

The risk management approach to macro-prudential policy*

*Sulkhan Chavleishvili,^(a) Robert F. Engle,^(b) Stephan Fahr,^(a)
Manfred Kremer,^(a) Simone Manganelli,^(a) Bernd Schwaab^(a)*

^(a) European Central Bank, Financial Research,

^(b) New York University Stern School of Business

April 15, 2021

Abstract

Macro-prudential authorities need to assess medium-term downside risks to the real economy, caused by severe financial shocks. Before activating policy measures, they also need to consider their short-term negative impact. This gives rise to a risk management problem, an inter-temporal trade-off between expected growth and downside risk. Predictive distributions are estimated with structural quantile vector autoregressive models that relate economic growth to measures of financial stress and the financial cycle. An empirical study with euro area and U.S. data shows how to construct indicators of macro-prudential policy stance and to assess when interventions may be beneficial.

Keywords: Growth-at-risk, stress testing, quantile vector autoregression, financial conditions, macro-prudential policy.

JEL classification: G21, C33.

*We would like to thank Jacopo Maria D'Andria, Albert Pierres Tejada, and Paul Delatte for excellent research support. E-mail contacts: sulkhan.chavleishvili@ecb.ecb.europa.eu; rengle@stern.nyu.edu; stephan.fahr@ecb.ecb.europa.eu; manfred.kremer@ecb.ecb.europa.eu; simone.manganelli@ecb.ecb.europa.eu; bernd.schwaab@ecb.ecb.europa.eu. The views expressed in this paper are those of the authors and they do not necessarily reflect the views or policies of the European Central Bank.

1 Introduction

The Stoic philosopher Seneca once observed that *“When pleasures have corrupted both mind and body, nothing seems tolerable – not because the suffering is hard, but because the sufferer is soft.”*¹

The quote nicely encapsulates a common view of macro-prudential policy: Make the financial system strong (hard) enough to withstand adverse shocks, taking advantage of good times to build up buffers and increase fortitude. According to this view, the activation of macro-prudential measures is recommended in case medium-term downside risks to the economy are deemed too severe. Such measures, however, are not necessarily without costs to the upside potential, or expected growth, of the economy. This paper proposes an econometric framework based on which macro-prudential authorities can optimally weigh the beneficial impact of their actions on future downside risks with the adverse impact of these actions on the upside potential of the economy.

There is a plethora of notions and techniques to measure downside risk. Yet, the question of how to make them operational for the conduct of macro-prudential policy has received much less attention. [Greenspan \(2003, p. 3\)](#), [Cecchetti \(2006\)](#), and [Kilian and Manganelli \(2008\)](#) propose to view a central bank’s decision as a risk management problem, requiring the central bank to optimally balance downside and upside risks to price stability. This paper extends their idea to the macro-prudential problem, where the relevant authority sets its policy by optimally balancing the inter-temporal trade-off between expected growth and downside risk to the economy. Related methodological approaches have recently been advocated by [Carney \(2020\)](#), [Suarez \(2021\)](#), and, from a general equilibrium perspective, [Mendicino et al. \(2018\)](#) and [Caballero and Simsek \(2020\)](#).

This paper uses a quantile vector autoregressive model (QVAR) to operationalize our risk management idea. QVAR was first proposed in unpublished work by [Cecchetti and Li \(2008\)](#). Independent work by [White et al. \(2015\)](#), [Chavleishvili and Manganelli \(2019\)](#) and [Montes-Rojas \(2019\)](#) has formalized the econometric model. These contributions fit into the broader literature on multivariate quantile regression, which is an active area of research (see, for instance, [Wei \(2008\)](#) or [Carlier et al. \(2016\)](#) and the references therein).

¹Seneca, De Ira, Liber II, XXV, 3.

We show how QVAR provides the ideal econometric tool to address many of the policy issues faced by macro-prudential authorities. It can be used to obtain: estimates of financial stability risks which are easy to communicate, a variable selection procedure for a multivariate model of downside risk, a stress testing assessment of the vulnerability of the financial system, and a measure of macro-prudential policy stance. The main body of the paper discusses the empirical findings for the euro area economy. A companion web appendix shows that the results are qualitatively similar for the U.S. economy.

A rapidly growing body of research has examined downside risk in macroeconomic outcomes. Most of this work has focused on the risk of significant declines in gross domestic product (GDP), brought about by a deterioration of financial conditions. In particular, growth-at-risk (GaR), the, say, 5% quantile of a predictive GDP distribution, has emerged as a popular measure of downside risk; see e.g. [Adrian et al. \(2019\)](#), [Prasad et al. \(2019\)](#), and [Caldara et al. \(2019\)](#). Both the International Monetary Fund (IMF) as well as the European Central Bank (ECB) now routinely publish GaR estimates for major world economies; see [IMF \(2017\)](#) and [ECB \(2019\)](#). These developments have motivated a proliferation of modeling frameworks to assess the severity of extreme events associated with key economic variables, including single-equation quantile regression (QR) models ([Adrian et al. \(2019\)](#)), panel QR models ([Adrian et al. \(2018\)](#), [Brandao-Marques et al. \(2020\)](#)), panel-GARCH models ([Brownlees and Souza \(2020\)](#)), fully non-parametric kernel regression models ([Adrian et al. \(2020\)](#)), combined linear vector autoregressive (VAR) and single-equation QR models ([Duprey and Ueberfeldt \(2020\)](#)), nonlinear Bayesian VAR models ([Caldara et al. \(2019\)](#), [Carriero et al. \(2020\)](#)), and quantile VAR models ([Chavleishvili and Manganelli \(2019\)](#)). [De Santis and van der Veken \(2020\)](#) show that even if a recession is due to an unforeseen real shock (as the recent Covid-19 recession) financial variables can still help policy makers by providing timely warnings about the severity of the crisis and the macroeconomic risks involved. [Plagborg-Moller et al. \(2020\)](#) provide a critical review of this literature.

This paper embeds the insights from this large empirical literature on downside risk into a macro econometric model that can be used to produce structural forecast distributions, taking into

account the strong asymmetries that characterize macro-financial interactions. The QVAR model allows us to quantify future risks to economic activity caused by elevated levels of financial stress as well as by economic vulnerability to shocks. We argue that our framework inherits the best features from both the VAR and QR strands of literature. The VAR permits all endogenous variables to interact over time, allows us to be transparent about the identification of structural shocks, and can be used to simulate from the model and compare different counterfactual policy scenarios. QR allows the dynamic properties of the system to differ across quantiles, capturing potential asymmetries in the propagation of structural shocks. As a welcome by-product, QR parameter estimates are less sensitive to outliers when compared to their least squares counterparts. This robustness feature can become relevant when financial variables are included in the model and the financial system and the economy face abrupt and large changes. Succinctly put, our QVAR model relates to the single-equation QR approach of [Adrian et al. \(2019\)](#) as the VAR model of [Sims \(1980\)](#) relates to the straightforward single-equation autoregressive approaches of e.g. [Koyck \(1954\)](#) and [Almon \(1965\)](#).

To relate to the problem faced by macro-prudential authorities, we include measures of financial stress and medium-term vulnerabilities alongside GDP growth in our baseline model. Indicators of financial stress serve as a macro-finance amplification mechanism that characterise financial crises. Measures of vulnerabilities represent intermediate target variables on which policy makers can act by activating their policy tools. The three-variable setup reflects the consideration that financial stability is of concern to policy makers if it is triggered by an impairment of the financial system and has real economic consequences, e.g. in terms of future employment, consumption, or overall economic activity.² Financial stress is proxied by the ECB's Composite Indicator of Systemic Stress (CISS; see [Hollo et al. \(2012\)](#)), while medium-term vulnerabilities are proxied by [Schüler et al. \(2020\)](#) real-time broad financial cycle indicator. Financial stress, the financial cycle,

²The 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 \(2009\)](#).

and GDP growth can interact freely in our preferred model specification, and can do so to different extents at different quantiles.

The empirical part of this paper applies our statistical model to both euro area and U.S. data. The paper focuses on euro area data between 1988Q3 and 2018Q4,³ and analogous tables and figures based on U.S. data between 1973Q1 and 2018Q4 are discussed in the web appendix. Our main findings are remarkably similar across the euro area and U.S. samples.

We focus on four empirical findings. First, a variable selection exercise suggests that central bank “intermediate target” variables, such as the financial cycle and (de-trended) money market interest rates, interact closely with GDP growth and financial stress across all quantiles. Other variables, such as the term spread may also play a role, particularly for U.S. data, but are not ranked as highly by model selection criteria. We focus on the financial cycle because it can be influenced, at least to some extent, by macro-prudential (and monetary) policy instruments ([Cerutti et al. \(2017\)](#)). Our downside risk estimates are not particularly sensitive to the exclusion of short-term interest rates; we therefore use the more parsimonious trivariate model specification for most of our results.

Second, the dynamic properties of the system differ significantly across quantiles. A formal Wald test rejects the pooling, or parameter homogeneity, restrictions implied by a linear VAR specification for our data at any reasonable confidence level. The QVAR is instead characterized by substantial asymmetries. In particular, a shock to financial stress shifts the left tail of future GDP towards more negative values, while leaving its conditional median and right tail approximately unaffected. As in [Adrian et al. \(2019\)](#), macro-financial interactions imply that the upper quantiles of predictive GDP growth distribution are less volatile than its lower quantiles. Our model-implied downside risk measures are strongly sensitive to the inclusion of financial variables. Not only does downside risk associated with the global financial crisis between 2008 and 2009 decline much later, and to a lesser extent, when financial stress is missing, but in addition the downside risk associated with the 2010–2012 euro area sovereign debt crisis is missed almost entirely when financial stress

³Counterfactual data preceding the formation of the euro area (pre-1999) is obtained from an internal ECB database.

is not included in the model.

