

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecastA time-varying skewness model for Growth-at-Risk[☆]

Martin Iseringhausen*

European Stability Mechanism, Luxembourg

ARTICLE INFO

Keywords:

Bayesian analysis
Downside risk
Macro-financial linkages
Quantile forecasts
Time variation

ABSTRACT

This paper studies macroeconomic risks in a panel of advanced economies based on a stochastic volatility model in which macro-financial conditions shape the predictive growth distribution. We find sizable time variation in the skewness of these distributions, conditional on the macro-financial environment. Tightening financial conditions signal increasing downside risk in the short term, but this link reverses at longer horizons. When forecasting downside risk, the proposed model, on average, outperforms existing approaches based on quantile regression and a GARCH model, especially at short horizons. In forecasting upside risk, it improves the average accuracy for several horizons up to four quarters ahead. The suggested approach can inform policymakers' assessment of macro-financial vulnerabilities by providing a timely signal of shifting risks and a quantification of their magnitude.

© 2023 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.

1. Introduction

Many economic policy institutions regularly publish forecasts of economic growth as part of their assessment and as a foundation of their policy advice. The baseline forecasts are surrounded by a varying degree of uncertainty, and the discussion of both downside and upside risks to the growth outlook is an essential part of macroeconomic forecasting. In recent years, several institutions, starting with the International Monetary Fund (IMF), have been using a new approach to quantify macro-financial risks to growth, which has become prominently known as

Growth-at-Risk (see e.g. Prasad et al., 2019). The present paper develops a new empirical model to analyze macroeconomic risks stemming from changing macro-financial conditions. In our proposed framework, the skewness of the predictive growth distribution, a measure of unbalanced risks, is driven by such conditions. Through changes in the skewness, macro-financial conditions therefore also impact to what extent growth is at risk.

In a seminal paper, Adrian, Boyarchenko, and Giannone (2019) show that future downside risk varies significantly depending on current financial conditions. Using quantile regression to study the dynamics of the quantiles of the conditional US GDP growth distribution, they show that downside risk increases as financial conditions become tighter.¹ Inspired by the financial market concept of Value-at-Risk, subsequent work has labeled the lower conditional p -quantile of the predictive growth distribution Growth-at-Risk (GaR^p) (Adrian et al., 2021). Quantile regression has become a standard tool to analyze risks to economic outcomes both in academia and in policy

[☆] The views expressed in this paper are those of the author and do not necessarily reflect those of the European Stability Mechanism (ESM). The author would like to thank Bruno Albuquerque, Antonello D'Agostino, Wouter Van der Veken, Alexander Raabe, Jemima Peppel-Srebrny, two anonymous referees and an associate editor, the editor Esther Ruiz, and seminar and conference participants at the ESM, the Bank of England and the 27th International Conference Computing in Economics and Finance for helpful comments, as well as Sergei Antoshin for sharing the Financial Conditions Indices (FCI) produced by the IMF.

* Correspondence to: European Stability Mechanism, 6a Circuit de la Foire Internationale, L-1347, Luxembourg.

E-mail address: m.iseringhausen@esm.europa.eu.

<https://doi.org/10.1016/j.ijforecast.2023.02.006>

0169-2070/© 2023 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.

¹ Extending this analysis to a larger set of advanced economies, Adrian, Grinberg, Liang, Malik, and Yu (2021) highlight an intertemporal tradeoff between short-term benefits and medium-term risks of loose financial conditions.

institutions (Caldera Sánchez & Röhn, 2016; Giglio, Kelly, & Pruitt, 2016; Prasad et al., 2019). While the approach was used in earlier studies to evaluate the link between financial variables and economic tail outcomes, Adrian et al. (2019) study the full conditional growth distribution by fitting, in a second step, a parametric distribution over the predicted quantiles. This allows them to obtain another common risk measure, namely expected shortfall (on the downside) and expected longrise (on the upside), defined as the expected growth rate conditional on the occurrence of a tail event (see also Adams, Adrian, Boyarchenko, & Giannone, 2021).

Quantile regression offers a simple tool to understand asymmetries of the growth distribution and the role that financial variables play in shaping this distribution, but the approach has recently been scrutinized. Several papers have questioned the ability of financial variables to inform the analysis of downside risk (e.g. Plagborg-Møller, Reichlin, Ricco, & Hasenzagl, 2020) or, at least, the stability of the relationship between these conditions and future growth vulnerability (Reichlin, Ricco, & Hasenzagl, 2020). Moreover, even if studies recognize that financial variables affect macroeconomic risks, the question remains whether this can benefit out-of-sample forecasting. Brownlees and Souza (2021) show that standard time-varying volatility (GARCH) models compete well with quantile regressions in forecasting Growth-at-Risk and often outperform them despite the fact that GARCH forecasts are only based on growth data. In addition, Carriero, Clark, and Marcellino (2020a) present evidence that multivariate time series (VAR) models with stochastic volatility—commonly used tools for both structural macroeconomic analysis and forecasting—effectively capture time-varying risks to growth. While in these models financial conditions also play an important role, they conclude that asymmetric conditional distributions, implied by the results of Adrian et al. (2019), do not have to be incorporated into empirical models to accurately model the quantiles of interest.

This paper revisits the question of whether conditional skewness, where a negative (positive) skewness value is interpreted as prevalent macroeconomic downside (upside) risk, is a relevant empirical feature that could help to improve Growth-at-Risk and expected shortfall/longrise (ES/EL) forecasts.² The contribution to the existing literature is both methodological and empirical. First, the paper proposes a parametric model that measures asymmetric risk characteristics of the conditional GDP growth distribution as a function of exogenous variables. The approach includes both time-varying volatility, which has been shown to be crucial for forecasting accuracy (see e.g. D'Agostino, Gambetti, & Giannone, 2013), and the possibility to evaluate the effect of macro-financial conditions on the presence of unbalanced risks. Second, it shows that the latter can further help to improve risk forecasts and increases the applicability in a policy environment,

where the goal is not just to predict risks, but also to identify the sources of risks. While most interest usually lies in forecasting macroeconomic downside risk, this paper considers both tails of the predictive growth distribution. Only by analyzing downside and upside risk jointly can we make statements about the 'balance of risks' around a given baseline outlook.

Methodologically, the model extends the stochastic volatility model with time-varying skewness developed in Iseringhausen (2020) for a single time series of daily exchange rate returns, to a panel of countries using low-frequency macroeconomic data. To capture evolving risks, the shocks are assumed to follow the noncentral t-distribution, where the asymmetry parameter of this distribution is a function of exogenous variables, namely macro-financial conditions. The model is estimated by an extension of the well-known Bayesian approach for stochastic volatility models developed in Kim, Shephard, and Chib (1998). This empirical model can be seen as a stochastic version of existing approaches that model the scale and shape parameters of a distribution as a purely deterministic function of explanatory variables. The first contributions in this literature are extensions of GARCH-type models to allow for higher-order dynamics (Hansen, 1994; Harvey & Siddique, 1999). Recently, Plagborg-Møller et al. (2020) used a model with skewed-t innovations, where the parameters of the distribution are driven by an aggregate economic factor and a financial factor. Their results suggest that for US output growth, moments beyond the conditional mean cannot be pinned down precisely in such a framework. By contrast, in a more general score-driven version of this model, Delle Monache, De Polissis, and Petrella (2020) find that both the volatility and skewness of US growth vary significantly over time. In a SVAR framework, Montes-Galdón and Ortega (2022) allow the distribution of structural shocks to be skewed and find that the identified demand, supply, and monetary policy shocks in the euro area feature time-varying asymmetry. The above-mentioned approaches, including the one developed in this paper, can also be viewed as alternatives to Markov-switching approaches (e.g. Caldara, Cascardi-Garcia, Cuba-Borda, & Loria, 2020; Hamilton, 1989; Morley & Piger, 2012) or threshold models (e.g. Alessandri & Mumtaz, 2017) to allow for non-linear effects of certain determinants on the distribution of macroeconomic outcomes.

When applying the proposed model to a panel of 11 OECD countries over the period from 1973:Q1–2019:Q4, we obtain the following results. First, in-sample financial conditions have a sizable impact on the skewness of the predictive growth distribution, which displays large time variation and is skewed to the left in most countries. In particular, tightening financial conditions are related to elevated future near-term downside risk, but this relationship reverses at longer horizons. Second, including a financial conditions index (FCI) in the process of the model's asymmetry parameter can improve average GaR and ES/EL forecasts compared to alternative approaches based on quantile regression or a GARCH model. For downside risk, the gains occur especially at short horizons up to two quarters ahead, and for upside risk

² Contributions providing theoretical support for a time-varying degree of asymmetry over the business cycle include, for example, Orlik and Veldkamp (2014), Salgado, Guvenen, and Bloom (2019), and Jensen, Petrella, Ravn, and Santoro (2020).

at various horizons.³ Lastly, including a measure of economic and policy uncertainty next to the FCI, or some of the prominent components of the index separately, the term spread or house price growth, does not generally improve (or adds only very little to) the accuracy of the model.

The remainder of the paper is structured as follows: Section 2 introduces a panel stochastic volatility model with time-varying skewness, which is driven by macro-financial conditions. Moreover, the estimation of the model using Bayesian methods is discussed. Section 3 presents the results both in-sample and when forecasting macroeconomic risk out-of-sample. Section 4 concludes. The appendix contains further methodological details and additional results.

2. A panel model with time-varying volatility and skewness

This section introduces an empirical specification to estimate the predictive distribution of GDP growth for a panel of countries. The model features time-varying skewness, where asymmetry is driven by a set of explanatory variables. It is an extension of the stochastic volatility–stochastic skewness model of Iseringhausen (2020), and the presentation of both the specification and the estimation approach closely follows that paper.

2.1. Empirical specification

The observed dependent variable y is assumed to be generated by the following panel stochastic volatility (SV) specification:

$$y_{it} = \mu_{it} + e^{h_{it}/2} \varepsilon_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \quad (1)$$

$$h_{it} = h_{i,t-1} + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, \sigma_h^2), \quad (2)$$

where y_{it} is GDP growth in country i and period t , and μ_{it} is the conditional mean, which is a linear function specified later with the corresponding vector of coefficients labeled γ_i . The country-specific (log-)volatility of growth shocks, denoted h_{it} , is assumed to evolve according to a random walk with innovation variance ('volatility of volatility') σ_h^2 that is pooled across countries.⁴ If the shocks ε_{it} are assumed to follow the standard normal distribution, the model is a panel version of the well-known standard normal SV model (e.g. Kim et al., 1998).

To analyze time-varying asymmetries in the growth distribution, we deviate from the assumption of Gaussian shocks and instead assume that these follow the (de-meant) noncentral t-distribution,

$$\varepsilon_{it} = u_{it} - \mathbb{E}[u_{it}], \quad \text{with } u_{it} \sim \mathcal{NCT}(\nu, \delta_{it}), \quad \text{and} \\ \mathbb{E}[u_{it}] = c_{11}(\nu)\delta_{it}, \quad \text{if } \nu > 1. \quad (3)$$

³ When comparing the proposed model with simpler stochastic volatility specifications, the results are somewhat weaker. However, time-varying skewness can still help in several cases to improve forecasts of both tails of the predictive growth distribution.

⁴ The random walk assumption for modeling (log-)volatility has become standard in empirical macroeconomics (see e.g. Antolin-Diaz, Drechsel, & Petrella, 2017; D'Agostino et al., 2013; Primiceri, 2005).

