A time-varying skewness model for Growth-at-Risk

This paper develops a new empirical model to study macro-financial risks and shows that changing financial conditions can signal risks to the path of future economic growth.

Martin Iseringhausen
European Stability Mechanism

European Stability Mechanism

A time-varying skewness model for Growth-at-Risk

Martin Iseringhausen[1] European Stability Mechanism

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 across all horizons up to four quarters ahead. The suggested approach can inform policy-makers' assessment of macro-financial vulnerabilities by providing a timely signal of shifting risks and a quantification of their magnitude.

Keywords: Bayesian analysis, downside risk, macro-financial linkages, time variation

JEL codes: C11, C23, C53, E44

[1] m.iseringhausen@esm.europa.eu

ISSN 2443-5503
ISBN 978-92-95223-06-6

doi:10.2852/176745
EU catalog number DW-AB-21-02-EN-N

A time-varying skewness model for Growth-at-Risk

Martin Iseringhausen*

European Stability Mechanism

June 2021

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 across all horizons up to four quarters ahead. The suggested approach can inform policy-makers' assessment of macro-financial vulnerabilities by providing a timely signal of shifting risks and a quantification of their magnitude.

JEL classification: C11, C23, C53, E44

Keywords: Bayesian analysis, downside risk, macro-financial linkages, time variation

*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 and the participants of an ESM Research Seminar for helpful comments, as well as Sergei Antoshin for sharing the Financial Conditions Indices (FCI) produced by the IMF. Correspondence to: Martin Iseringhausen, European Stability Mechanism, 6a Circuit de la Foire Internationale, L-1347 Luxembourg. E-mail: m.iseringhausen@esm.europa.eu.

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 analyse 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 et al. (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 analyse risks to economic outcomes both in academia and in policy institutions (Caldera Sánchez and Röhn, 2016; Giglio et al., 2016; Prasad et al., 2019). While the approach has been 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 makes it possible 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 et al., 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 been recently scrutinised. Several papers have questioned the ability of financial variables to inform the analysis of downside risk (e.g. Plagborg-Møller et al., 2020) or, at least, the stability of the relationship between these conditions and future growth vulnerability (Reichlin et al., 2020). Moreover, even if studies identify an impact of financial variables on macroeconomic risks, the question remains whether this can benefit out-of-sample forecasting. Brownlees and Souza (2021) show that standard time-varying volatility (GARCH)

[1] Extending this analysis to a larger set of advanced economies, Adrian et al. (2021) highlight an intertemporal trade-off between short-term benefits and medium-term risks of loose financial conditions.

models compete well with quantile regressions in forecasting Growth-at-Risk and often outperform those despite the fact that GARCH forecasts are only based on growth data. In addition, Carriero et al. (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 also in these models financial conditions 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 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] Our proposed empirical model measures such 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 crucial for forecasting accuracy (see e.g. D'Agostino et al., 2013), and the possibility to evaluate the effect of macro-financial conditions on the presence of risks. The latter further helps to improve forecasting results and increases the applicability in a policy environment, where the goal is to identify sources of risk, in addition to predicting it.

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 et al. (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 and Siddique, 1999). Recently, Plagborg-Møller et al. (2020) use a model with skewed-t innovations where the parameters of the distribution are driven by an aggregate economic 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. In contrast, in a more general score-driven version of this model, Delle Monache et al. (2020) find both volatility and skewness of US growth to vary significantly over time. The

[2]Contributions providing theoretical support for a time-varying degree of asymmetry over the business cycle include, for example, Orlik and Veldkamp (2014), Salgado et al. (2019), and Jensen et al. (2020).

before-mentioned approaches, including the one developed in this paper, can also be viewed as alternatives to Markov-switching approaches (e.g. Hamilton, 1989; Morley and Piger, 2012; Caldara et al., 2020) or threshold models (e.g. Alessandri and 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 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 across horizons. 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 generally not improve or only adds 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 this paper.

2.1 Empirical specification

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

$$y_{it} = \mu_{it} + e^{h_{it}/2}\varepsilon_{it}, \qquad i = 1, ..., N, \quad t = 1, ..., T, \tag{1}$$

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

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

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

The shape of the noncentral t-distribution is driven by two parameters, the degrees of freedom ν, that are assumed 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 has been 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 either modeled 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,

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

where X_{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_{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:

$$\begin{aligned} \sigma_h^2 &\sim \mathcal{IG}(c_{h0}, C_{h0}), & \sigma_\delta^2 &\sim \mathcal{IG}(c_{\delta 0}, C_{\delta 0}), & \nu &\sim \mathcal{U}(0, \bar{\nu}), \\ \beta &\sim \mathcal{N}(\beta_0, \sigma_{\beta 0}^2 I_K), & \phi &\sim \mathcal{TN}_{(-1,1)}(\phi_0, \sigma_{\phi 0}^2). \end{aligned} \tag{5}$$