Third, we find that the euro area economy is not equally resilient to the same sequence of adverse financial shocks at all times. The asymmetries uncovered in the data suggest that our QVAR model provides a natural environment to perform repeated model-based macro-prudential stress tests for the economy as a whole. We here understand stress testing as a forecast of what would happen to all variables in the system should it be subjected to a fixed sequence of adverse shocks. We find that downside risk conditional on future adverse real and financial shocks spikes during crises. Our model-based stress testing outcomes can be used as a quantitative yardstick to help calibrate the size of macro-prudential capital and liquidity buffers.

Fourth, our risk management framework and its underlying econometric model can be used to guide financial stability policies. In QVAR-based stress tests, the impact of a shock to financial conditions (stress) depends not only on its initial severity, but also on the endogenous, asymmetric responses of all other variables in the system. Allowing for such feedback and asymmetries is crucial when subjecting the system to a sequence of tail shocks. To counteract the feedback and the asymmetries, we argue that macro-prudential policy should act in a counter-cyclical fashion by releasing buffers when downside risk is exceptionally high, and increasing them when downside risk is exceptionally low (Van der Ghote (2021)). Welfare calculations from stabilizing the financial cycle can be based on a suitably chosen objective function. The associated welfare gains can be positive or negative, and are most positive in exuberant times when the financial cycle is above its conditional median. The QVAR estimates therefore provide a metric to assess whether the macro-prudential stance is too tight or too loose. Model-based stress testing and scenario analysis can complement early warning indicators such as the credit-to-GDP-ratio gap that has traditionally informed the calibration of the counter-cyclical capital buffer. Having complementary approaches available may help overcome a potential “inactivity-bias,” according to which few jurisdictions have set their counter-cyclical capital buffers to above-zero levels from the buffer’s inception in 2014 to late 2019.⁴

⁴At the end of 2018, 19 out of 28 European Union countries, and 15 out of 19 euro area countries, had counter-cyclical capital buffers set at zero; see Web Appendix A for details.

We proceed as follows. Section 2 defines our downside risk measures, introduces the risk management framework, and presents the statistical model. Section 3 describes our data. Section 4 applies the model to euro area and U.S. data. Section 5 concludes. A web appendix provides further technical and empirical results.

2 The risk management framework

This section starts by introducing measures of downside risk borrowed from the financial risk management literature. Next, it shortly summarizes the QVAR model and shows how it can be used for forecasting, for semi-parametric risk measurement and for counterfactual analysis. It ends by pulling these elements together into an encompassing risk management framework.

2.1 Measures of downside risk and upside potential

We define three measures of downside risk, which are well-known in the risk management literature. Each measure is of interest in different settings.

2.1.1 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

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

where y_t denotes the quarterly annualized real GDP growth rate between time $t - 1$ and t , and $h = 1, \dots, H$ denotes a certain prediction horizon. The information set \mathcal{F}_{1t} contains all data known at time t ; see Section 2.2 below. 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 .

2.1.2 Growth shortfall

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

$$\begin{aligned} \text{GS}_{t,t+h}^\tau &= \int_{-\infty}^{\tau} y_{t+h} dF_{t,t+h}(y_{t+h}) \\ &= \mathbb{E} [y_{t+h} | y_{t+h} < \tau, \mathcal{F}_{1t}] \times \mathbb{P} [y_{t+h} < \tau | \mathcal{F}_{1t}], \end{aligned} \quad (2)$$

where $F_{t,t+h}$ is a time- t conditional cumulative distribution function (cdf), $\mathbb{E} [\cdot | \mathcal{F}_{1t}]$ denotes a time- t conditional expectation, and 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) coincides with the familiar notion of expected shortfall; see e.g. McNeil et al. (2005, Ch. 2). Alternatively, it could be set to a certain unconditional quantile, or be set to zero.

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.⁵ While both components can be studied separately and can be of interest in their own right, such as in stress test or macroeconomic modeling, GS summarizes them tractably into one metric and is easily obtained from a QVAR model. When $\tau = 0$, GS corresponds to the economic question: what is the time t -expected contraction of the economy at time $t + h$?

2.1.3 Average growth shortfall

Our final measure of adverse real economic impact is the *average future growth shortfall* (AGS) between $t + 1$ and $t + H$, defined as

$$\text{AGS}_{t,t+1:t+H}^\tau = H^{-1} \sum_{h=1}^H \text{GS}_{t,t+h}^\tau. \quad (3)$$

If $\tau = 0$, then the AGS corresponds to the question: what is the average future expected contraction of the economy between $t + 1$ and $t + H$. Since it is an average of future GS, AGS retains all the

⁵To see this, note that $\mathbb{E} [y_{t+h} | y_{t+h} < \tau, \mathcal{F}_{1t}] \equiv \frac{\int_{-\infty}^{\tau} y_{t+h} \cdot 1_{\{y_{t+h} < \tau\}} dF_{t,t+h}(y_{t+h})}{\int_{-\infty}^{\tau} 1_{\{y_{t+h} < \tau\}} dF_{t,t+h}(y_{t+h})} = \frac{\int_{-\infty}^{\tau} y_{t+h} dF_{t,t+h}(y_{t+h})}{\mathbb{P}[y_{t+h} < \tau | \mathcal{F}_{1t}]}$.

statistical properties of GS.

All above risk measures are economically intuitive and straightforward to communicate. Risk measures (2) and (3), however, have theoretical and practical advantages over (1). First, expected shortfall-based measures are coherent risk measures, while any single 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, wants to study sector contributions to aggregate GDP at risk. Second, while all above risk measures (1) – (3) can take into account the asymmetric impact of financial variables on the economy, only (2) and (3) take into account the entire left tail.

When considering financial stability policies aimed at containing downside risk, then the expected growth rate of the economy, as well as the upper quantiles of future GDP growth, should not be unduly affected. For setting up the risk management framework later in section 2.3, we consider two measures of upside potential that are symmetric to the measures of downside risk just defined.

2.1.4 Growth longrise

We define the *growth longrise*⁶ (GL) as the complement to GS,

$$\begin{aligned} \text{GL}_{t,t+h}^\tau &= \int_\tau^\infty y_{t+h} dF_{t,t+h}(y_{t+h}) \\ &= \mathbb{E}[y_{t+h} | y_{t+h} > \tau, \mathcal{F}_{1t}] \times \mathbb{P}[y_{t+h} > \tau | \mathcal{F}_{1t}]. \end{aligned} \quad (4)$$

If $\tau = 0$, then (4) corresponds to the question: what is the time- t expected expansion of the economy between $t + h - 1$ and $t + h$? Similarly to GS, the growth longrise (4) captures the expected growth given an expansion, and the conditional probability of experiencing an expansion.

⁶The term longrise was coined by Adrian et al. (2019) as the antonym to shortfall.

2.1.5 Average growth longrise

Analogously to (3), we also define the *average growth longrise* (AGL) between $t + 1$ and $t + H$ as

$$\text{AGL}_{t,t+1:t+H}^\tau = H^{-1} \sum_{h=1}^H \text{GL}_{t,t+h}^\tau. \quad (5)$$

Given the complementarity between GS and GL, their sum equals the expected growth rate of the economy between $t + h - 1$ and $t + h$,

$$\begin{aligned} \mathbb{E}[y_{t+h} | \mathcal{F}_{1t}] &= \int_{-\infty}^{\infty} y_{t+h} dF_{t,t+h}(y_{t+h}) \\ &= \int_{-\infty}^{\tau} y_{t+h} dF_{t,t+h}(y_{t+h}) + \int_{\tau}^{\infty} y_{t+h} dF_{t,t+h}(y_{t+h}) \\ &= \text{GS}_{t,t+h}^\tau + \text{GL}_{t,t+h}^\tau. \end{aligned}$$

Furthermore, let $\bar{y}_{t,t+1:t+H} = H^{-1} \sum_{h=1}^H y_{t,t+h}$ be the average future economic growth rate between $t + 1$ and $t + H$. Since (2) and (4) are linear, the expected future growth rate of the economy between $t + 1$ and $t + H$ is $\mathbb{E}[\bar{y}_{t,t+1:t+H} | \mathcal{F}_{1t}] = \text{AGS}_{t,t+1:t+H}^\tau + \text{AGL}_{t,t+1:t+H}^\tau$. As a result, expected average future growth can be read off any figure reporting $\text{AGS}_{t,t+1:t+H}^\tau$ and $\text{AGL}_{t,t+1:t+H}^\tau$ by adding the two lines.

2.2 Quantile vector autoregression

This section provides a concise exposition of the structural quantile vector autoregressive (QVAR) model of [Chavleishvili and Manganelli \(2019\)](#) and shows how it can be used to obtain semi-parametric estimates of (1) – (5).⁷ The structural identification is obtained by imposing triangular restrictions and can be thought of as a macro-econometric application of the approach proposed by [Wei \(2008\)](#). Intuitively, the QVAR model provides the forecast of the quantiles of the distribution of the endogenous variables at any period ahead. The quantile forecasts can be treated as the

⁷See also [Montes-Rojas \(2019\)](#) and [Ruzicka \(2020\)](#) for related work on QVAR parameter and quantile impulse response function estimation.

equivalent of an empirical distribution and can therefore be used to approximate the risk quantities discussed in the previous section. Furthermore, since we attach a structural interpretation to the model, the structural shocks can be recovered and used to perform counterfactual exercises, as one would do with a standard VAR.