The shape of the noncentral t-distribution is driven by two parameters: the degrees of freedom ν , which are assumed to be homogeneous across countries; and the time-varying noncentrality parameter δ_{it} . The functional form of the coefficient $c_{11}(\nu)$ can be found in Appendix A. For $\delta > 0$, this distribution is skewed to the right, whereas for $\delta < 0$, the distribution is skewed to the left. The noncentral t-distribution is introduced into a univariate stochastic volatility framework by Tsiotas (2012). Iseringhausen (2020) extends this model to allow for a time-varying asymmetry parameter δ_t , modeled either as a stationary autoregressive process or, alternatively, as a random walk. Using the model, Iseringhausen (2020) documents time-varying skewness, which can be interpreted as changing downside risk, in daily exchange rate returns. This paper imposes a different structure on the evolution of the asymmetry parameter by assuming that δ_{it} is driven explicitly by a vector of explanatory variables,

$$\delta_{it} = \phi\delta_{i,t-1} + X_{\delta,it}\beta + \omega_{it}, \quad \omega_{it} \sim \mathcal{N}(0, \sigma_\delta^2), \quad |\phi| < 1, \quad (4)$$

where $X_{\delta,it}$ is of dimension $1 \times K$ and contains the explanatory variables including a constant, and β is the corresponding $K \times 1$ vector of coefficients, pooled across countries. Since the relation between δ_{it} and $X_{\delta,it}$ is likely not exact, a zero mean error term ω_{it} with pooled variance σ_δ^2 is added.

Since the model will be estimated by Bayesian methods, the specification of this stochastic volatility model with time-varying skewness is completed by assuming the following prior distributions for the parameters σ_h^2 , σ_δ^2 , ν , β , and ϕ :

$$\sigma_h^2 \sim \mathcal{IG}(c_{h0}, C_{h0}), \quad \sigma_\delta^2 \sim \mathcal{IG}(c_{\delta0}, C_{\delta0}), \quad \nu \sim \mathcal{U}(0, \bar{\nu}), \quad (5)$$

$$\beta \sim \mathcal{N}(\beta_0, \sigma_{\beta0}^2 I_K), \quad \phi \sim \mathcal{TN}_{(-1,1)}(\phi_0, \sigma_{\phi0}^2).$$

To obtain closed-form expressions of the model-implied variance and skewness for each country and period, we rely on the result of Hogben, Pinkham, and Wilk (1961), stating that the central moments of a noncentral t-distributed random variable, $X \sim \mathcal{NCT}(\nu, \delta)$, can be written as polynomials of δ whose coefficients are functions of ν . Specifically, Hogben et al. (1961) derive expressions for the second and third central moments of the noncentral t-distribution, which in turn can be used to obtain formulas for the time-varying variance and skewness of GDP growth implied by the panel model introduced in this section:

$$\text{Var}[y_{it} | h_{it}, \delta_{it}, \nu] = e^{h_{it}} [c_{22}(\nu)\delta_{it}^2 + c_{20}(\nu)], \quad \text{if } \nu > 2, \quad (6)$$

$$\text{Skew}[y_{it} | \delta_{it}, \nu] = \frac{c_{33}(\nu)\delta_{it}^3 + c_{31}(\nu)\delta_{it}}{[c_{22}(\nu)\delta_{it}^2 + c_{20}(\nu)]^{3/2}}, \quad \text{if } \nu > 3. \quad (7)$$

The functional expressions of the coefficients $c_{20}(\nu)$, $c_{22}(\nu)$, $c_{31}(\nu)$, and $c_{33}(\nu)$ can again be found in Appendix A.

We want to point out an important aspect of the presented specification. As already discussed by Iseringhausen (2020), the error term in Eq. (3) has zero mean—i.e. δ_{it} does not affect the conditional mean of y_{it} —but it is not standardized to have unit variance, implying that δ_{it} appears in the conditional variance equation. Put differently, the scaling factor h_{it} is not equal to the conditional log-variance. This has practical reasons related to the implementation of the proposed MCMC algorithm. Specifically, it would no longer be possible using standard algebra to derive a state-space representation with an observation equation that is linear in δ_{it} , which is required to apply the proposed sampling algorithm.⁵ While the current model specification implies more complicated formulas for the conditional variance and conditional skewness, this specification has great advantages in terms of practical implementation. Moreover, conditional variance and skewness can still move independently, since two state variables, h_{it} and δ_{it} , are driving the two moments. For later use, we define $y = (y_1, \dots, y_T)'$, $\mu = (\mu_1, \dots, \mu_T)'$, $h = (h_1, \dots, h_T)'$, $\delta = (\delta_1, \dots, \delta_T)'$, and $\lambda = (\lambda_1, \dots, \lambda_T)'$, where each element, y_t , μ_t , h_t , δ_t , and λ_t , is an $N \times 1$ vector. Finally, define $X_\delta = (X_{\delta,1}, \dots, X_{\delta,N})'$, where each $X_{\delta,i}$ is of dimension $T \times K$.

2.2. Bayesian estimation: Building blocks and MCMC algorithm

The model developed in this paper is estimated using Markov chain Monte Carlo (MCMC) methods. Before outlining the detailed algorithm, a few preliminary aspects need to be discussed that are crucial for the derivation of some of the conditional posterior distributions that constitute the key components of the Gibbs sampling algorithm (see also Iseringhausen, 2020, for more details).

Location-scale mixture representation of the noncentral t-distribution

First, in order to avoid working directly with the complex probability density function of the noncentral t-distribution, we make use of the fact that this type of distribution can be written as a location-scale mixture of normal distributions (Johnson, Kotz, & Balakrishnan, 1995; Tsionas, 2002). Specifically, if the random variable X follows the noncentral t-distribution with noncentrality parameter δ and degrees of freedom ν , it has the following representation:

$$X = \sqrt{\lambda}(z + \delta), \quad \text{where } \lambda \sim \mathcal{IG}(\nu/2, \nu/2) \text{ and } z \sim \mathcal{N}(0, 1). \quad (8)$$

When applying Eq. (8) to the time-varying skewness model, the observation equation, obtained by merging Eqs. (1) and (3), can be written as

$$y_{it} = \mu_{it} + e^{h_{it}/2} \varepsilon_{it} = \mu_{it} + e^{h_{it}/2} \left(\sqrt{\lambda_{it}}(z_{it} + \delta_{it}) - c_{11}(\nu)\delta_{it} \right), \quad (9)$$

where again $\lambda_{it} \sim \mathcal{IG}(\nu/2, \nu/2)$ and $z_{it} \sim \mathcal{N}(0, 1)$. This representation of the model can be viewed as a type of data augmentation in the sense of Tanner and Wong (1987), where introducing the additional latent variable λ_{it} facilitates the derivation of an observation equation that is linear in δ_{it} . This is essential for the implementation of the MCMC algorithm presented later in this section.

Augmented auxiliary sampler

Having discussed above how to obtain linearity in δ_{it} , we now turn to how to linearize the model for the purpose of sampling the unobserved (log-)volatility h_{it} . The approach taken here directly follows (Iseringhausen, 2020), who extends the so-called auxiliary sampler for the estimation of the Gaussian stochastic volatility model developed by Kim et al. (1998) to the noncentral-t model with a time-varying asymmetry parameter. Consider the following transformation of the previously derived observation Eq. (9), obtained after squaring both sides and taking the natural logarithm:

$$\log((y_{it} - \mu_{it})^2 + c) = h_{it} + \tilde{\varepsilon}_{it}, \quad (10)$$

where $c = 10^{-6}$ is an offset constant, and where the transformed error term $\tilde{\varepsilon}_{it}$ is

$$\tilde{\varepsilon}_{it} = \log(\varepsilon_{it}^2) = \log \left[\left(\sqrt{\lambda_{it}}(z_{it} + \delta_{it}) - c_{11}(\nu)\delta_{it} \right)^2 \right]. \quad (11)$$

This transformation allows h_{it} to now enter the model in a linear manner. If the error term in Eq. (10) would be Gaussian, standard algorithms for linear Gaussian state-space models, based on the Kalman filter, could be used to estimate h . The non-normal distribution of $\tilde{\varepsilon}_{it}$ requires an extension to the standard estimation approach. We follow Iseringhausen (2020), who, based on the idea developed in Kim et al. (1998), approximates the log-noncentral-t-squared distribution of $\tilde{\varepsilon}_{it}$ by a mixture of normal distributions,

$$f(\tilde{\varepsilon}_{it} | \nu, \delta_{it}) = \sum_{j=1}^M q_j(\nu, \delta_{it}) f_{\mathcal{N}}(\tilde{\varepsilon}_{it} | m_j(\nu, \delta_{it}), v_j^2(\nu, \delta_{it})). \quad (12)$$

In this expression, $q_j(\nu, \delta_{it})$ is the corresponding probability of a specific normal distribution with mean $m_j(\nu, \delta_{it})$ and variance $v_j^2(\nu, \delta_{it})$. In contrast to Kim et al. (1998), where a single mixture distribution is sufficient to approximate the target distribution, here one needs a large grid of mixture distributions, since the specific target distribution to be approximated depends on the values of ν and δ_{it} . To this end, a large number of mixture distributions are ‘pre-fitted’, such that this does not involve any additional computational costs in the estimation. In setting the number of mixture components M , we follow Omori, Chib, Shephard, and Nakajima (2007) and Iseringhausen (2020) and choose $M = 10$.⁶ The mixture representation in Eq. (12) can be reformulated based on

⁶ Additional details on this approach can be found in Iseringhausen (2020), and the mixture components (means, variances, and probabilities) for each combination of ν and δ can be obtained from the author.

⁵ See Equation (B-17) in Appendix B.

the probabilities of the Gaussian components,

$$\tilde{\varepsilon}_{it} | s_{it} = j, v, \delta_{it} \sim \mathcal{N}(m_j(v, \delta_{it}), v_j^2(v, \delta_{it})), \quad (13)$$

$$\Pr(s_{it} = j | v, \delta_{it}) = q_j(v, \delta_{it}),$$

where the mixture indicators $s_{it} \in [1, \dots, 10]$ are unobserved and can be sampled jointly with the remaining parameters.

MCMC algorithm

Equipped with these two statistical tools, we are now ready to develop a Gibbs sampler that simulates draws from the intractable joint and marginal posterior distributions of the parameters and unobserved states by only exploiting conditional distributions. In some cases, where these are not members of standard distributional families, a Metropolis–Hastings step is added to the respective block of the sampler. Appendix B contains a more detailed presentation of the MCMC approach. The algorithm is an adjusted version of the one developed by Iseringhausen (2020) and loops over the following blocks:

- | | |
|--|--|
| 1. Conditional mean coefficients: | draw γ from $p(\gamma y, h, \delta, \lambda, v)$; |
| 2. Mixture indicators: | draw s from $p(s y, h, \delta, v, \gamma)$; |
| 3. (Log-)volatility: | draw h from $p(h y, s, \delta, v, \gamma, \sigma_h^2)$; |
| 4. Location-scale weights: | draw λ from $p(\lambda y, h, \delta, v, \gamma)$; |
| 5. Degrees of freedom: | draw v from $p(v \lambda)$; |
| 6. Noncentrality parameter: | draw δ from $p(\delta y, h, \lambda, v, \gamma, X_\delta, \phi, \beta, \sigma_\delta^2)$; |
| 7. Coefficients in δ -equation: | draw ϕ and β from $p(\phi, \beta \delta, X_\delta, \sigma_\delta^2)$; |
| 8. Innovation variances: | draw σ_h^2 from $p(\sigma_h^2 h)$ and σ_δ^2 from $p(\sigma_\delta^2 \delta, X_\delta, \phi, \beta)$. |

In terms of practical implementation, blocks 1, 2, 3, 4, and 6 are executed country-by-country, while blocks 5, 7, and 8 draw the pooled components. Block 1 generates a draw of the conditional mean μ , which in our application has a linear specification and includes a constant and lagged values of the dependent variable. Specifically, we sample the corresponding $k \times 1$ vector of country-specific regression coefficients γ_i as in Tsonas (2002). The conditional mean is then simply retrieved as $\mu_{it} = X_{\mu,it} \gamma_i$, where $X_{\mu,it}$ is of dimension $1 \times k$ and contains the conditional mean regressors. Block 2 samples the mixture indicators via the inverse-transform method (Kim et al., 1998). The different mixture components for each period t , each country i , and each Gibbs iteration j are selected depending on the corresponding (rounded) values of v_j and $\delta_{i,t,j}$. The (log-)volatilities h (block 3) and noncentrality parameters δ (block 6) are sampled using sparse matrix algorithms (Chan & Hsiao, 2014; Chan & Jeliazkov, 2009) that can yield large computational efficiency gains compared to more standard algorithms for estimating linear Gaussian state-space models (e.g. Carter & Kohn, 1994). The conditional posterior distributions of λ (block 4) and v (block 5) are non-standard and a Metropolis–Hastings step needs to be included, as described by Tsonas (2002) and Chan and Hsiao (2014),

respectively. Sampling the coefficients of the asymmetry process, ϕ and β (block 7), is relatively standard, as they follow (truncated) normal distributions, where an acceptance–rejection step is included to ensure that ϕ remains in the stationary region. Finally, the innovation variances σ_h^2 and σ_δ^2 (block 8) follow inverse-gamma distributions and sampling is standard. The algorithm is initialized with an arbitrary set of starting values.

