of the City of London, on 14 June 2012. <https://www.gov.uk/>.
- Plagborg-Møller, M., L. Reichlin, G. Ricco, and T. Hasenzagl (2020). When is growth at risk? *Brookings Papers on Economic Activity*.
- Plagborg-Møller, M. and C. K. Wolf (2021). Local projections and VARs estimate the same impulse responses. *Econometrica* 89(2), 955 – 980.
- Prasad, A., S. Elekdag, P. Jeasakul, R. Lafarguette, A. Alter, A. F. Xiaochen, and C. Wang (2019). Growth at risk: Concept and applications in IMF country surveillance. *IMF Working Paper*.
- Rajan, R. G. (2006). Has finance made the world riskier? *European Financial Management* 12(4), 499–533.
- Schüler, Y. S. (2020). The impact of uncertainty and certainty shocks. *Deutsche Bundesbank Discussion Paper* 14, 1–47.
- Sims, C. A. (1992). Interpreting the macroeconomic time series facts: The effects of monetary policy. *European Economic Review* 36, 975–1000.
- 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.
- Van der Ghote, A. (2021). Interactions and coordination between monetary and macro-prudential policies. *American Economic Journal: Macroeconomics* 13(1), 1–34.
- Wei, Y. (2008). An approach to multivariate covariate-dependent quantile contours with application to bivariate conditional growth charts. *Journal of the American Statistical Association* 103(481), 397 – 409.
- Yu, K. and R. A. Moyeed (2001). Bayesian quantile regression. *Statistics & Probability Letters* 54, 437–447.