2.2.1 The model

We observe a series of random variables $\{\tilde{x}_t : t = 1, \dots, T\}$, where $\tilde{x}_t \in \mathbb{R}^n$ is an n -vector with i^{th} element denoted by \tilde{x}_{it} for $i = 1, \dots, n$ and $n \in \mathbb{N}$. For any arbitrary but fixed quantile γ , the QVAR model of order 1 is given by

$$\tilde{x}_{t+1} = \omega^\gamma + A_0^\gamma \tilde{x}_{t+1} + A_1^\gamma \tilde{x}_t + \epsilon_{t+1}^\gamma \quad (6)$$

$$\mathbb{P}(\epsilon_{i,t+1}^\gamma < 0 | \mathcal{F}_{it}) = \gamma, \quad \text{for } i = 1, \dots, n, \quad (7)$$

where the vector of structural quantile residuals is given by $\epsilon_t^\gamma \equiv [\epsilon_{1t}^\gamma, \dots, \epsilon_{nt}^\gamma]'$. Recursive identification is achieved by restricting the $[n \times n]$ matrix A_0^γ to be lower triangular with zeros along the main diagonal. The presence of contemporaneous dependent variables on the right-hand side of (6) requires us to be precise about the available information at any time and for each variable. We work with a recursive information set that increases one scalar observation at a time,

$$\mathcal{F}_{1t} = \{\tilde{x}_t, \tilde{x}_{t-1}, \dots\} \quad (8)$$

$$\mathcal{F}_{it} = \{\tilde{x}_{i-1,t+1}, \mathcal{F}_{i-1,t}\} \text{ for } i \in \{2, \dots, n\}. \quad (9)$$

In words, \mathcal{F}_{1t} contains only variables observed up to time t . The information sets \mathcal{F}_{it} for $i > 1$ contain increasingly more information about variables observed at $t + 1$.⁸

We may wish to consider multiple quantiles of multiple variables at the same time. To do this in a compact way, we consider p distinct quantiles $0 < \gamma_1 < \dots < \gamma_p < 1$, for $p \in \mathbb{N}$, not necessarily equidistant. In addition, we let $x_t \equiv [\iota_p \otimes \tilde{x}_t]$ denote the vector stacking p times the dependent

⁸For a similar incremental conditioning approach in a different setting see e.g. [Koopman and Durbin \(2000\)](#).

variables \tilde{x}_t , where ι_p is a p -vector of ones. The stacked QVAR model of order 1 is then given by

$$x_{t+1} = \omega + A_0 x_{t+1} + A_1 x_t + \epsilon_{t+1} \quad (10)$$

$$\mathbb{P}(\epsilon_{i,t+1}^{\gamma_j} < 0 | \mathcal{F}_{it}) = \gamma_j, \quad \text{for } i = 1, \dots, n, \quad j = 1, \dots, p \quad (11)$$

where the vector of structural quantile residuals is given by $\epsilon_t \equiv [\epsilon_{1t}^{\gamma_1}, \dots, \epsilon_{nt}^{\gamma_1}, \dots, \epsilon_{1t}^{\gamma_p}, \dots, \epsilon_{nt}^{\gamma_p}]'$.

The $[np \times np]$ matrices A_0 and A_1 are block diagonal to avoid trivial multicollinearity problems.

The model (10) – (11) is essentially a convenient way to stack p quantile-specific QVAR models (6) – (7).

An explicit example may be instructive. While the baseline empirical model in Section 4 considers three variables, we here develop intuition based on a simpler bivariate model for the data vector $\tilde{x}_t = (y_t, s_t)'$, where y_t is the quarterly annualized real GDP growth between $t - 1$ and t , and s_t is a coincident indicator of systemic financial stress. Let us consider $p = 2$ quantiles for simplicity, 0.10 and 0.90. The system (10) – (11) can then be written as

$$\begin{aligned} \begin{bmatrix} y_{t+1} \\ s_{t+1} \\ y_{t+1} \\ s_{t+1} \end{bmatrix} &= \begin{bmatrix} \omega_y^1 \\ \omega_s^1 \\ \omega_y^9 \\ \omega_s^9 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ a_{021}^1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & a_{021}^9 & 0 \end{bmatrix} \begin{bmatrix} y_{t+1} \\ s_{t+1} \\ y_{t+1} \\ s_{t+1} \end{bmatrix} \\ &+ \begin{bmatrix} a_{11}^1 & a_{12}^1 & 0 & 0 \\ a_{21}^1 & a_{22}^1 & 0 & 0 \\ 0 & 0 & a_{11}^9 & a_{12}^9 \\ 0 & 0 & a_{21}^9 & a_{22}^9 \end{bmatrix} \begin{bmatrix} y_t \\ s_t \\ y_t \\ s_t \end{bmatrix} + \begin{bmatrix} \epsilon_{y,t+1}^1 \\ \epsilon_{s,t+1}^1 \\ \epsilon_{y,t+1}^9 \\ \epsilon_{s,t+1}^9 \end{bmatrix} \quad (12) \end{aligned}$$

Here, the ordering of the observations in (12) reflects the assumption that the financial stress variable s_t can react contemporaneously to macroeconomic shocks, while real output growth y_t can react to financial shocks only with a lag. Such triangular identification assumptions are standard in

the empirical literature; see e.g. [Christiano et al. \(1999\)](#), [Kilian \(2009\)](#), and [Gilchrist and Zakrajsek \(2012\)](#), among many others.

2.2.2 Forecasting

This section explains how forecasts can be generated from the stacked QVAR model (10) – (11) without invoking parametric assumptions on ϵ_{t+1} .

It is helpful to introduce the conditional quantile operator $Q_{it}^{\gamma_j}(x_{k,t+1})$, where $x_{k,t+1}$ is the k -th element of x_{t+1} , $k = 1, \dots, np$. Given information set \mathcal{F}_{it} , the operator is implicitly defined by

$$\mathbb{P}(x_{k,t+1} < Q_{it}^{\gamma_j}(x_{k,t+1}) | \mathcal{F}_{it}) = \gamma_j, \quad \text{for } j = 1, \dots, p.$$

In words, $Q_{it}^{\gamma_j}(x_{k,t+1})$ returns the γ_j quantile of random variable $x_{k,t+1}$ conditional on \mathcal{F}_{it} . The element $x_{k,t+1}$ is random because it depends on its own shock at time $t + 1$, but also on shocks to earlier elements $x_{1,t+1}, \dots, x_{k-1,t+1}$.

To build intuition first, let us return to the simple bivariate example (12) with $n = p = 2$. Let's assume we are interested in forecasting, say, the 0.9 quantile of the financial stress variable s_{t+1} . The fourth equation of (12), corresponding to the 0.9 quantile of s_{t+1} , is

$$\begin{aligned} s_{t+1} &= \omega_s^{\cdot 9} + a_{021}^{\cdot 9}[\omega_y^{\cdot 9} + a_{11}^{\cdot 9}y_t + a_{12}^{\cdot 9}s_t + \epsilon_{y,t+1}^{\cdot 9}] + a_{21}^{\cdot 9}y_t + a_{22}^{\cdot 9}s_t + \epsilon_{s,t+1}^{\cdot 9} \\ &= \omega_s^{\cdot 9} + a_{021}^{\cdot 9}\omega_y^{\cdot 9} + (a_{021}^{\cdot 9}a_{11}^{\cdot 9} + a_{21}^{\cdot 9})y_t + (a_{021}^{\cdot 9}a_{12}^{\cdot 9} + a_{22}^{\cdot 9})s_t + a_{021}^{\cdot 9}\epsilon_{y,t+1}^{\cdot 9} + \epsilon_{s,t+1}^{\cdot 9} \\ &= q_{st}^{\cdot 9} + a_{021}^{\cdot 9}\epsilon_{y,t+1}^{\cdot 9} + \epsilon_{s,t+1}^{\cdot 9} \end{aligned} \tag{13}$$

where $q_{st}^{\cdot 9} \equiv \omega_s^{\cdot 9} + a_{021}^{\cdot 9}\omega_y^{\cdot 9} + (a_{021}^{\cdot 9}a_{11}^{\cdot 9} + a_{21}^{\cdot 9})y_t + (a_{021}^{\cdot 9}a_{12}^{\cdot 9} + a_{22}^{\cdot 9})s_t$ depends only on deterministic parameters to be estimated and variables observed at time t . We note that $Q_{st}^{\cdot 9}(\epsilon_{s,t+1}^{\cdot 9}) = 0$ because of the identifying restriction (11), stating $\mathbb{P}(\epsilon_{s,t+1}^{\cdot 9} < 0 | \mathcal{F}_{st}) = 0.9$ when $\mathcal{F}_{st} = \{y_{t+1}, y_t, s_t, \dots\}$. In addition, $q_{st}^{\cdot 9} + a_{021}^{\cdot 9}\epsilon_{y,t+1}^{\cdot 9} | \mathcal{F}_{st}$ is non-random. As a result, $Q_{st}^{\cdot 9}(s_{t+1}) = q_{st}^{\cdot 9} + a_{021}^{\cdot 9}\epsilon_{y,t+1}^{\cdot 9}$ is still a random variable at time t . To eliminate this randomness, we keep on taking quantiles. Using the identifying restriction (11) again, $Q_{yt}^{\cdot 9}(\epsilon_{y,t+1}^{\cdot 9}) = 0$ yields $Q_{yt}^{\cdot 9}(Q_{st}^{\cdot 9}(s_{t+1})) = q_{st}^{\cdot 9}$. As a result, $q_{st}^{\cdot 9}$

is our sought-after forecast of the 0.9 quantile of s_{t+1} , and is easily computed. This approach of iterated quantiles can be repeated for any potentially remaining variables in x_{t+1} . Following that, the approach can be repeated for future variables in x_{t+h} for $h > 1$.⁹