After iterating the algorithm for an initial burn-in period of length B followed by another J iterations, the sequence of draws $(B + 1, \dots, J)$ can be taken as a sample from the joint posterior distribution $f(h, \delta, \lambda, v, \gamma, \phi, \beta, \sigma_h^2, \sigma_\delta^2 | y, X_\delta)$. In the following analysis, we discard the initial 20,000 draws as burn-in to ensure convergence to the ergodic distribution. Afterwards, we iterate the MCMC algorithm another 2,000,000 times, keeping only every 20th draw.⁷ The results presented in the following sections are thus based on 100,000 effective draws.

In order to evaluate the performance of this MCMC algorithm, Appendix C presents a small Monte Carlo simulation exercise. In particular, the results show that the algorithm can successfully recover the true data-generating model parameters.⁸ Moreover, Appendix C comments on the convergence properties and efficiency of the sampler.

3. Results

After describing the dataset, this section proceeds with an in-sample predictive analysis to measure the impact of macro-financial conditions on the conditional growth distribution. Finally, it contains an out-of-sample forecasting exercise to compare the proposed model with alternative approaches from the literature.

3.1. Data

The dataset used for the analysis is a balanced panel of 11 OECD countries over the period from 1973:Q1–2019:Q4. The countries included are Australia (AU), Canada (CA), France (FR), Germany (DE), Italy (IT), Japan (JP), Spain (ES), Sweden (SE), Switzerland (CH), the United Kingdom (GB), and the United States (US). Economic growth is measured by the quarter-over-quarter growth rate of seasonally adjusted real GDP obtained from the OECD database.

The country-specific variables entering the state equation of the asymmetry parameter are selected based on the existing literature. We consider a financial conditions index (FCI), which is calculated by the IMF (2017) until 2016:Q4 from a large set of domestic and global financial indicators using the common factor approach of Koop and Korobilis (2014). FCIs have emerged in the literature as the single most important determinant of Growth-at-Risk (see for example Brownlees & Souza, 2021). The

⁷ Applying this ‘thinning’ has only computational reasons, i.e. memory limits (see Gelman & Shirley, 2011).

⁸ In addition, a simulation exercise in Iseringhausen (2020) (Appendix D) shows that the model does not indicate time-varying skewness if the data-generating process is in fact symmetric.

IMF's methodology to compute the FCI has been simplified and the set of underlying components somewhat changed in IMF (2018), but this revised series only starts in 1996. To obtain a longer sample, we splice both series in 2016:Q4 by re-scaling the new series and removing a potential level difference so that both series have identical values in 2016:Q4.

In addition, we include variables that have been among the more promising predictors considered in the study of Brownlees and Souza (2021) or that regularly draw the attention of policymakers: the term spread (TS) and real house price growth (HP), both obtained from the OECD, and an index of economic and policy uncertainty (WUI), following Ahir, Bloom, and Furceri (2018).⁹ All variables are standardized by subtracting their panel-wide mean and dividing by their panel-wide standard deviation to facilitate a straightforward interpretation and comparison of the estimated coefficients. Moreover, this also allows for a sensible prior configuration that is uniform across variables.

For around half of the countries, there are some missing values for the FCI and the term spread. We follow Brownlees and Souza (2021) in imputing these observations and take the imputed values directly from the dataset made available by these authors. Details on the dataset, including the imputation procedure (see also Brownlees & Souza, 2021), as well as time series plots of the explanatory variables, can be found in Appendix D. As a final remark, we try to limit distortions due to imputed values by including only advanced economies for which an individual FCI is available for at least most of the sample period.

3.2. In-sample parameter estimates and unobserved components

The presentation of the results starts with an in-sample analysis of the predictive growth distribution. In this section, for the sake of brevity, we focus mostly on results for one-step-ahead predictions ($h = 1$), but also discuss selected important insights for longer horizons ($h = 4$ and $h = 8$). The in-sample results presented here stem from a direct multi-step-ahead analysis using the model shown in Eqs. (1)–(4), where the dependent variable y_{it} is simply shifted by h periods.¹⁰ Thus, as in Brownlees and Souza (2021), throughout the paper we predict the h -step-ahead quarter-over-quarter growth rate of real GDP, instead of

⁹ The inclusion of an uncertainty measure is also motivated by the findings in Hengge (2019) and Jovanovic and Ma (2021), who, using the forecast error based uncertainty measure of Jurado, Ludvigson, and Ng (2015), find a strong effect of uncertainty on the lower quantiles of the conditional US output (GDP/IP) growth distribution. Unfortunately, this uncertainty measure is not readily available for all countries in our sample.

¹⁰ The term 'predictive' in the in-sample analysis refers to the timing of the variables in the conditional mean and asymmetry equation. For the (log-)volatility series h , we report the actual estimated posterior mean for each period $t+h$. Alternatively, Carriero et al. (2020a) present a way to account for some degree of uncertainty around the path of the latent volatility series in an in-sample predictive analysis. Finally, the AR term in $\delta_{i,t+h}$ is always lagged by one instead of h periods, i.e. $t+h-1$.

the average cumulative growth (for $h > 1$) as in Adrian et al. (2019). The conditional mean specification $\mu_{i,t+h|t}$ includes a constant, current GDP growth in period t , and three additional lags.¹¹

The prior configurations used for both the in-sample and the out-of-sample analysis can be considered largely uninformative. For the regression coefficients in the mean equation and the asymmetry equation, we use $\gamma_i \sim \mathcal{N}(\mathbf{0}, 10 \times I_k)$, $\phi \sim \mathcal{TN}_{(-1,1)}(\mathbf{0}, 10)$, and $\beta \sim \mathcal{N}(\mathbf{0}, 10 \times I_k)$. The priors on the innovation variances σ_h^2 and σ_δ^2 are taken from Kim et al. (1998), $\sigma^2 \sim \mathcal{IG}(2.5, 0.025)$. This implies a prior expectation of 0.017 and a prior standard deviation of 0.024. The upper bound of the degrees-of-freedom parameter is set to $\bar{\nu} = 30$, since for this value, the (noncentral) t-distribution becomes nearly indistinguishable from the standard normal distribution.

Before turning to the analysis of skewness, Fig. 1 shows the estimated model-implied variance of the one-step-ahead predictive distribution for each country, calculated using Eq. (6). These plots clearly show the well-known Great Moderation, i.e. a significant decline in output volatility in most industrialized economies with some differences in the timing and magnitude across countries (Del Negro & Otrok, 2008; Stock & Watson, 2003; Summers, 2005). Some visible spikes in volatility, which contrast with the even smoother volatility estimates that are usually obtained from SV models when applied to macroeconomic data, are the result of the fact that the explanatory variables in the asymmetry equation also affect the conditional variance, as explained in the previous section.

Fig. 2 presents one of the key results of the paper: the estimated conditional skewness of the one-step-ahead predictive GDP growth distribution across countries. First, there is significant movement in the evolution of the skewness series over time and across countries. For the US, Adrian et al. (2019) present similar evidence using quantile regression and Delle Monache et al. (2020) based on their score-driven time-varying skewed-t model. Second, for most countries and periods, skewness is negative, indicating that the left tail of the one-step-ahead predictive growth distribution is usually longer than the right tail. This is a first piece of evidence that suggests that the existence of sizable downside risk is a common phenomenon across countries and over time. In particular, during the Great Recession, the left tails of conditional growth distributions are pronounced, pointing to a highly vulnerable macroeconomic environment. Importantly, the pooled approach chosen in this paper to estimate the asymmetry coefficients, including the intercept, could mask cross-country heterogeneity, but is necessitated by the relatively low frequency of macroeconomic time series compared to, for example, financial market data.

To analyze the role of financial conditions and uncertainty in determining the shape of the predictive growth

¹¹ The autoregressive conditional mean specification does not reduce the effective sample size, since quarterly growth rates for the 11 countries are available for a sufficient number of periods prior to 1973:Q1.

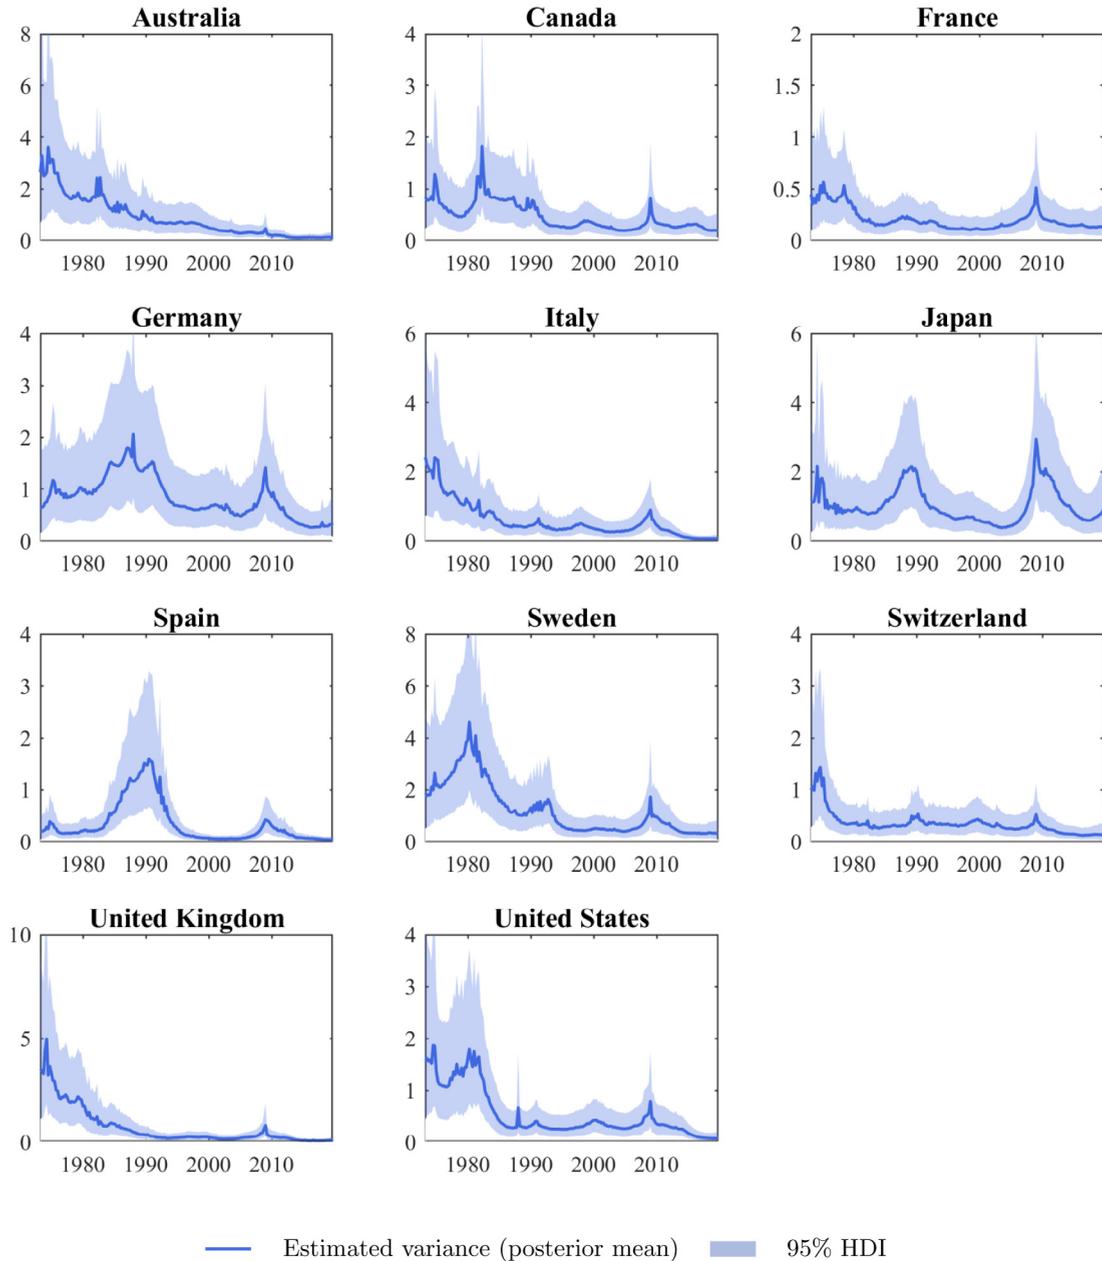


Fig. 1. Variance of in-sample one-step-ahead predictive distribution.

distribution, Table 1 presents the estimated model parameters, with the exception of the conditional mean coefficients. We start by discussing the (pooled) coefficients on the explanatory variables in the asymmetry process. The exogenous variables are assumed to impact the asymmetry coefficient δ in a linear fashion according to Eq. (4).

However, conditional skewness is a non-linear function of δ , determined by Eq. (7). Therefore, to interpret the direct effect of the regressors on skewness, we report the

marginal effects at the average (MEA). Since the data are standardized, these represent the effect of a one standard deviation increase in the respective variable on skewness, assuming that all variables are at their sample means. The MEAs are calculated by plugging Eq. (4) into the skewness Eq. (7) and obtaining the first derivative of this expression with regard to the variable of interest. This derivative is then evaluated at the sample means of X .

When looking at the estimated coefficients for the one-step-ahead analysis ($h = 1$), the estimates broadly

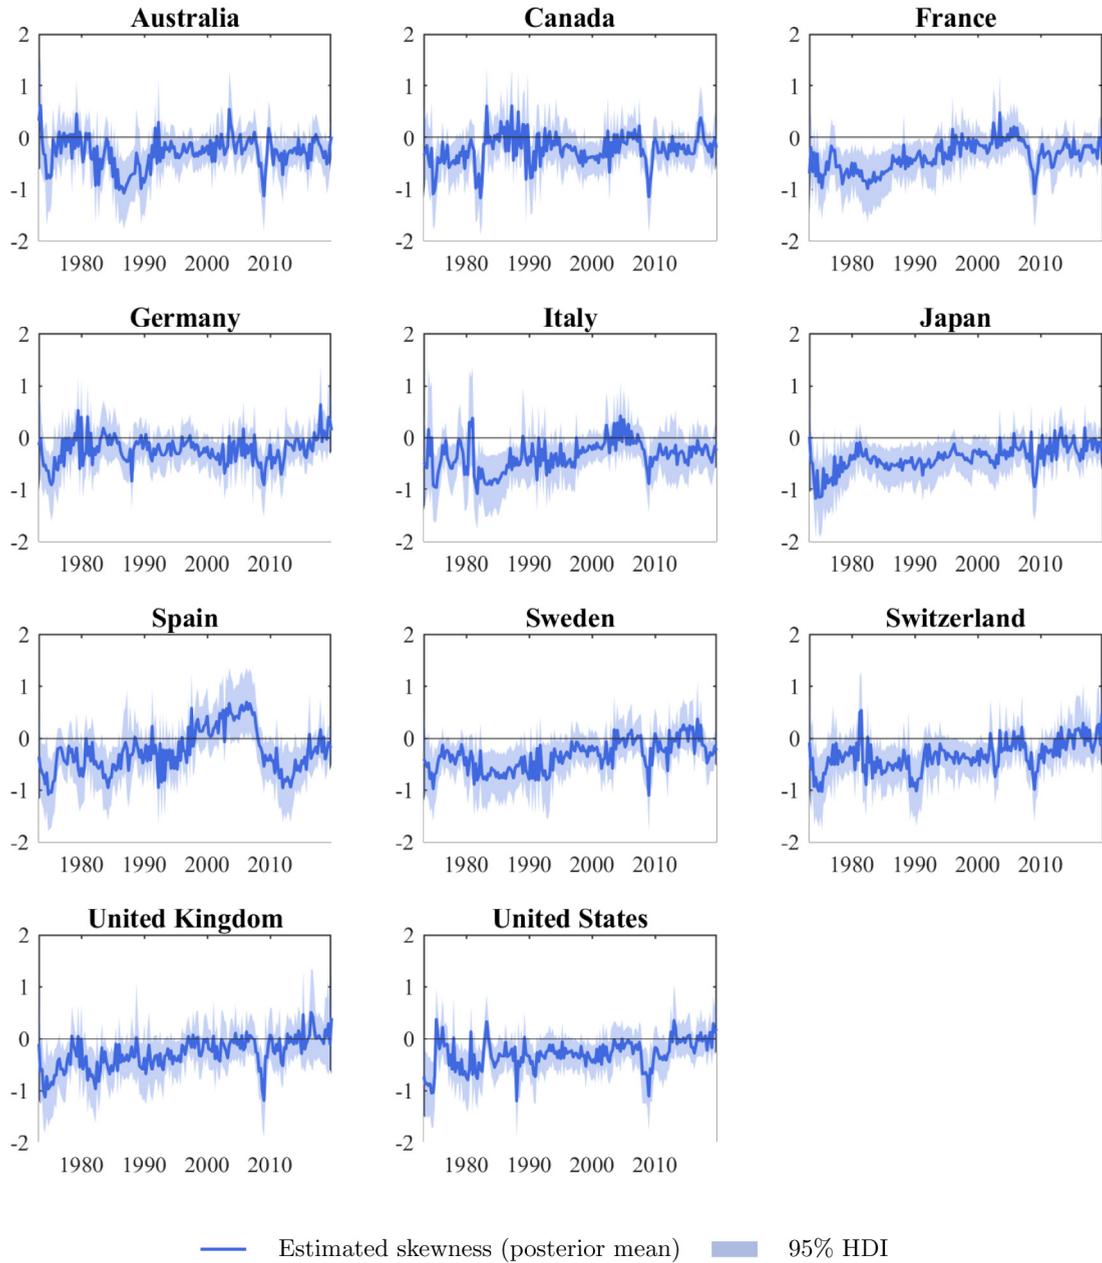


Fig. 2. Skewness of in-sample one-step-ahead predictive distribution.

confirm the findings of previous work and suggest an important role of financial conditions in shaping the growth distribution. Worsening financial conditions, as reflected by an increase in the FCI, are linked to short-term downside risk; i.e. the predictive growth distribution becomes skewed to the left (Adrian et al., 2019). Overall, the results on the remaining variables are ambiguous but do not generally indicate that they play any prominent role in explaining downside risk. However, this does not imply

that they are irrelevant, as the term spread and house price growth are also among the many factors from which the FCI is extracted. Based on the prominent role that these variables play in a policy environment, coupled with the findings in Brownlees and Souza (2021), we therefore also include them as separate determinants.

At the short horizon ($h = 1$), the term spread does not seem to predict downside risk. Interestingly, in the very near term, house price increases seem to go along

with upside risks to the economy, indicating that such increases are usually a phenomenon during expansionary periods. As mentioned above, the overall signal of these variables is difficult to assess, due to their additional indirect impact through the FCI. The results on economic and policy uncertainty are counter-intuitive, as a more uncertain environment in the current quarter relates to a somewhat higher chance of an upside growth surprise in the next quarter. However, posterior dispersion is very large. This broadly aligns with the results reported by Hengge (2019), showing that economic policy uncertainty has limited power to explain growth vulnerabilities.

At the one-year-ahead horizon ($h = 4$), tightening financial conditions, as measured by the FCI, no longer appear to signal downside risks. The posterior mean changes sign while the 95% posterior interval includes zero. At the two-year-ahead horizon ($h = 8$), this reversed effect becomes stronger and the posterior interval no longer includes zero. Qualitatively, these results support the findings of Adrian et al. (2021), i.e. the existence of a tradeoff between short-term benefits (risks) and medium-term risks (benefits) when loosening (tightening) financial conditions. In addition, with an increasing horizon, the term spread seems to play a more important role in predicting risks, and the posterior mean of the coefficient has the expected sign. This supports the findings of Estrella and Hardouvelis (1991), who show that the slope of the yield curve can help to predict recessions around one to two years ahead. Again, the posterior distribution remains relatively wide. Finally, asymmetries of the predictive growth distribution appear to be moderately negatively autocorrelated. In line with Fig. 2, the estimated intercept of the asymmetry process is clearly negative.

Growth-at-risk and expected shortfall/longrise

Throughout the paper, we focus on the 5% (95%) quantile to analyze downside (upside) risk. Formally defined, GaR is the p %-quantile of the predictive growth distribution,¹²

$$\Pr(y_{i,t+h} \leq \text{GaR}_{i,t+h|t}^p) = p. \quad (14)$$

Based on these quantiles, ES/EL are then defined as the conditional expectation of the distribution beyond the GaR level,

$$\begin{aligned} \text{ES}_{i,t+h|t}^p &= \mathbb{E}(y_{i,t+h} | y_{i,t+h} \leq \text{GaR}_{i,t+h|t}^p), \\ \text{EL}_{i,t+h|t}^p &= \mathbb{E}(y_{i,t+h} | y_{i,t+h} \geq \text{GaR}_{i,t+h|t}^p). \end{aligned} \quad (15)$$

Fig. 3 shows the model-implied one-step-ahead Growth-at-Risk and expected shortfall/longrise values over time along with the realized growth rates. To generate these, in each MCMC iteration, we draw from the one-step-ahead growth distribution of each country and period, and compute the relevant quantiles (GaR) and expected values beyond these quantiles (ES/EL) using all these draws. This presentation of the results complements and merges the insights from the previously shown

volatility and skewness Figs. 1 and 2. Based on a visual inspection, the model seems to capture the dynamics of the conditional GDP growth distribution appropriately in most countries. While both upside and downside risk vary over time in most countries, downside risk seems to be generally more volatile, which is in line with the results reported in Adrian et al. (2019).¹³

It is worth discussing this point a bit further. In the original Adrian et al. (2019) paper that uses quantile regression, a larger variability of downside risk relative to upside risk emerges since financial conditions affect the lower quantiles of the predictive growth distribution more than the upper ones. In our modeling approach, this result stems from the fact that financial conditions directly affect the skewness of the predictive growth distribution. Notably, Carriero et al. (2020a) show that obtaining this result does not require a model that produces asymmetric conditional growth distributions. In particular, a symmetric VAR model with time-varying volatility that allows for simultaneous shifts in the conditional mean and variance can also produce this result. They note that including financial conditions in the system is still crucial for observing the larger variability of downside risk than of upside risk.

However, the possibility to observe this one particular pattern in both models with symmetric and asymmetric conditional distributions does not imply that the latter is an irrelevant characteristic when trying to fit empirical models to GDP growth data. Specifically, the next section focuses on whether the specification suggested in this paper, which models the conditional skewness of the growth distribution as a function of macro-financial conditions, can help to improve the out-of-sample forecasting accuracy compared to alternative approaches.

Accuracy measures for quantile and expected shortfall forecasts

Table 1 also reports various measures to assess the accuracy of GaR and ES/EL predictions. These accuracy measures are mostly used in the out-of-sample forecasting exercise described in the next section and are reported here to allow for a comparison of both in-sample and out-of-sample results.

To assess the accuracy of the GaR and ES/EL forecasts, we rely on different measures that are regularly used in the literature on backtesting quantile and expected shortfall forecasts. First, to assess the accuracy of GaR forecasts, we report two versions of the dynamic quantile test of Engle and Manganelli (2004). The first, DQ_{uc} , tests whether the GaR forecasts have unconditionally correct coverage, i.e. whether the share of GaR violations equals the nominal coverage. The second, DQ_{hits} , tests whether GaR violations are optimal when including lagged violations in the test equation, i.e. whether the violations are independent. These tests were also used by Brownlees and Souza (2021), and we refer the reader there for a more detailed explanation.

¹² While GaR usually refers to the lower quantile, for simplicity, here we also use the label for the upper quantile.