The above reasoning can be formalized. The scalar operators $Q_{it}^{\gamma_j}(x_{k,t+1})$ can be combined into a vector version, with quantile operators nesting each other up to n times. The vector operators can again be sequentially combined, up to H times. In the end, the $[np \times 1]$ -vector of quantile forecasts at time t associated with process (10), for $h = 1, \dots, H$, can be obtained quite straightforwardly as

$$\hat{x}_{t+h} = \sum_{j=0}^{h-1} B^j \nu + B^h x_t, \quad (14)$$

where $\nu = (I_{np} - A_0)^{-1} \omega$ and $B = (I_{np} - A_0)^{-1} A_1$. It is easily verified that $\hat{x}_{4,t+1}$ (i.e., the fourth element of \hat{x}_{t+1} , obtained using (14)) coincides with q_{st}^9 as defined below (13).

2.2.3 Semi-parametric risk measurement

This section explains how we obtain the time- t downside risk measures introduced in Section 2.1 from our semi-parametric structural QVAR model (10) – (11) using simulation methods. To this end we rely on a growing literature on simulation methods for quantile regression; see e.g. [Hahn \(1995\)](#) and [Koenker \(2005, Ch. 2.6\)](#).

When we defined the structural QVAR model for an arbitrary quantile γ as (6) – (7), and insisted that the model holds for all $\gamma \in (0, 1)$, we effectively specified a complete stochastic mechanism for generating the one-step ahead variable \tilde{x}_{t+1} conditional on time- t information and deterministic parameters. Recall that any scalar response variable $\tilde{x}_{i,t+1}$, $i = 1, \dots, n$, with conditional cdf $F_{i,t,t+1}$, can be simulated by generating a uniform random variable $u_{i,t+1} \sim U[0, 1]$, and then setting $\tilde{x}_{i,t+1} = F_{i,t,t+1}^{-1}(u_{i,t+1})$. Thus, in model (6) – (7), $\tilde{x}_{i,t+1}$ can be simulated setting $\tilde{x}_{i,t+1} = \omega_i^{u_{i,t+1}} + A_{0,i}^{u_{i,t+1}} \tilde{x}_{t+1} + A_{1,i}^{u_{i,t+1}} \tilde{x}_t$, where $\omega_i^{(\cdot)}$, $A_{0,i}^{(\cdot)}$, and $A_{1,i}^{(\cdot)}$ denote the i -th row of ω^γ , A_0^γ ,

⁹ This example implicitly assumes that a_{021}^9 is positive. If not, then the 0.9 conditional quantile and the 0.1 conditional quantile cross. This is because, when $a_{021}^9 < 0$, then $\mathbb{P}(a_{021}^9 \epsilon_{y,t+1}^9 < 0) = \mathbb{P}(\epsilon_{y,t+1}^9 > 0) = 1 - \mathbb{P}(\epsilon_{y,t+1}^9 < 0) = 1 - 0.9 = 0.1$. If this happens we reorder (relabel) the quantiles accordingly.

and A_1^γ , respectively, evaluated at $\gamma = u_{i,t+1}$.¹⁰ This procedure allows us to generate the $\tilde{x}_{i,t+h}$, recursively, for $i = 1, \dots, n$ and $h = 1, \dots, H$, conditional on the relevant information sets.

We sketch our simulation algorithm here, and refer to Web Appendix B.4 for details. Let $t = 1, \dots, T$ denote any time in our sample. With n variables, p quantiles and H steps ahead, there are p^{nH} possible paths, a number which quickly becomes computationally unmanageable. We resort instead to simulation, by randomly generating S potential future paths for all n variables in $\tilde{x}_{i,t+h}$, $h = 1, \dots, H$ quarters ahead.¹¹ The simulations are based on inverse cdf-sampling by drawing $S = 10,000$ sequences of $nH = 3 \times 8 = 24$ uniform random variables with support $[\gamma_1, \dots, \gamma_p]$, and use the one-step-ahead recursion (6). At each $t + h$, we calculate $\text{GS}_{t,t+h}^\tau$ and $\text{GL}_{t,t+h}^\tau$ by evaluating the sample analogues of (2) and (4). At the end, we average across H to obtain downside risk measures $\text{AGS}_{t,t+1:t+h}^\tau$ and $\text{AGL}_{t,t+1:t+h}^\tau$; see (3) and (5).¹²

2.2.4 Counterfactual scenarios

Web Appendix B.5 explains in detail how counterfactual scenarios can be obtained from the QVAR model (10) – (11). Rather than moving through the complete tree of potential future values of \tilde{x}_{t+h} at random, as explained in Section 2.2.3, we then consider only one path in isolation. Such a path can be thought of as a ‘counterfactual scenario,’ or model-based thought experiment that conditions on an arbitrary but fixed sequence of future shocks. We use such counterfactual scenarios when considering a market-based stress test in Section 4.3, and when studying the benefits vs. cost from tightening macro-prudential policy stance in Sections 4.4.

¹⁰Recall that A_0 is lower triangular.

¹¹The simulation approach is no panacea, as S still needs to be chosen large enough to sufficiently explore the tree. Our risk estimates presented in Section 4 are insensitive to the initial random seed and to variations in the number of simulations.

¹²Rather than re-estimating the model parameters within each simulation and for each variable using $\gamma_i = u_{i,t+h}$, it is computationally advantageous to discretize the support of the standard uniform random variable with an appropriately chosen grid $0 < \gamma_1 < \dots < \gamma_p < 1$, and to estimate all parameters once and for all in the beginning based on the full sample. We then use the parameter estimates associated with the closest selected quantile in any simulation. We use $p = 20$ grid-points for this purpose, $0 < 0.025, 0.075, \dots, 0.925, 0.975 < 1$, each at the midpoint of 1/20th of the unit interval. These grid-points are symmetric around the median, and yield equi-probable simulation paths. Crossing quantiles (see footnote 9) are not an issue since we move through the tree at random. Our downside risk estimates reported in Section 4.2 are robust to increasing the number of grid-points, and to interpolating parameter estimates between quantiles.

2.3 Putting it all together in a risk management framework

A key question for a policy maker is to what extent a policy intervention reduces downside risk to the economy and what risk it imposes in terms of reduced growth. In other words, how is a policy maker to assess the change in forecast distributions triggered by its actions? The risk management framework answers this question by requiring the policy maker to be explicit about the trade-off between downside risk and upside potential. It is equivalent to requiring the decision maker to provide a loss function, and setting the policy variables to levels that minimize such a loss. We believe that framing the problem in terms of risk management facilitates the elicitation of the preferences of the policy maker and the communication of the policy decision.

Suppose the macro-prudential authority has an instrument (or vector of instruments) c_t that can be used to influence the predictive growth distribution. This influence can be direct ($c_t \rightarrow y_{t+1}$) or indirect (e.g., $c_t \rightarrow s_t \rightarrow y_{t+1}$). The QVAR structure allows us to capture both types of transmission. A convenient way to penalize downside risk is given by specifying the utility maximization problem as

$$\max_{\{c_{t+h}\}_{h=1}^{\infty}} \sum_{h=1}^{\infty} \beta^h \left(GL_{t,t+h}(y_{t+h}(c_{t:t+h})) + \lambda GS_{t,t+h}(y_{t+h}(c_{t:t+h})) \right) \quad (15)$$

where $\lambda > 1$ is a weight determining the aversion to negative realisations of output growth, β is an intertemporal discount factor, $c_{t:t+h} = (c_t, \dots, c_{t+h})'$, and $GS_{t,t+h}$ is always a negative number.

The objective function (15) is reminiscent of the mean with downside risk model in asset allocation; see e.g. [Fishburn \(1977\)](#).¹³ Since $\mathbb{E}[y_{t+h} | \mathcal{F}_{1t}] = GS_{t,t+h}^{\tau} + GL_{t,t+h}^{\tau}$, see Section 2.1, (15) can be rewritten in terms of expected future economic growth instead of upper quantiles to future growth. The objective function (15) is then equal to the expression suggested by [Carney \(2020\)](#),

$$\max_{\{c_{t+h}\}_{h=1}^{\infty}} \sum_{h=1}^{\infty} \beta^h \left(\mathbb{E}_t(y_{t+h}(c_{t:t+h})) + (\lambda - 1)GS_{t,t+h}(y_{t+h}(c_{t:t+h})) \right), \quad (16)$$

¹³[Kilian and Manganelli \(2008\)](#) show that most of the currently existing downside risk measures are special cases of the downside risk notion proposed by [Fishburn \(1977\)](#) in the context of portfolio allocation.

trading off future trend growth against downside risks to the economy. We refer to [Suarez \(2021\)](#) for a micro-foundation of a similar objective function based on a representative agent with a CARA utility function on GDP.¹⁴ We use (16) to study the benefits from adopting an active financial stability policy in Section 4.4 below.