¹³ The average standard deviation across countries for downside (upside) risk, as shown in Fig. 3, is 0.59 (0.52) and 0.78 (0.64) for Growth-at-Risk and expected shortfall/longrise, respectively.

Table 1
In-sample posterior parameter estimates and accuracy measures.

		h = 1			h = 4			h = 8			
		Mean	Perc _{2.5}	Perc _{97.5}	Mean	Perc _{2.5}	Perc _{97.5}	Mean	Perc _{2.5}	Perc _{97.5}	
FCI	β	-0.85	-1.49	-0.25	0.30	-0.42	0.99	0.76	0.07	1.74	
	MEA	-1.00	-2.70	-0.20	0.11	-0.13	0.51	0.42	0.02	1.31	
TS	β	0.01	-0.46	0.45	0.40	-1.28	1.40	0.50	-0.29	1.33	
	MEA	0.04	-0.52	0.72	0.19	-0.21	0.84	0.29	-0.09	1.04	
HP	β	0.31	-0.21	0.82	0.37	-0.44	1.29	0.01	-0.65	1.06	
	MEA	0.42	-0.17	1.57	0.13	-0.13	0.58	-0.03	-0.49	0.39	
WUI	β	0.34	-0.14	0.91	0.52	-0.27	1.57	0.04	-0.53	0.74	
	MEA	0.37	-0.18	1.18	0.14	-0.13	0.45	0.00	-0.41	0.34	
δ_{t+h-1}	ϕ	-0.39	-0.78	0.13	-0.25	-0.80	0.43	-0.21	-0.76	0.33	
Const.	β	-1.05	-2.07	-0.40	-1.47	-3.24	-0.50	-1.41	-2.89	-0.47	
		σ_h^2	0.040	0.024	0.062	0.069	0.042	0.102	0.057	0.032	0.089
		σ_δ^2	0.016	0.004	0.055	0.017	0.004	0.061	0.017	0.004	0.068
		ν	9.85	6.19	16.58	19.56	9.10	29.53	15.45	8.10	28.15
			DQ _{uc}	DQ _{hits}	TL	DQ _{uc}	DQ _{hits}	TL	DQ _{uc}	DQ _{hits}	TL
Growth-at-Risk _{5%}			100	100	0.076	100	91	0.079	91	100	0.080
Growth-at-Risk _{95%}			100	100	0.067	82	91	0.069	100	100	0.068
			EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
Expected shortfall _{5%}			0.171	0.355		0.177	0.382		0.173	0.387	
Expected longrise _{95%}			0.147	0.620		0.215	0.629		0.191	0.629	

Note: This table contains the means and percentiles of the parameters' posterior distributions. 'MEA' refers to the marginal effect of an explanatory variable on the skewness of the predictive distribution evaluated at the average values of the remaining regressors. Under 'DQ_{uc}' and 'DQ_{hits}', we report the share of country series for which adequacy of the quantile forecasts is not rejected at the 5% level using two versions of the dynamic quantile test developed by [Engle and Manganelli \(2004\)](#) (for details see also [Brownlees & Souza, 2021](#)). 'TL' is the tick loss, 'EKP' refers to the expected shortfall precision measure of [Embrechts, Kaufmann, and Patie \(2005\)](#), and the VaR-ES score is the measure developed by [Fissler, Ziegel, and Gneiting \(2015\)](#).

To further quantitatively assess the results, we report the tick loss, which is the standard loss function to evaluate quantile estimates. In particular, the values reported in the tables are the average tick loss across countries, computed as

$$TL_p = \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{T} \sum_{t=1}^T (y_{i,t+h} - GaR_{i,t+h|t}^p) (p - \mathbb{I}_{\{y_{i,t+h} < GaR_{i,t+h|t}^p\}}) \right). \quad (16)$$

When evaluating expected shortfall/longrise predictions, it needs to be noted that the accuracy of these predictions inherently depends on the accuracy of the corresponding quantile predictions. While expected shortfall lacks the so-called elicibility property ([Fissler & Ziegel, 2016](#)), quantile and shortfall forecasts can be evaluated jointly. Specifically, as in [Carriero et al. \(2020a\)](#), we rely on the VaR-ES score of [Fissler et al. \(2015\)](#):

$$\begin{aligned} VaR-ES \text{ score}_p = & \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{T} \sum_{t=1}^T ((GaR_{i,t+h|t}^p (\mathbb{I}_{\{y_{i,t+h} < GaR_{i,t+h|t}^p\}} - p) \right. \\ & - y_{i,t+h} \mathbb{I}_{\{y_{i,t+h} < GaR_{i,t+h|t}^p\}} + \frac{e^{ES_{i,t+h|t}^p}}{1 + e^{ES_{i,t+h|t}^p}} (ES_{i,t+h|t}^p - GaR_{i,t+h|t}^p) \\ & \left. + p^{-1} (GaR_{i,t+h|t}^p - y_{i,t+h} \mathbb{I}_{\{y_{i,t+h} < GaR_{i,t+h|t}^p\}}) + \ln \frac{2}{1 + e^{ES_{i,t+h|t}^p}}) \right), \end{aligned} \quad (17)$$

where a smaller value indicates a better joint GaR/ES forecast. To compute this loss measure for the upper quantile, we follow [Carriero et al. \(2020a\)](#) and multiply the quantile and longrise series, as well as the data by -1, and apply the formula for the 5%-quantile.

Finally, as a second approach to evaluate the expected shortfall/longrise predictions, we report the measure developed by [Embrechts et al. \(2005\)](#). This measure has been used by [Nakajima \(2013\)](#) and [Iseringhausen \(2020\)](#) and is explained in more detail in those references. In particular, it includes a penalty term for the accuracy of the shortfall/longrise forecast, depending on the precision of the quantile forecast. We report the average across countries, where a smaller value of the *EKP* measure indicates a more precise ES/EL forecast.

3.3. Out-of-sample forecasting of macroeconomic risks

This section uses the model proposed in this paper along with existing approaches from the literature to forecast Growth-at-Risk and expected shortfall/longrise out-of-sample. We estimate the models initially over the period from 1973:Q1–1984:Q4 and then generate the *h*-step-ahead, $h \in \{1, \dots, 4\}$, $GaR^p/ES^p/EL^p$ forecasts.¹⁴ Starting from 1985:Q1, we then re-estimate the models by extending the in-sample period by *N* observations in each quarter and generating a new set of forecasts. This recursive procedure is continued until the end of the sample. As a result, we obtain for each combination of model, forecasting horizon, and quantile level a country-specific set of *GaR* and *ES/EL* forecasts of size $T_{oos}^h = 140 - (h - 1)$.

¹⁴ As in [Brownlees and Souza \(2021\)](#) and [Carriero et al. \(2020a\)](#), for the out-of-sample exercise, we do not consider horizons beyond one year. In addition, we do not use real-time vintages of growth data. Further, the FCIs are not constructed in real time and include a look-ahead bias ([Brownlees & Souza, 2021](#)).

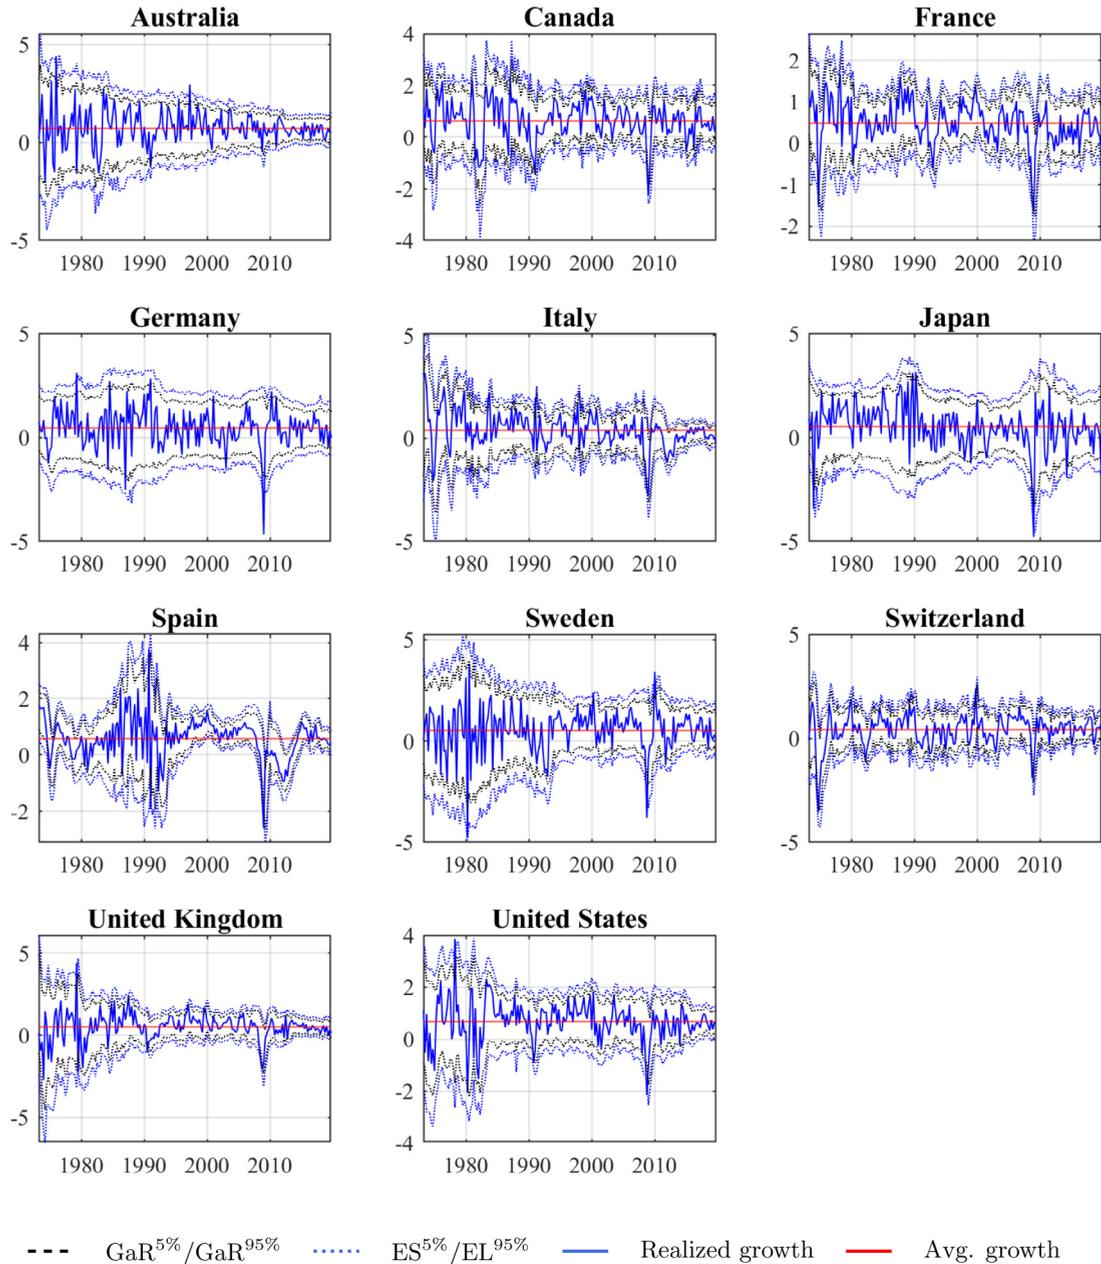


Fig. 3. In-sample one-step-ahead Growth-at-Risk and expected shortfall/longrise.

Competing models

The competitors used in the out-of-sample forecasting exercise are prominent models discussed in the literature: the quantile regression approach put forward by [Adrian et al. \(2019\)](#) and one of the generalized autoregressive conditional heteroskedasticity (GARCH) models applied by [Brownlees and Souza \(2021\)](#).¹⁵ Lastly, we include as a

benchmark the 'naive' historical forecast and a symmetric version of the model proposed in this paper, i.e. one without skewness. The conditional mean for the GARCH and SV models (both with and without time-varying skewness) includes a constant, current GDP growth, and three additional lags. Section 3.4 discusses the forecasting results across models when, in addition, including financial conditions in the conditional mean specification.