3 Data

3.1 Macroeconomic data pre-1999

Structural QVAR models require a sufficiently large sample size to ensure that its parameters can be estimated with adequate precision. At least two challenges are present, however, when working with euro area macro data in practice. First, the euro area celebrated its 20th anniversary merely in 2019. When working with quarterly data, $T = 4 \times 20 = 80$ is at the lower end of what is required for a meaningful empirical study of macro-financial interactions at different quantiles. Second, euro area membership has been expanding over time, from initially 11 countries in 1999 to 19 countries in 2015. Changes in euro area aggregate data stemming from new countries joining, rather than, say, from changes in financial conditions or growing vulnerabilities, would severely complicate any empirical analysis.

Fortunately, both problems can be addressed. During the ECB's early years, pre-1999 macro-financial time series data were urgently needed for monetary policy analysis. Against this background counterfactual data were constructed "as if" the euro area had already consisted earlier; see e.g. [Fagan et al. \(2001\)](#). Such pre-1999 euro area data is publicly available.¹⁵ We obtain real GDP growth data from 1988Q3 to 2018Q4 from this source, resulting in $T = 121$, and refer to [Web Appendix C.1](#) for a time series plot.

¹⁴Our empirical results presented in Sections 4.1 to 4.3 do not depend on the objective function. In Section 4.4 any other objective function could be used, if so desired, including complicated nonlinear specifications.

¹⁵<https://eabcn.org/page/area-wide-model>. In its most recent version, the database adopts a fixed euro area composition approach, constructing aggregate data series as if the euro area had always consisted of its current (end-of-sample) 19 members. Most variables are available from 1970Q1 onwards. The further back, however, the more uncertain the data quality.

3.2 Composite indicator of systemic stress

The ECB's composite indicator of systemic stress (CISS) is a summary measure of the level of financial distress. Web Appendix C.1 provides a summary of the methodology, a time series plot, and a listing of all included data series. The CISS is computed for the euro area as a whole and includes 15 raw, market-based financial indicators that are split equally into five categories: financial intermediaries, money markets, equity markets, bond markets, and foreign exchange markets. Each category is summarized by a sub-index. The sub-indices are subsequently aggregated to a single time series in a way that takes their time-varying cross-correlations into account. As a result, the CISS takes higher values when stress prevails in several 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 segments of the financial system. The CISS is updated regularly and publicly available.¹⁶

3.3 Real-time estimates of the financial cycle

The real-time financial cycle indicator used in the empirical analysis is based on Schüller et al. (2020). The construction of the indicator mirrors that of the CISS; see Web Appendix C.2 for details and a time series plot. Their indicator takes high values when *i*) total non-financial credit volumes grow at an unusually fast pace (proxying a credit boom), and *ii*) real estate, equity, and bond prices grow at an unusually fast pace as well *at the same time* (proxying asset price inflation). In this sense, their financial cycle indicator is not a measurement of credit growth, which can be beneficial, but of *bad*, or excess, credit growth that coincides with asset price inflation.

The financial cycle indicator is available for the euro area and the U.S. from the authors. Their indicator took 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. Their indicator took particularly low values in 2009 and 2011, times associated with crisis-induced fire sales and financial system deleveraging.

¹⁶<https://sdw.ecb.europa.eu/>

4 Implementing the macro-prudential risk management approach

This section uses an estimated QVAR model to study macro-prudential policy stance in the presence of substantial asymmetries and tail interactions. It first discusses model selection, parameter estimates for a baseline specification, and the outcome of specification tests. Second, it quantifies downside risks to, and the upside potential of, the euro area economy stemming from financial stress and vulnerabilities. Third, it reports the outcome of a model stress testing exercise, assessing whether the euro area economy was at all times equally vulnerable to a fixed sequence of adverse shocks. Finally, it asks whether it pays off to adopt an active macro-prudential policy, producing a risk management-based metric of macro-prudential policy stance. We focus our discussion on the euro area, and report analogous tables and figures for U.S. data in Web Appendix F.

4.1 QVAR estimates

In this section, we report the estimation results of our favorite QVAR model specification. We discuss the variable selection procedure, the characteristics of the parameters estimation, specification tests, and quantile impulse response functions.

4.1.1 Variable selection exercise

A two-variable QVAR model for quarterly real GDP growth and the CISS provides a minimal system to study downside risks to the real economy. GDP growth is required to quantify downside risks, and the CISS significantly impacts the left tail of the predictive GDP growth distribution; see Section 4.1.2 below. This minimal system, however, may miss important interactions with other economic variables. In addition, it misses a variable that can be influenced directly through financial stability policies.

Web Appendix D discusses a systematic search over potential additional endogenous variables to be included in a QVAR. Two variables stand out as interacting closely with euro area GDP growth and financial stress at all nine quantiles. Both are related to central bank policy instru-

ments. The de-trended three-months EURIBOR rate, a measure of monetary policy, is ranked first, impacting both future GDP growth as well as current financial conditions. [Schüler et al. \(2020\)](#) broad financial cycle indicator (see Section 3.3) is ranked second, followed by the euro area’s capacity utilization rate. Capacity utilization is a business cycle indicator, and as such highly correlated with GDP growth, and arguably of lesser interest in a financial stability context.

Web Appendix F.2 reports analogous results for U.S. data. Approximately similar variables are selected.

4.1.2 Model specification and parameter estimates

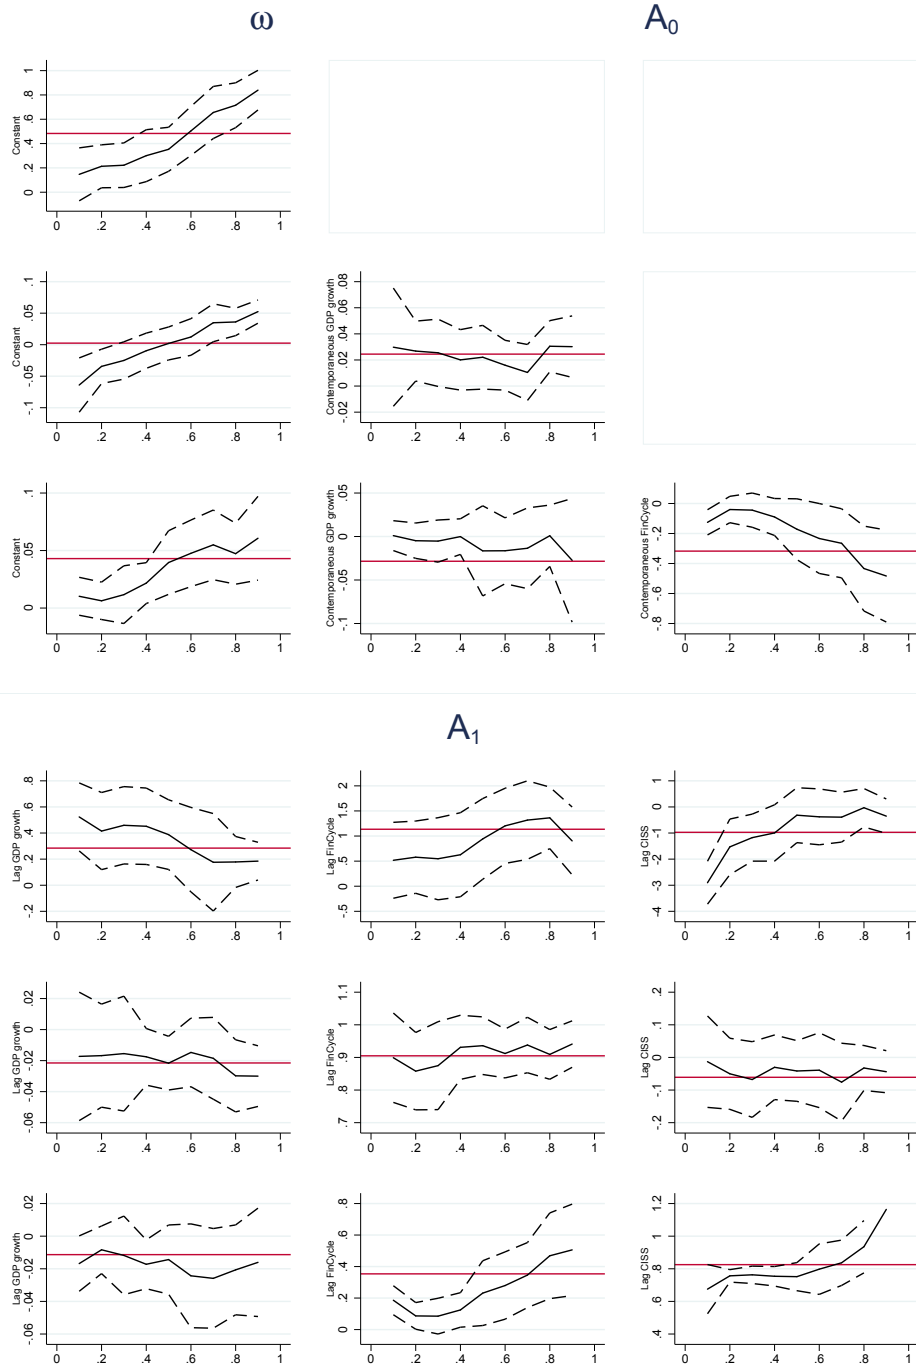
We choose a trivariate QVAR specification as our benchmark model. Our benchmark model consists of annualized quarterly real GDP growth y_t , the financial cycle indicator c_t , and the CISS s_t . We therefore consider $\tilde{x}_t = (y_t, c_t, s_t)'$.