¹⁵ The robustness Section 3.4 also discusses the performance of another conditionally heteroskedastic model ([Adrian et al., 2019](#)), in which both the conditional mean and volatility are driven by financial

conditions. For all of these models, we use the MATLAB routines made available by [Adrian et al. \(2019\)](#) and [Brownlees and Souza \(2021\)](#).

Historical benchmark

As a simple benchmark to forecast Growth-at-Risk and expected shortfall/longrise, we use the unconditional historical measure. In this case, the h -step-ahead GaR and ES/EL forecasts for period $t + h$ are simply the empirical quantile and the empirical expected shortfall/longrise, respectively, calculated using all growth observations of a particular country available up to period t .

Quantile regression

Growth-at-Risk analysis through quantile regression was popularized by [Adrian et al. \(2019\)](#) and has become a regular monitoring exercise at the IMF and other organizations. Quantile regression was developed by [Koenker and Bassett \(1978\)](#) and we refer the reader to this work for details on the estimation. It can be viewed as a generalization of a standard (mean) regression where the p %-quantile of the dependent variable, $Q_p(y)$, is directly modeled as a function of explanatory variables:

$$Q_p(y_{i,t+h|t}) = \alpha_{0i}^p + \alpha_{1i}^p y_{it} + \alpha_{2i}^p FCI_{it}. \tag{18}$$

The baseline quantile regressions for the forecasting exercise include a constant, current growth, and the FCI.¹⁶ The corresponding country-specific quantile regression coefficients α_{0i}^p , α_{1i}^p , and α_{2i}^p can vary across quantiles. While quantile regression directly delivers GaR forecasts, we apply the two-step approach of [Adrian et al. \(2019\)](#) to obtain ES/EL forecasts. Specifically, [Adrian et al. \(2019\)](#) fit the skewed t-distribution of [Azzalini and Capitanio \(2003\)](#) over the predicted 5%, 25%, 75%, and 95% quantiles. The forecasts for ES/EL are then calculated using the cumulative distribution function of the fitted skewed t-distribution. In considering country-specific quantile regressions, we follow [Brownlees and Souza \(2021\)](#), who find these to generally have better forecasting performance than their panel counterparts.

Panel-GARCH

This model is one of the preferred specifications discussed by [Brownlees and Souza \(2021\)](#). In essence, their approach is a GARCH(1,1) model with a flexible non-parametric form used to model the standardized growth distribution. The model has the following form:

$$y_{i,t+1} = \mu_{i,t+1|t} + \sqrt{\sigma_{i,t+1|t}^2} \varepsilon_{i,t+1},$$

$$\varepsilon_{i,t+1} \sim \mathcal{D}_{\varepsilon_i}(\mathbf{0}, 1), \tag{19}$$

$$\sigma_{i,t+1|t}^2 = \sigma_i^2(1 - \alpha - \beta) + \alpha u_{it}^2 + \beta \sigma_{it}^2,$$

$$\alpha, \beta > 0, (\alpha + \beta) < 1, \tag{20}$$

where $\sigma_{i,t+1|t}^2$ is the conditional (deterministic) one-step-ahead variance, σ_i^2 is the unconditional variance, α and β are the GARCH parameters pooled across countries, and $u_{it} = \sigma_{it} \varepsilon_{it}$ is the non-standardized residual. The model is estimated by so-called composite likelihood methods, following [Pakel, Shephard, and Sheppard \(2011\)](#). Since the h -step-ahead predictive distribution (for $h > 1$) is not available in closed form, the authors rely on bootstrap

techniques to generate a large number of iterated one-step-ahead forecast paths. The forecasts of GaR and ES/EL at each horizon are then calculated using these paths. Note that the GARCH approach relies on iterated one-step-ahead forecasts for horizons $h > 1$ (including for the conditional mean) while both the quantile regressions and the SV models are used in a direct multi-step-ahead setting. For further methodological details, we refer the reader to [Brownlees and Souza \(2021\)](#).

SV models with time-varying skewness and with symmetric (t-distributed) shocks

For the model specification with time-varying skewness, we consider two sets of explanatory variables: the full set discussed in previous sections (SV (*skew, full*)) and, motivated by the in-sample analysis, a set that only includes the FCI (SV (*skew, FCI*)). Both specifications include a constant and an autoregressive term in the asymmetry equation.¹⁷

The general approach to GaR and ES/EL forecasting is identical for all stochastic volatility specifications. To generate h -quarter-ahead GaR and ES/EL forecasts, we obtain the h -quarter-ahead forecasts of the conditional mean and the latent variables h and δ for each country. Since (log-)volatility is assumed to follow a random walk, the forecast for h is simply, in each MCMC iteration, the draw for the last in-sample period. To obtain a forecast of δ , we iterate forward Eq. (4) by h quarters, starting by using the draw for the last in-sample period together with the $(h - 1)$ -period-lagged explanatory variables. This process is continued to obtain the forecast of δ for $h > 1$. Importantly, since X_δ in Eq. (4) is lagged by h quarters, we only use information available at the end of the in-sample period to obtain the h -quarter-ahead forecast of δ . Using the forecasts of the conditional mean and latent variables, we then generate a draw from the corresponding predictive density in each MCMC iteration and for each country. GaR and ES/EL are then estimated from these draws.

Finally, to assess the importance of the information contained in the variables in the asymmetry equation, we also consider a simplified version of the time-varying skewness model (SV (*sym.*)). This model is an SV model with t-distributed growth shocks obtained by imposing the restriction $\delta = 0$ for each country.

Forecasting results

The results of the out-of-sample forecasting exercise for horizons up to four quarters ahead are presented in [Table 2](#). We start by assessing the GaR forecasts. First, in terms of the DQ tests, quantile regression clearly outperforms the historical benchmark at the upper quantile. However, this is not the case for the lower quantile. By contrast, and in line with the results in [Brownlees and Souza \(2021\)](#), the GARCH model generally outperforms both the historical benchmark and quantile regression for the lower and the upper quantiles. Lastly, all three

¹⁶ Quantile regressions that include all the explanatory variables generally perform worse at predicting downside risk, and these results are discussed in the robustness section.

¹⁷ For computational reasons, in the forecasting exercise the Gibbs sampler for all SV models is only run with 1,000,000 iterations and saving every 10th iteration, thus leaving the number of effective draws unchanged.

Table 2
Results of out-of-sample forecasting exercise.

h	Model	GaR and ES/EL (5%)					GaR and ES/EL (95%)				
		DQ _{lit}	DQ _{hits}	TL	VaR-ES sc.	EKP	DQ _{lit}	DQ _{hits}	TL	VaR-ES sc.	EKP
1	Historical	73	18	0.101	0.575	0.452	45	64	0.081	0.690	0.446
	QR (FCI)	73	45	0.090	0.497	0.276	91	82	0.079	0.676	0.272
	GARCH(1,1)	91	64	0.086	0.458	0.211	100	82	0.071	0.653	0.290
	SV (sym.)	100	64	0.083	0.427	0.203	100	100	0.068	0.628	0.174
	SV (skew, FCI)	100	73	0.082	0.418	0.162	100	100	0.068	0.628	0.152
	SV (skew, full)	100	64	0.082	0.417	0.190	100	100	0.068	0.628	0.152
2	Historical	73	82	0.102	0.588	0.455	45	73	0.081	0.691	0.450
	QR (FCI)	82	55	0.096	0.538	0.387	64	73	0.080	0.698	0.334
	GARCH(1,1)	82	91	0.095	0.535	0.409	82	91	0.072	0.657	0.316
	SV (sym.)	100	73	0.092	0.520	0.303	100	82	0.070	0.639	0.203
	SV (skew, FCI)	100	82	0.092	0.515	0.283	100	82	0.070	0.637	0.174
	SV (skew, full)	91	82	0.092	0.520	0.317	91	91	0.071	0.642	0.195
3	Historical	73	82	0.103	0.593	0.470	45	64	0.082	0.691	0.454
	QR (FCI)	64	73	0.102	0.622	0.589	64	91	0.082	0.697	0.240
	GARCH(1,1)	91	100	0.102	0.617	0.462	82	82	0.077	0.669	0.339
	SV (sym.)	91	82	0.100	0.616	0.439	100	82	0.074	0.656	0.178
	SV (skew, FCI)	100	91	0.100	0.611	0.416	100	82	0.074	0.654	0.160
	SV (skew, full)	100	82	0.099	0.612	0.449	100	82	0.075	0.657	0.184
4	Historical	73	82	0.103	0.596	0.473	45	82	0.082	0.692	0.457
	QR (FCI)	64	64	0.109	0.701	0.519	64	82	0.082	0.693	0.250
	GARCH(1,1)	91	100	0.105	0.655	0.511	64	73	0.080	0.682	0.412
	SV (sym.)	91	100	0.102	0.665	0.439	100	73	0.077	0.666	0.202
	SV (skew, FCI)	91	100	0.103	0.667	0.429	100	73	0.077	0.664	0.184
	SV (skew, full)	91	82	0.102	0.662	0.448	82	73	0.078	0.668	0.200

Note: This table contains the results of the out-of-sample forecasting exercise. Under 'DQ_{lit}' and 'DQ_{hits}', we report the share of country series for which adequacy of the quantile forecasts is not rejected at the 5% level using two versions of the dynamic quantile test developed by Engle and Manganelli (2004). For details, see also (Brownlees & Souza, 2021). 'TL' is the tick loss, the VaR-ES score refers to the measure developed by Fissler et al. (2015), and 'EKP' is the expected shortfall precision measure of Embrechts et al. (2005). Bold numbers indicate the model with the highest average accuracy for each measure.

stochastic volatility models perform similarly, on average, compared to the GARCH approach, with only small differences among them.

Second, in terms of the tick loss (TL), the stochastic volatility models are the best-performing models across quantiles and horizons. However, the differences in average tick losses are very small across the stochastic volatility specifications.

Third, when evaluating quantile and shortfall/longrise forecasts simultaneously using the VaR-ES score, the SV models with time-varying skewness prove most successful at short horizons up to $h = 2$, and when forecasting upside risk at horizons $h > 1$. For $h > 2$, the historical benchmark dominates for the left tail of the distribution, highlighting that precise forecasts of downside risk at longer horizons are difficult for econometric models. In addition, adding other variables next to the FCI does not meaningfully improve the average forecasting accuracy.

Fourth, when evaluating the expected shortfall/longrise forecasts using the measure of Embrechts et al. (2005), financial conditions help to improve forecasts of downside risk and upside risk across horizons. For $h = 1$, the EKP value of the time-varying skewness model that only includes the FCI is clearly smaller than those of the strongest competitors when measuring downside risk. Again, including variables beyond the FCI does not add much value in terms of shortfall/longrise forecasts when looking at the EKP measure.

Table 3 complements the results presented in Table 2 by showing the outcome of tests for superior forecasting

performance (Diebold & Mariano, 1995).¹⁸ Specifically, as in Brownlees and Souza (2021), we conduct pair-wise tests with the null hypothesis being equal predictive ability against a one-sided alternative. The tests are based on the series of loss differences using two different models, defined as the differences in the tick loss (for GaR) and the VaR-ES score (for ES/EL). We note upfront that the existing literature often finds limited evidence of statistically significant superior predictive ability for Growth-at-Risk and expected shortfall/longrise forecasts (Brownlees & Souza, 2021; Carriero et al., 2020a). In our analysis, we find that for downside risk at the short horizon ($h = 1$), the SV model with time-varying skewness ('FCI only') is superior to quantile regression and the GARCH model for 3–4 (out of 11) countries, and for 1–3 countries when evaluated against the symmetric SV model. In turn, our proposed model is outperformed by other models for at most 1–2 countries. Overall, and in line with Brownlees and Souza (2021), the evidence becomes weaker as h increases, and for $h > 2$, most models have difficulty outperforming even the historical benchmark for more than 1–2 countries. When predicting upside risk, the evidence of superior forecasting accuracy is stronger across horizons, and the time-varying skewness model that includes

¹⁸ The various SV models are nested and the test can be invalid in this case (see e.g. Clark & McCracken, 2001). However, it is often also used for nested models, due to the complexity of alternative approaches.

Table 3
Tests for superior forecasting performance.