Figure 1 reports parameter and standard error estimates for our baseline specification. Parameter point estimates are obtained equation-by-equation via np univariate quantile regressions. The appropriate standard error bands around the parameter point estimates, however, do not coincide with the equation-by-equation estimates as supplied by common software packages. The standard errors reported in Figure 1 take cross-equation restrictions at common quantiles into account, see Web Appendix B.1 for details, and can be tighter or wider compared to the equation-by-equation standard error estimates.

We discuss the parameter estimates from top left to bottom right. Each of the panels presents the parameter estimates across nine deciles together with 95% confidence bands and the corresponding least squares estimate. The arrangement of panels in Figure 1 corresponds to the ordering of variables in (6). Overall, the quantile regression estimates differ substantially across quantiles, as well as from their least squares counterparts. Each intercept estimate in ω increases monotonically in the considered quantile. This pattern is by construction, and reflects the fact that quantile shocks are not centered around zero; see (11).

Figure 1: Parameter estimates for baseline QVAR model

Parameter estimates from a trivariate QVAR model estimated for $p = 9$ quantiles from 0.1 to 0.9. Variables are ordered GDP growth (respective first row), financial cycle (second row), and CISS (third row). Parameter estimates are obtained equation-by-equation while standard error estimates take cross-equation restrictions into account; see Web Appendix B.1. Standard error bands are dashed and at a 95% confidence level. Red horizontal lines indicate least squares estimates. Estimation sample is 1988Q3 to 2018Q4.



All contemporaneous effects are visible from matrix A_0 . The contemporaneous impact of GDP growth on the financial cycle (element [2,1]) as well as on the CISS (element [3,1]) is small and rarely statistically significant. The [3,2]-element of A_0 points to a positive contemporaneous impact of the financial cycle on the CISS at its lower quantiles. This element is a mirror image of the [3,2]-element in A_1 . Taken together, they suggest that the CISS is high when the financial cycle falls (or vice versa), a pattern that also shows up in the respective impulse response function shown in Figure 2. This is intuitive, as financial sector deleveraging and financial stress tend to go hand-in-hand.

All lagged effects are visible from matrix A_1 . The [1,3]-element signals the presence of substantial asymmetries in the impact that financial stress (CISS) has on future GDP growth. The [3,3]-element of A_1 captures the autoregressive coefficient associated with the CISS. The estimate exceeds one at the 0.9 quantile, pointing to a local non-stationarity in the rightmost tail. Local non-stationarity is not uncommon in QAR models, and does not imply global non-stationarity; see [Koenker \(2005, Ch. 8.3\)](#). Indeed, conditional quantiles simulated from our QVAR model at estimated parameters converge to their unconditional counterparts. The standard errors around the locally non-stationary estimate are, however, not reliable, and not reported for this reason.

Three specifications have been run for robustness. First, the variable selection exercise in Section 4.1.1 suggested that short-term interbank rates can be a useful additional variable to consider in a QVAR. Web Appendix E.1 studies a five-variable monetary structural QVAR model. This model additionally contains the three-month EURIBOR rate as well as quarterly changes in the GDP deflator (inflation). This monetary structural QVAR model is of considerable interest in its own right. It yields, however, broadly similar predictions in terms of downside risks and measures of macro-prudential policy stance. We therefore proceed with the above more parsimonious trivariate model for simplicity.

Second, Web Appendix E.2 extends our baseline model with an additional, annual lag for all variables. Information criteria prefer the more parsimonious version. The average future growth shortfall responds more quickly, and more severely, to contemporaneous financial stress when

Table 1: Wald test of parameter homogeneity.

Wald tests statistics. The test's null hypothesis states that the quantile regression estimates, across $p = 9$ quantiles, are equal to the median regression parameter estimates. We consider our baseline trivariate QVAR model, estimated decile-by-decile, ranging from 0.1 to 0.9; see Figure 1. The test statistic is χ^2 -distributed. The appropriate degrees-of-freedom (df) are given by the number of right-hand-side variables per equation (excluding the constant, 3, 4, and 5, respectively), times the number of imposed restrictions ($9 - 1 = 8$).

	df	test statistic	p-value
real GDP growth, y_t	24	209.71	0.00
financial cycle indicator, c_t	32	26.12	0.76
CISS Financial stress index, s_t	40	79.52	0.00

based on a single-lag specification. We therefore proceed with the single-lag specification.

Finally, Web Appendix E.3 presents our baseline QVAR parameter estimates when the estimation sample is restricted to exclude counterfactual pre-1999 euro area data. The point estimates are more noisy but overall similar. The standard error bands are wider, suggesting less precise parameter estimates.

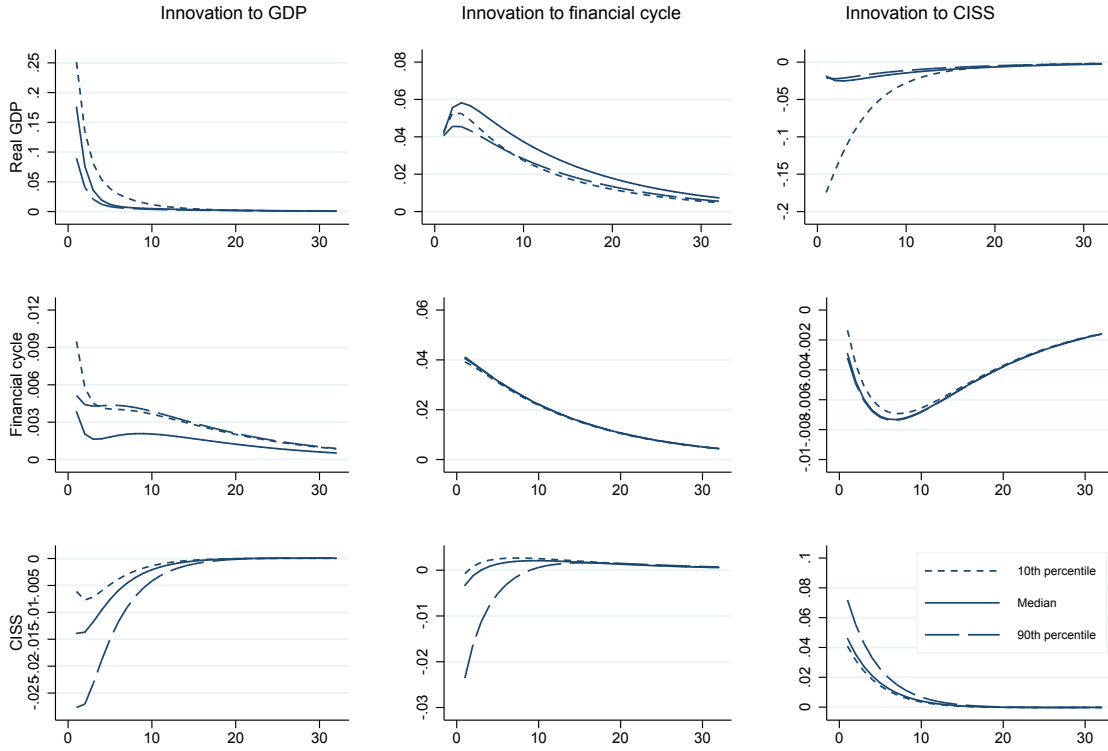
Web Appendix F.3 reports parameter and standard error estimates based on our baseline QVAR model for U.S. data. The parameter estimates are broadly in line with those for the euro area: Growing financial vulnerabilities shift the right tail of the U.S. CISS towards more positive values. A shock to the U.S. CISS shifts the left tail of the predictive GDP growth distribution towards more negative values, while leaving the right tail less affected.

4.1.3 Wald test and quantile impulse response functions

Table 1 reports the outcome of three Wald χ^2 tests of parameter homogeneity across quantiles. We proceed equation by equation for $i = 1, 2, 3$. Each Wald test is implemented as explained in Koenker (2005, Ch. 3.3.2); see also Koenker and Basset (1982) and Web Appendix B.2. The test rejects the parameter equality restrictions implied by a linear VAR for two of our three variables, GDP growth and CISS. Parameter homogeneity is most forcefully rejected for the GDP growth equation. The test outcomes are intuitive given the parameter and standard error estimates reported

Figure 2: Quantile impulse response functions

Impulse response functions implied by the parameter estimates reported in Figure 1. Variables are ordered as GDP growth (respective first row), financial cycle (second row), and CISS (third row). Estimation sample is 1988Q3 to 2018Q4.



in Figure 1.

Figure 2 plots quantile impulse response functions (QIRF) as implied by the parameter estimates in Figure 1. We refer to Web Appendix B.3 for a precise definition and derivation of QIRF in our modeling context. We define the QIRF as the change in the conditional quantile forecast q_{t+h}^γ , at any $\gamma \in (0, 1)$, when a one standard deviation (of $\epsilon_{it}^{\cdot 5}$) sized shock is applied to the structural shocks ϵ_{it}^γ . In a standard VAR, there is only the mean forecast to be characterized. In a quantile VAR, one could study all the possible combinations of quantiles. In the figure, we report the impact of the shocks on different quantiles of each variable, conditioning on a median evolution of the other variables. If the data generating process were linear then the conditional median responses

reported in Figure 2 would coincide with standard IRFs from a linear VAR.