		Downside risk (5%)						Upside risk (95%)					
		Tick loss						VaR-ES score					
h		Hist.	QR	GARCH	SV (sym.)	SV (FCI)	SV (full)	Hist.	QR	GARCH	SV (sym.)	SV (FCI)	SV (full)
1	Hist.	-	36 (36)	55 (55)	55 (73)	64 (73)	64 (73)	-	36 (45)	36 (55)	55 (55)	55 (64)	55 (55)
	QR	0 (0)	-	27 (27)	27 (27)	27 (27)	27 (27)	0 (0)	-	18 (27)	27 (45)	36 (45)	36 (45)
	GARCH	0 (0)	9 (18)	-	27 (27)	27 (36)	36 (45)	0 (0)	9 (9)	-	27 (45)	36 (36)	36 (45)
	SV (sym.)	0 (0)	9 (18)	0 (9)	-	9 (27)	9 (18)	0 (0)	0 (18)	0 (9)	-	0 (18)	9 (18)
	SV (FCI)	0 (0)	9 (18)	0 (0)	0 (9)	-	9 (9)	0 (0)	0 (9)	0 (0)	0 (0)	-	0 (9)
	SV (full)	0 (0)	9 (18)	0 (0)	0 (0)	0 (0)	-	0 (0)	9 (18)	0 (0)	0 (0)	0 (0)	-
2	Hist.	-	36 (36)	45 (64)	36 (55)	45 (55)	36 (55)	-	18 (18)	18 (36)	18 (36)	18 (36)	18 (36)
	QR	0 (0)	-	18 (27)	27 (27)	27 (27)	18 (27)	0 (0)	-	18 (27)	18 (18)	18 (18)	18 (18)
	GARCH	9 (9)	0 (9)	-	9 (18)	9 (27)	18 (18)	0 (0)	0 (9)	-	0 (9)	0 (18)	0 (9)
	SV (sym.)	9 (9)	0 (9)	9 (9)	-	0 (9)	9 (18)	0 (0)	0 (0)	0 (0)	-	0 (18)	0 (9)
	SV (FCI)	9 (9)	0 (9)	9 (9)	9 (9)	-	9 (18)	0 (0)	0 (0)	0 (0)	0 (0)	-	0 (0)
	SV (full)	9 (9)	0 (9)	9 (9)	0 (0)	9 (9)	-	0 (0)	0 (0)	0 (9)	0 (0)	9 (9)	-
3	Hist.	-	27 (36)	9 (9)	9 (9)	9 (9)	9 (27)	-	27 (36)	9 (9)	9 (9)	9 (9)	9 (9)
	QR	18 (27)	-	18 (27)	9 (18)	9 (18)	9 (18)	18 (36)	-	9 (27)	9 (9)	9 (18)	9 (9)
	GARCH	9 (18)	0 (0)	-	9 (36)	18 (27)	27 (27)	0 (18)	0 (0)	-	0 (9)	0 (0)	0 (0)
	SV (sym.)	9 (18)	0 (0)	9 (9)	-	0 (9)	9 (27)	0 (9)	0 (0)	0 (9)	-	0 (18)	0 (18)
	SV (FCI)	9 (18)	0 (0)	9 (9)	18 (18)	-	9 (36)	0 (9)	0 (0)	0 (9)	9 (18)	-	0 (9)
	SV (full)	9 (18)	0 (0)	9 (9)	9 (9)	0 (9)	-	0 (9)	0 (0)	0 (9)	9 (9)	0 (9)	-
4	Hist.	-	9 (9)	9 (9)	18 (18)	18 (18)	18 (18)	-	0 (18)	9 (9)	18 (18)	18 (18)	9 (18)
	QR	9 (36)	-	27 (36)	27 (36)	27 (36)	27 (36)	9 (45)	-	18 (27)	18 (18)	9 (18)	18 (18)
	GARCH	27 (45)	0 (0)	-	18 (27)	18 (27)	18 (27)	9 (18)	0 (0)	-	18 (18)	18 (18)	18 (18)
	SV (sym.)	18 (27)	0 (0)	9 (9)	-	0 (0)	0 (9)	9 (18)	0 (0)	0 (0)	-	0 (9)	0 (27)
	SV (FCI)	18 (27)	0 (0)	9 (9)	36 (45)	-	9 (27)	18 (18)	0 (0)	0 (0)	27 (27)	-	9 (18)
	SV (full)	27 (27)	0 (0)	9 (9)	0 (0)	0 (0)	-	0 (18)	0 (0)	0 (0)	0 (0)	0 (0)	-

Note: This table contains the results of Diebold and Mariano (1995) tests (using (Newey & West, 1987) HAC standard errors with lag truncation $h = 1$) to compare the forecasts generated by the different models. Similarly to Brownlees and Souza (2021), we report the share of countries for which the model in the column produces more precise Growth-at-Risk and expected shortfall/longrise forecasts than the model in the respective row at the 5% (10%) significance level (one-sided test).

only the FCI performs particularly well. While these results are largely consistent with the cross-country average measures presented in Table 2, especially among the stochastic volatility models, the differences in the average losses are often not large enough to achieve statistical significance.

In summary, the results confirm the previous literature in some ways while adding new insights in others. In line with Brownlees and Souza (2021), we find commonly used quantile regression to forecast poorly compared to standard time-varying volatility models, such as a GARCH-type specification that only requires data on economic growth as input. However, a relatively simple specification from the family of symmetric stochastic volatility models seems to perform even better than the GARCH model.

When comparing SV models with standard quantile regressions, Carriero, Clark, and Marcellino (2020b) also

find that the former perform significantly better at out-of-sample forecasting.¹⁹ Importantly from a policy perspective, we find that a time-varying skewness model that includes the FCI can, in some cases, help to improve average forecasts of macroeconomic risk. However, the gains compared to the symmetric SV model—and, to a lesser extent, the GARCH model—are often small. A larger model that also includes the term spread, house price growth, and an index of economic and policy uncertainty occasionally helps to improve single measures at some horizons but overall does not perform better than the ‘FCI only’ specification.

¹⁹ Carriero et al. (2020b) include additional variables in the conditional mean specification of the SV model. Our results suggest that even an SV model that only includes autoregressive terms in the mean equation outperforms quantile regressions.

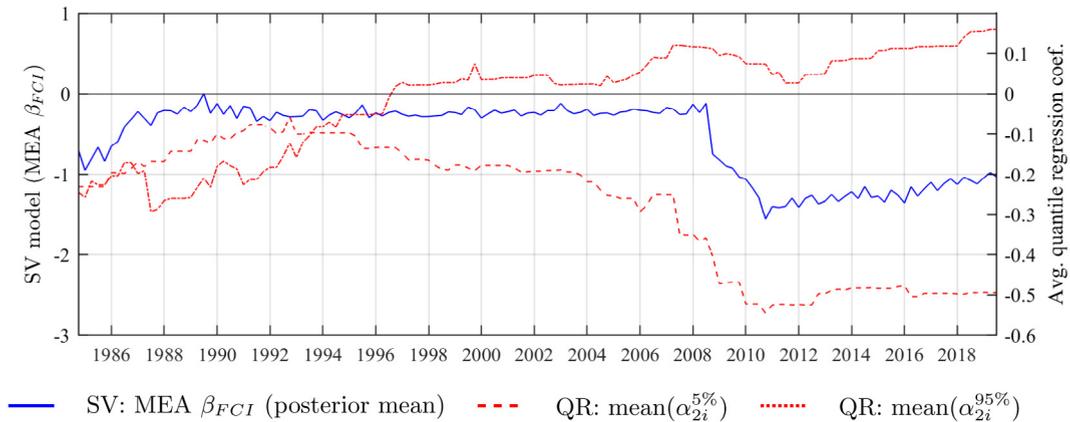


Fig. 4. Recursive estimates of the impact of FCI on asymmetry (one quarter ahead). Note: The dates refer to the last observation of the respective estimation (in-sample) period.

Appendix E contains figures showing the out-of-sample one-step-ahead GaR and ES/EL forecasts obtained from quantile regressions, panel-GARCH, and the SV model with time-varying skewness that only includes the FCI. Overall, the two time-varying volatility models (GARCH and SV) produce one-step-ahead forecasts of upside and downside risk that look generally similar across countries. However, the forecasts from quantile regressions are quite different and can be wider (e.g. in the case of Australia) or narrower (e.g. in the case of Japan), and are often more erratic. The latter confirms the findings by [Carriero et al. \(2020a\)](#) for the US, and this is not surprising, as the forecasts from the time-varying volatility models have a stronger dependence, due to the assumptions on the respective volatility processes.

To conclude the forecasting section, [Fig. 4](#) evaluates the time variation in the impact of financial conditions on the asymmetry of one-step-ahead conditional growth distributions. Specifically, we plot the recursive estimate of the marginal effect at the average (MEA) for the FCI from the ‘FCI only’ stochastic volatility model, computed as explained in [Section 3.2](#). We also plot the cross-country average quantile regression coefficients for the FCI and the 5% and 95% quantiles. First, the marginal effect of an increase of the FCI on skewness in the SV model is negative for all recursive estimations. Second, the Great Recession marks a structural break, after which the impact of financial conditions is estimated to be much stronger. Third, a similar pattern can be observed for the average estimated 5%-quantile regression coefficient. While the magnitudes are not directly comparable, the average 5%-quantile coefficient of the FCI co-moves closely with the MEA obtained from the time-varying skewness model. Interestingly, according to the quantile regression estimates, in the first part of the sample, financial conditions seem to have had, on average, a stronger impact on the upper quantile than the lower quantile. Future work could more explicitly consider this apparent time variation in

the effect of financial conditions on the asymmetry of predictive growth distributions.

3.4. Robustness checks and alternative modeling choices

This section reports additional results with a focus on the out-of-sample forecasting performance to assess the robustness of the baseline results and to discuss the impact of alternative modeling choices. The detailed results are presented in [Table E-1](#) in [Appendix E](#).

First, to ensure that skewness, which is driven by macro-financial conditions, is a useful feature for forecasting, we consider variations of the different models where the FCI is also included in the conditional mean (with country-specific coefficients). Specifically, we augment the previously considered GARCH model, the symmetric SV model, and the SV model with time-varying skewness driven by the FCI.²⁰ Moreover, for this exercise we also consider an SV model with a constant degree of skewness, where the asymmetry parameter δ is pooled across countries, i.e. where $\delta_{it} = \delta$. Including this model offers additional insights into the relevance of time-varying skewness in addition to the FCI driving the conditional mean. Allowing financial conditions to also impact the conditional mean overall improves the average forecasting performance further for $h = 1$ in terms of the loss measures. At the remaining horizons, accuracy gains mostly occur when predicting downside risk, whereas precision in the right tail of the distribution decreases on average. When forecasting downside risk for the next quarter, the SV model with time-varying skewness and the FCI in the mean equation maintains, on average, small advantages compared to the remaining

²⁰ Since the GARCH model relies on iterated one-step-ahead forecasts for $h > 1$, including for the conditional mean (see [Section 3.3](#)), this requires forecasts of the exogenous variables. We obtain these forecasts from an AR(1) model fitted to the FCI of each country.

models. At longer horizons and when forecasting upside risk, the evidence is more mixed. In particular, the constant skewness specification with the 'FCI-in-mean' often performs very well. In addition, the 'FCI-in-mean' GARCH model proves useful for downside risk at horizons $h > 2$, but the average VaR-ES scores remain above those of the historical benchmark (see Table 2). Table E-2 shows the results of tests for superior predictive accuracy for the 'FCI-in-mean' augmented specifications, which overall confirm this assessment.

Second, the quantile regression considered as a competitor model is close to the original specification of Adrian et al. (2019) and only includes the financial conditions index. We also consider a quantile regression specification that includes the full set of explanatory variables. However, the results indicate that in most cases this leads to a deterioration in forecasting performance compared to the more parsimonious specification.

Third, to allow for an effect of current growth rates on future asymmetry, we consider a specification where y_t enters the asymmetry equation as an additional variable. While there are some occasions where this can improve the accuracy of GaR and ES/EL forecasts, the results remain overall comparable to the baseline specification.

Fourth, we consider a specification that allows for more flexibility in the cross-country dynamics of volatility by allowing the innovation variance of the (log-)volatility process h , σ_h^2 , to vary across countries. Table E-1 shows that in some cases this helps both the symmetric SV model and the SV model with time-varying skewness to achieve gains in forecasting accuracy compared to their baseline (pooled) counterparts. Again, this does not change the ranking among the two models, and the time-varying skewness model still produces, on average, somewhat more precise risk forecasts.

Finally, we compare our model to another (country-specific) conditionally heteroskedastic model, in which the conditional mean and volatility are functions of current growth and the FCI.²¹ Table E-1 shows that in terms of both DQ statistics, this model performs worse overall compared to the baseline time-varying skewness model (with the FCI included in the mean equation). For the remaining loss measures, this simple model generates some gains when forecasting downside risk at longer horizons, but remains inferior overall at short horizons, and across horizons in the case of upside risk. In summary, while certain alternative modeling choices can help to further improve the performance compared to the baseline specification, these additional tests show that time-varying skewness driven directly by macro-financial conditions remains a useful feature in several cases.

²¹ Following Adrian et al. (2019), this model is given by

$$y_{i,t+h} = \gamma_{0i} + \gamma_{1i}y_{it} + \gamma_{2i}FCI_{it} + \sigma_{it}\varepsilon_{i,t+h},$$

$$\log(\sigma_{it}^2) = \delta_{0i} + \delta_{1i}y_{it} + \delta_{2i}FCI_{it},$$

with $\varepsilon_{i,t+h} \sim \mathcal{N}(0, 1)$. The model parameters can be estimated by maximum likelihood.

4. Conclusion

Economic policymakers have a long-standing interest in measuring and assessing downside risks to economic growth stemming from macro-financial conditions, and the academic literature has recently provided the Growth-at-Risk approach for this purpose. The study of growth vulnerability remains an active area of research, and several recent contributions have scrutinized various aspects of the Growth-at-Risk framework. In particular, the question of to what extent financial variables can help to inform this analysis and improve Growth-at-Risk forecasts has been extensively discussed.

This paper makes both a methodological and an empirical contribution to this discussion. On the methodological side, it proposes a new parametric model to measure the evolving asymmetry of the predictive GDP growth distribution, which can be interpreted as changing macroeconomic risk. The methodological basis for this approach is a stochastic volatility model in which the asymmetry parameter of the shock distribution varies as a function of macro-financial conditions. Thus, the model allows the conditional growth distribution to feature time-varying skewness, which reflects unbalanced risks surrounding the baseline macroeconomic outlook. On the empirical side, the paper supports earlier findings in the literature on the link between financial conditions and macroeconomic risk and suggests that the proposed model can, in several situations, help to achieve gains when forecasting risks.

For a panel of 11 OECD countries over the period from 1973:Q1–2019:Q4, this model provides the following insights. First, the estimated effect of financial conditions, as measured by the IMF's financial conditions index (FCI), on the skewness of the predictive growth distribution is in line with a growing number of studies in the literature. Tightening financial conditions skew the short-term predictive growth distribution to the left, but the opposite effect emerges at longer horizons. Second, when forecasting Growth-at-Risk and expected shortfall/longrise out-of-sample, the proposed model competes well with existing approaches and proves particularly useful at short horizons for downside risk, and at several horizons when forecasting upside risk. In addition, including a measure of economic and policy uncertainty, or some of the prominent components of the FCI separately, does not generally improve (or adds only very little to) the forecasting performance of the model.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.ijforecast.2023.02.006>.

References

- Adams, P. A., Adrian, T., Boyarchenko, N., & Giannone, D. (2021). Forecasting macroeconomic risks. *International Journal of Forecasting*, 37(3), 1173–1191.
- Adrian, T., Boyarchenko, N., & Giannone, D. (2019). Vulnerable growth. *American Economic Review*, 109(4), 1263–1289.
- Adrian, T., Grinberg, F., Liang, N., Malik, S., & Yu, J. (2021). The term structure of Growth-at-Risk. *American Economic Journal: Macroeconomics*, 24(3), 283–323.
- Ahir, H., Bloom, N., & Furceri, D. (2018). The world uncertainty index. mimeo.
- Alessandri, P., & Mumtaz, H. (2017). Financial conditions and density forecasts for US output and inflation. *Review of Economic Dynamics*, 24, 66–78.
- Antolin-Diaz, J., Drechsel, T., & Petrella, I. (2017). Tracking the slow-down in long-run GDP growth. *The Review of Economics and Statistics*, 99(2), 343–356.
- Azzalini, A., & Capitanio, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 65(2), 367–389.
- Brownlees, C. T., & Souza, A. (2021). Backtesting global Growth-at-Risk. *Journal of Monetary Economics*, 118, 312–330.
- Caldara, D., Cascaldi-Garcia, D., Cuba-Borda, P., & Loria, F. (2020). Understanding Growth-at-Risk: A Markov-switching approach. mimeo.
- Caldera Sánchez, A., & Röhn, O. (2016). *How do policies influence GDP tail risks?* OECD Economics Department working papers. No. 1339, OECD.
- Carriero, A., Clark, T. E., & Marcellino, M. G. (2020a). *Capturing macroeconomic tail risks with Bayesian vector autoregressions: Federal Reserve Bank of Cleveland working paper. No. 20-02R*, Federal Reserve Bank of Cleveland.
- Carriero, A., Clark, T. E., & Marcellino, M. G. (2020b). *Nowcasting tail risks to economic activity with many indicators: Federal Reserve Bank of Cleveland working paper. No. 20-13R*, Federal Reserve Bank of Cleveland.
- Carter, C. K., & Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3), 541–553.
- Chan, J. C., & Hsiao, C. Y. (2014). Estimation of stochastic volatility models with heavy tails and serial dependence. In I. Jeliazkov, & X.-S. Yang (Eds.), *Bayesian inference in the social sciences* (pp. 155–176). Wiley-Blackwell.
- Chan, J. C., & Jeliazkov, I. (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1–2), 101–120.
- Clark, T. E., & McCracken, M. W. (2001). Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, 105(1), 85–110.
- D'Agostino, A., Gambetti, L., & Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics*, 28(1), 82–101.
- Del Negro, M., & Otrok, C. (2008). *Dynamic factor models with time-varying parameters: Measuring changes in international business cycles: Staff Reports 326*, Federal Reserve Bank of New York.
- Delle Monache, D., De Polis, A., & Petrella, I. (2020). *Modeling and forecasting macroeconomic downside risk: CEPR discussion papers. DP15109*, Centre for Economic Policy Research.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Embrechts, P., Kaufmann, R., & Patie, P. (2005). Strategic long-term financial risks: Single risk factors. *Computational Optimization and Applications*, 32(1–2), 61–90.
- Engle, R. F., & Manganelli, S. (2004). CAViAR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4), 367–381.
- Estrella, A., & Hardouvelis, G. A. (1991). The term structure as a predictor of real economic activity. *The Journal of Finance*, 46(2), 555–576.
- Fissler, T., & Ziegel, J. (2016). Higher order elicibility and Osband's principle. *The Annals of Statistics*, 44(4), 1680–1707.
- Fissler, T., Ziegel, J., & Gneiting, T. (2015). Expected shortfall is jointly elicitable with value at risk – implications for backtesting. *Risk*.
- Gelman, A., & Shirley, K. (2011). Inference from simulations and monitoring convergence. In *Handbook of Markov chain Monte Carlo: vol. 6*, (pp. 163–174). CRC Press Boca Raton, FL.
- Giglio, S., Kelly, B., & Pruitt, S. (2016). Systemic risk and the macroeconomy: An empirical evaluation. *Journal of Financial Economics*, 119(3), 457–471.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.
- Hansen, B. E. (1994). Autoregressive conditional density estimation. *International Economic Review*, 35(3), 705–730.
- Harvey, C. R., & Siddique, A. (1999). Autoregressive conditional skewness. *Journal of Financial and Quantitative Analysis*, 34(4), 465–487.
- Hengge, M. (2019). *Uncertainty as a predictor of economic activity: Working paper series. HEIDWP19-2019*, Graduate Institute of International and Development Studies.
- Hogben, D., Pinkham, R., & Wilk, M. (1961). The moments of the non-central t-distribution. *Biometrika*, 48(3/4), 465–468.
- IMF (2017). *Is growth at risk?: Global financial stability report (Chapter 3), October 2017*, International Monetary Fund.
- IMF (2018). *A decade after the global financial crisis: Are we safer?: Global financial stability report (Chapter 1), October 2018*, International Monetary Fund.
- Iseringhausen, M. (2020). The time-varying asymmetry of exchange rate returns: A stochastic volatility – stochastic skewness model. *Journal of Empirical Finance*, 58, 275–292.
- Jensen, H., Petrella, I., Ravn, S. H., & Santoro, E. (2020). Leverage and deepening business-cycle skewness. *American Economic Journal: Macroeconomics*, 12(1), 245–281.
- Johnson, N. L., Kotz, S., & Balakrishnan, N. (1995). *Wiley series in probability and mathematical statistics: Applied probability and statistics: vol. 2, Continuous univariate distributions*. New York: Wiley.
- Jovanovic, B., & Ma, S. (2021). Uncertainty and growth disasters. *Review of Economic Dynamics*, 44, 33–64.
- Jurado, K., Ludvigson, S. C., & Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3), 1177–1216.
- Kim, S., Shephard, N., & Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3), 361–393.
- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50.
- Koop, G., & Korobilis, D. (2014). A new index of financial conditions. *European Economic Review*, 71, 101–116.
- Montes-Galdón, C., & Ortega, E. (2022). Skewed SVARS: Tracking the structural sources of macroeconomic tail risks. In *Essays in Honour of Fabio Canova (Advances in Econometrics, Vol. 44A)* (pp. 177–210). Emerald Publishing Limited.
- Morley, J., & Piger, J. (2012). The asymmetric business cycle. *The Review of Economics and Statistics*, 94(1), 208–221.
- Nakajima, J. (2013). Stochastic volatility model with regime-switching skewness in heavy-tailed errors for exchange rate returns. *Studies in Nonlinear Dynamics & Econometrics*, 17(5), 499–520.
- Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- Omori, Y., Chib, S., Shephard, N., & Nakajima, J. (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140(2), 425–449.
- Orlik, A., & Veldkamp, L. (2014). *Understanding uncertainty shocks and the role of black swans: NBER working papers. No. 20445*, National Bureau of Economic Research.
- Pakel, C., Shephard, N., & Sheppard, K. (2011). Nuisance parameters, composite likelihoods and a panel of GARCH models. *Statistica Sinica*, 21(1), 307–329.
- Plagborg-Møller, M., Reichlin, L., Ricco, G., & Hasenzagl, T. (2020). When is growth at risk? *Brookings Papers on Economic Activity, Spring 2020*.
- Prasad, A., Elekdag, S., Jeasakul, P., Lafarguette, R., Alter, A., Feng, A. X., et al. (2019). *Growth at risk: Concept and application in IMF country surveillance: IMF working papers. No. 19/36*, International Monetary Fund.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3), 821–852.

- Reichlin, L., Ricco, G., & Hasenzagl, T. (2020). *Financial variables as predictors of real growth vulnerability: CEPR discussion papers DP14322*, Centre for Economic Policy Research.
- Salgado, S., Guvenen, F., & Bloom, N. (2019). *Skewed business cycles: NBER working papers. No. 26565*, National Bureau of Economic Research.
- Stock, J. H., & Watson, M. W. (2003). Has the business cycle changed and why? In *NBER Chapters, NBER Macroeconomics Annual 2002, Volume 17* (pp. 159–230). National Bureau of Economic Research, Inc.
- Summers, P. M. (2005). What caused the great moderation? : Some cross-country evidence. *Economic Review*, (Q III), 5–32.
- Tanner, M. A., & Wong, W. H. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association*, 82(398), 528–540.
- Tsionas, E. G. (2002). Bayesian inference in the noncentral student-t model. *Journal of Computational and Graphical Statistics*, 11(1), 208–221.
- Tsiotas, G. (2012). On generalised asymmetric stochastic volatility models. *Computational Statistics & Data Analysis*, 56(1), 151–172.