The asymmetries implied by the Wald test outcomes are clearly visible in the shapes of the QIRF. As expected, the real GDP response to a shock to the CISS depends markedly on the quantile of interest. The bottom 0.1 quantile of real GDP responds much more strongly than its upper 0.9 quantile. This is not surprising, and in line with [Adrian et al. \(2020\)](#). The response of the CISS to a shock to the financial cycle is highly asymmetric. A shock to the financial cycle does not move the CISS much in most parts of the CISS's distribution. The upper 0.9 quantile of the CISS, however, displays a marked negative response in the short term that disappears after a year or so. Vice versa, positive shocks to the CISS depress the financial cycle approximately uniformly across the financial cycle's distribution. The interactions between the CISS and the financial cycle suggest a vicious circle to emerge during crisis periods, with an increase in financial stress triggering lower financial activity, which in turn causes stress to further increase, and so forth. Such dynamics may reflect, for example, negative externalities from financial institutions' deleveraging efforts spilling over to other segments of the financial system, with potential feedback effects, during financial turmoil.

Web Appendix [F.4](#) discusses the analogous results for U.S. data. The Wald test outcomes and impulse response function estimates are remarkably similar across both sets of data.

4.2 Estimates of downside risk and upside potential

This section discusses our downside risk and upside potential estimates as introduced in Section [2.1](#). Web Appendix [F.5](#) discusses the U.S. case.

Figure 3 plots the average future growth shortfall (AGS) and longrise (AGL) for the euro area. The risk estimates are based on full-sample parameter estimates, but are otherwise conditional on variables observed up to time t only. Growth shortfall and longrise are forward-looking, and averaged over $t + 1$ and $t + 8$; see [\(3\)](#) and [\(5\)](#). To study the importance of including current financial conditions and medium-term vulnerabilities we compare our baseline QVAR model to a much simpler, univariate quantile autoregressive (QAR) model for GDP growth only. The QAR

model does not include the financial cycle nor the CISS.

We focus on three findings. First, accounting for financial conditions is crucial. There is a pronounced difference between the downside risk (AGS) estimate implied by the trivariate QVAR and the univariate QAR. During the global financial crisis (GFC) between 2008 and 2009, the QAR-based downside risk estimate declines much later, and by much less, compared to the QVAR-based estimate. The sovereign debt crisis between 2010 and 2012 is missed almost entirely based on the QAR model.


Second, as a result of macro-financial interactions, the QVAR's lower quantiles for future GDP growth are more volatile than its upper quantiles. This observation mirrors those of [Adrian et al. \(2019\)](#), who focus on single quantiles in isolation. Web Appendix E.4 plots the tail conditional expectation and expected probability of a contraction underlying $AGS_{t,t+1:t+H}^{\tau}$ and $AGL_{t,t+1:t+H}^{\tau}$ separately; see the first and second term in (2) and (4), respectively. Most of the variability in $AGS_{t,t+1:t+8}^{\tau}$ comes from the changes in growth conditional on being in a recession-term, with an additional contribution from increasing the probability of a contraction in bad times.

Lastly, the downside risks implied by the QVAR model can be economically large. The GFC implied an extreme AGS over eight quarters of approximately -3.5%. This corresponds to a $(1 - 0.035/4)^8 - 1 \approx -6.8\%$ reduction in real living standards. This is a substantial expected contraction, reflecting severe downside risks from a deterioration of financial conditions. During median times, the estimated AGS is approximately -0.5% and corresponds to a more moderate risk of a $(1 - 0.005/4)^8 - 1 \approx -1.0\%$ reduction in real living standards.

From a risk management perspective the AGS can be compared to the AGL as the latter provides an indication of the upside for the economy. The GFC did not only generate an extremely low value for the AGS, but also for the AGL. With a value of only 0.4%, the average expected expansion of the economy over the following eight quarters would have been approximately 0.8%. This compares to an average of approximately 4% over the entire sample. The GFC thus reduced living standards especially because of the contraction, but also persistently muted the upside potential of the economy, and did so until early 2015.

Figure 3: Euro area AGS and AGL estimates

Time- t average future growth shortfall ($AGS_{t,t+1:t+8}^{\tau}$) and average future growth longrise ($AGL_{t,t+1:t+8}^{\tau}$) estimates evaluated at $\tau = 0$; see (3) and (5). The trivariate estimate is based on our baseline QVAR model (dashed line, scale on left axis) that allows for macro-financial interactions. The univariate estimate is based on a one-equation restricted model with a constant and lagged GDP growth as the only right-hand-side variables (dotted line, scale on left axis). Each model is estimated for $p = 20$ quantiles ranging from 0.025 to 0.975. We compare these estimates to quarterly annualized real GDP growth (solid line, scale on right axis). Shaded areas indicate euro area recessions as determined by the CEPR business cycle dating committee. The estimation sample is 1988Q3 to 2018Q4.



4.3 Model-based stress testing

Our structural QVAR model provides a natural environment to perform model-based stress testing exercises. We here understand stress testing as a forecast of what would happen to \tilde{x}_t conditional on the system being subjected to a certain sequence of adverse shocks. We refer to such a sequence of adverse shocks as a stress scenario. For the computation of forecasts conditional on such scenarios we refer to Section 2.2.4. Our stress testing approach is different from supervisory stress tests in that our main variable of interest is not banking sector health but real economic (GDP) impact.

Figure 4 reports the time- t conditional forecast of average future real GDP growth $\tilde{y}_{t,t+1:t+4}$ between time t and $t + 4$ as implied by our trivariate model. The forecast is conditional on a 0.1 (conditional) quantile realization for GDP growth y_{t+h} , a 0.1 quantile realization of the financial

cycle c_{t+h} , and a 0.8 quantile realization for CISS s_{t+h} , consecutively for $h = 1, \dots, 4$. The magnitude of these shocks is approximately in line with the four observed quantile realizations for all variables between 2008Q2 and 2009Q2. The stress test is repeated at each $t = 1, \dots, T$, and always based on the same (full sample) parameter estimates. As a result, the figure is informative about the impact of GFC-sized real and financial shocks on real living standards at any time in our sample.¹⁷

We observe that the euro area economy is not equally resilient to the same sequence of equally unlikely adverse financial shocks at all times. This is a direct consequence of the asymmetries (nonlinearities) inherent in the estimated QVAR model. When financial imbalances and financial stress are high, real GDP growth is particularly vulnerable.

Figure 4 can be informative when assessing macro-prudential policy stance. An unusually high level of vulnerability to future real and financial shocks — a value of $\bar{y}_{t,t+1:t+4}$ below its own 10% quantile, say — indicates that large shocks have materialized and macro-prudential buffers should be released. In the euro area, such values are observed during the financial crisis of 2008 – 2009 and the sovereign debt crisis in 2011 – 2012. Low to moderate levels of vulnerability indicate times when macro-prudential buffers could be built up. Gradually growing macro-prudential capital buffers help increase banking sector resilience, lean against bad credit growth, improve incentives, and are available to be released later whenever necessary.

Web Appendix F.6 discusses the analogous figure for U.S. data. Similar observations hold true for these data as well.


4.4 Towards a metric for macro-prudential policy stance

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

¹⁷Alternatively, one could define stress in absolute size; see e.g. [Brownlees and Engle \(2017\)](#) for a discussion. We prefer the quantile-based approach because the probability of the stress materializing remains constant over time regardless of current levels of volatility. The severity of stress would be much higher in periods of low volatility, as it would take a sequence of more severe shocks to reach the same level of impact. On the other hand, low volatility does not necessarily imply that the tipping point is also low.

Figure 4: Vulnerability to GFC-sized shocks

Dashed line: euro area annualized quarterly real GDP growth. Solid line: predicted average annualized quarterly real GDP growth $\hat{y}_{t,t+1:t+4}$ one year ahead conditional on consecutive 0.1 quantile realizations for GDP growth y_t , 0.1 quantile realizations of the financial cycle c_t , and 0.8 quantile realizations for CISS s_t . Predictions are based on full sample parameter estimates. Estimations sample 1988Q3 – 2018Q4. Horizontal lines refer to 0.1, 0.5, and 0.9 empirical quantiles of $\hat{y}_{t,t+1:t+4}$.



conditional on currently-available information? We use the risk management framework of Section 2.3 and the associated objective function (16) to make a step towards addressing this question. We assume that macro-prudential instruments can be used to control the financial cycle and ask the question when it is optimal to use them.¹⁸ The thought experiment of this section is the following: How does the objective function of the macro-prudential authority change if the financial cycle is marginally reduced now, to be released later on should a financial crisis occur? If the change is positive, we conclude that the macro-prudential stance is too loose (as it would benefit from a less

¹⁸A complete answer to this question would require including policy instruments into our baseline model, such as bank capital and various interest rates. This can be done, at the cost of decreased model parsimony, parameter estimation precision, and identification credibility. Web Appendix E.1 suggests that our baseline downside risk and upside potential estimates are not affected, to first order, by an extension of the QVAR information set.

buoyant financial cycle). If on the other hand the answer is negative, macro-prudential policy is too tight.

Table 2 summarizes our policy experiment by contrasting two counterfactual scenarios. Each scenario looks three years into the future, equally split into two periods of $H/2 = 6$ quarters. The first six quarters are normal times during which the financial cycle could be marginally reduced. The second six quarters refer to a financial crisis during which the CISS takes high values and the financial cycle takes low values.

The bottom panel of Table 2 sets out an active, or marginally less passive, macro-prudential policy scenario. It is identical to the passive scenario in the top panel, except that the policy maker marginally reduces the financial cycle in the first period, for example by requiring higher counter-cyclical capital buffer requirements. During the financial crisis these buffers can be released, leading to a marginally less vicious collapse of the financial cycle. We simulate this policy by setting the financial cycle to its 0.5 quantile during $h = 1, \dots, 6$, instead of 0.6, and to its 0.2 conditional quantile during $h = 7, \dots, 12$, instead of 0.1. The evolution of GDP growth is always unrestricted. Doing so allows us to simulate forward the GDP growth rate, y_{t+h} , and growth shortfall, $GS_{t,t+h}$, at any time $t + h$, $h = 1, \dots, 12$.

Each policy scenario is evaluated as

$$u_t(\text{Scenario}) = \hat{y}_{t+1:t+12}(\text{Scenario}) + 0.50 \cdot \widehat{AGS}_{t,t+1:t+12}(\text{Scenario}), \quad (17)$$

where the mean growth estimate $\hat{y}_{t+1:t+12}$ and average future growth shortfall estimate $\widehat{AGS}_{t,t+1:t+12}$ are obtained from 10,000 simulations of potential future y_{t+h} . The utility function (17) operationalizes (16) by choosing parameters as $\beta = 1$ for $h = 1, \dots, 12$ and $\beta = 0$ thereafter, $\lambda = 1.50$, and $\tau = 0$. Our choice of the penalty term λ implies that the policy maker cares twice as much (positively) about future trend growth than she cares (negatively) about downside risk. We evaluate (17) twice, once for the active scenario and once for the passive scenario, and study the difference between the two at any time t .

Table 2: Passive vs. active macroprudential policy

The top and bottom panels report selected quantiles for GDP growth (y_t), financial cycle (c_t), and CISS (s_t) under a passive and active macroprudential policy benchmark, respectively. Multiple quantiles 0.1 – 0.9 mean that the quantile is picked at random. The first six quarters are normal times during which the financial cycle could, in principle, be marginally reduced. The second six quarters refer to a financial crisis during which the CISS takes high values and the financial cycle takes low values. If the financial cycle is actively managed in the first period then it does not have to contract as much during the crisis.


		first six quarters “normal times”	second six quarters “financial crisis”
passive	y_t	0.1 – 0.9	0.1 – 0.9
benchmark	c_t	0.6	0.1
	s_t	0.1 – 0.9	0.9
active	y_t	0.1 – 0.9	0.1 – 0.9
macro-pru	c_t	0.5	0.2
policy	s_t	0.1 – 0.9	0.9

Figure 5 plots the utility difference $\Delta u_t = u_t(\text{active}) - u_t(\text{passive})$ associated with adopting the active macro-prudential policy. Adopting the active policy is the preferred option most of the time. This is not surprising as we condition on a severe financial crisis in the second period. Adopting the active policy, however, 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 global financial crisis in 2007. This is intuitive, as the financial system was buoyant during these times, arguably seeding the respective busts later on. The benefits from leaning against the financial cycle are estimated to be the most negative following the global financial crisis of 2008, and during the euro area sovereign debt crisis between 2010 and 2012. This is again intuitive, as the financial system was already deleveraging during these times, and requiring more would add insult to injury. The utility difference Δu_t is mildly correlated with the euro area financial cycle, suggesting that it is a valuable variable to track to inform macroprudential policy discussions.

Web Appendix F.7 derives a stance metric for U.S. data. Adopting an active financial stability policy remains the preferred policy when the financial cycle is buoyant.

Figure 5: The benefits from active macro-prudential policy

The benefit of adopting an active macro-prudential policy stance in utility terms, $\Delta u_t = u_t(\text{active}) - u_t(\text{passive})$; see (17). Parameters are chosen as $\beta = 1$, $\lambda = 1.5$, $\tau = 0$, and $H = 12$. The difference is based on full sample estimates. Estimation sample is 1988Q3 to 2018Q4. Shaded areas indicate euro area recessions.



5 Conclusion

We proposed a risk management approach to macro-prudential policy that relates downside risks and upside potential of the economy to measures of financial stress and medium-term vulnerabilities. In an empirical study of euro area and U.S. data we found evidence of substantial asymmetries in the conditional distribution of GDP. The left quantiles of the predictive GDP growth distribution are related to a contemporaneous indicator of systemic stress, whose right quantiles are related to financial vulnerabilities. Counterfactual exercises allow us to perform model based stress testing, to construct urgently-needed indicators of macro-prudential policy stance, and to assess when macro-prudential interventions are relatively more likely to be beneficial.

References

- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review, forthcoming* 109(4), 1263–89.
- Adrian, T., N. Boyarchenko, and D. Giannone (2020). Multimodality in macro-financial dynamics. *NY Fed staff reports* 903, 1–54.
- Adrian, T., F. Grinberg, N. Liang, and S. Malik (2018). The term structure of growth-at-risk. *IMF working paper*.
- Aikman, D., J. Bridges, S. H. Hoke, C. O’Neill, and A. Raja (2019). Credit, capital and crisis: a GDP-at-risk approach. *Bank of England working paper* 824.
- Almon, S. (1965). The distributed lag between capital appropriations and net expenditures. *Econometrica* 33, 178–196.
- Artzner, P., F. Delbaen, J. M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance* 9(3), 203–228.
- 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 R. F. Engle (2017). SRISK: A Conditional Capital Shortfall Measure of Systemic Risk. *The Review of Financial Studies* 30(1), 48–79.
- Brownlees, C. and A. B. M. Souza (2020). Backtesting global growth-at-risk. *Journal of Monetary Economics, forthcoming*.
- Caballero, R. and A. Simsek (2020). Prudential monetary policy. *NBER WP* (25977).
- Caldara, D., C. Scotti, and M. Zhong (2019). Macroeconomic and financial risks: A tale of volatility. *mimeo*.
- Carlier, G., V. Chernozhukov, and A. Galichon (2016). Vector quantile regression: An optimal transport approach. *The Annals of Statistics* 44(3), 1165 – 1192.

- 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 (2020). Capturing macroeconomic tail risks with Bayesian vector autoregressions. *working paper 2020*, 1–68.
- Cecchetti, S. G. (2006). Measuring the macroeconomic risks posed by asset price booms. *NBER working paper No. 12542*, 1–31.
- Cecchetti, S. G. and H. Li (2008). Measuring the impact of asset price booms using Quantile Vector Autoregressions. *unpublished manuscript*, 1–31.
- 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 S. Manganelli (2019). Forecasting and stress testing with quantile vector autoregression. *ECB working paper 2330*, 1 – 43.
- 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.
- De Santis, R. and W. van der Veken (2020). Forecasting Macroeconomic Risk in Real Time: Great and Covid-19 Recessions. *ECB Working Paper forthcoming*.
- 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. .
- 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.

- 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*.
- 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.
- Hahn, J. (1995). Bootstrapping quantile regression estimators. *Econometric Theory* 11(1), 105–121.
- Hollo, D., M. Kremer, and M. L. Duca (2012). CISS – A composite indicator of systemic stress in the financial system. *ECB Working Paper 1426*.
- IMF (2017). Global financial stability report: Is growth at risk? *IMF Global Financial Stability Report 3*, 91–118.
- 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 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. Basset (1982). Robust tests for heteroscedasticity based on regression quantiles. *Econometrica* 50, 43 – 61.
- Koopman, S. J. and J. D. Durbin (2000). Fast filtering and smoothing for multivariate state space models. *Journal of Time Series Analysis* 21, 281–296.
- Koyck, L. M. (1954). *Distributed lags and investment analysis*. North-Holland, Amsterdam.
- McNeil, A., R. Frey, and P. Embrechts (2005). *Quantitative Risk Management*. Princeton: Princeton University Press.

- Mendicino, C., K. Nikolov, J. Suarez, and D. Supera (2018). Optimal dynamic capital requirements. *Journal of Money, Credit and Banking* 50(6), 1271–1297.
- Montes-Rojas, G. (2019). Multivariate quantile impulse response functions. *Journal of Time Series Analysis* 40, 739 – 752.
- Plagborg-Moller, M., L. Reichlin, G. Ricco, and T. Hasenzagl (2020). When is growth at risk? *Brookings Papers on Economic Activity*.
- 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*.
- Ruzicka, J. (2020). Quantile structural vector autoregression. *Universidad Carlos III working paper*, 1 – 70.
- Schüler, Y., P. Hiebert, and T. A. Peltonen (2020). Financial cycles: Characterisation and real-time measurement. *Journal of International Money and Finance* 100, 82–102.
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica* 48, 1–48.
- Suarez, J. (2021). Growth-at-risk and macroprudential policy design. *CEMFI mimeo*.
- 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.
- White, H., T. H. Kim, and S. Manganelli (2015). VAR for VaR: Measuring tail dependence using multivariate regression quantiles. *Journal of Econometrics* 187, 169–188.