To obtain closed-form expressions of the model-implied variance and skewness for each country and period, we rely on the result of Hogben et al. (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 moment 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,

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

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

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 in Iseringhausen (2020), the error term in Equation (3) has zero mean but is not standardised 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.[3] 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 $y = (y_1, ..., y_T)'$, $\mu = (\mu_1, ..., \mu_T)'$, $h = (h_1, ..., h_T)'$, $\delta = (\delta_1, ..., \delta_T)'$, and $\lambda = (\lambda_1, ..., \lambda_T)'$, where each element y_t, μ_t, h_t, δ_t, and λ_t is a $N \times 1$ vector. Finally, define $X = (X_1, ..., X_N)'$ where each X_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, 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).

[3] See Equation (B-17) in Appendix B.

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 et al., 1995; Tsionas, 2002). Specifically, if the random variable X follows the noncentral t-distribution, it has the following representation,

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

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

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

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 towards how to linearise the model for the purpose of sampling the unobserved (log-)volatility h_{it}. The approach taken here directly follows Iseringhausen (2020), which 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 Equation (9), obtained after squaring both sides and taking the natural logarithm,

$$log\left((y_{it} - \mu_{it})^2 + c\right) = h_{it} + \widetilde{\varepsilon}_{it}, \tag{10}$$

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

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

This transformation allows h_{it} to now enter the model in a linear manner. If the error term in Equation (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 $\widetilde{\varepsilon}_{it}$ 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 $\widetilde{\varepsilon}_{it}$ by a mixture of normal distributions,

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

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 is '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 et al. (2007) and Iseringhausen (2020) and choose $M = 10$. Additional details on this approach can be found in Iseringhausen (2020). The mixture representation in Equation (12) can be reformulated based on the probabilities of the Gaussian components,

$$\widetilde{\varepsilon}_{it}|s_{it} = j, \nu, \delta_{it} \sim \mathcal{N}\left(m_j(\nu, \delta_{it}), v_j^2(\nu, \delta_{it})\right), \qquad Pr(s_{it} = j|\nu, \delta_{it}) = q_j(\nu, \delta_{it}), \tag{13}$$

where the mixture indicators $s_{it} \in [1, ..., 10]$ 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 in Iseringhausen (2020) and loops over the following blocks:

1. Conditional mean coefficients: draw γ from $p(\gamma|y, h, \delta, \lambda, \nu)$;
2. Mixture indicators: draw s from $p(s|y, h, \delta, \nu, \gamma)$;
3. (Log-)volatility: draw h from $p(h|y, s, \delta, \nu, \gamma, \sigma_h^2)$;
4. Location-scale weights: draw λ from $p(\lambda|y, h, \delta, \nu, \gamma)$;

5. Degrees of freedom: draw v from $p(\nu|\lambda)$;
6. Non-centrality parameter: draw δ from $p(\delta|y, h, \lambda, \nu, \gamma, X, \phi, \beta, \sigma_\delta^2)$;
7. Coefficients in δ-equation: draw ϕ and β from $p(\phi, \beta|\delta, X, \sigma_\delta^2)$;
8. Innovation variances: draw σ_h^2 from $p(\sigma_h^2|h)$ and σ_δ^2 from $p(\sigma_\delta^2|\delta, X, \phi, \beta)$.

In terms of practical implementation, the 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 Tsionas (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 ν_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 and Hsiao, 2014; Chan and Jeliazkov, 2009) that can yield large computational efficiency gains compared to more standard algorithms for estimating linear Gaussian state space models (e.g. Carter and Kohn, 1994). The conditional posterior distributions of λ (block 4) and ν (block 5) are non-standard and a Metropolis-Hastings step needs to be included as described in Tsionas (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 initialised 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 $(B+1, ..., J)$ can be taken as a sample from the joint posterior distribution $f(h, \delta, \lambda, \nu, \gamma, \phi, \beta, \sigma_h^2, \sigma_\delta^2|y, X)$. 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 for another 2,000,000 times, keeping only every 20th draw.[4] 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. Moreover, Appendix C comments on the convergence properties and efficiency of the sampler.

[4]Applying this 'thinning' has only computational reasons, i.e. memory limits (see Gelman et al., 2011).

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 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), United Kingdom (GB), and 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 and Souza, 2021). The 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 are regularly drawing the attention of policy-makers: 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 et al. (2018).[5] All variables are standardised 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

[5]The inclusion of an uncertainty measure is also motivated by the findings in Hengge (2019) and Jovanovic and Ma (2020) who, using the forecast error based uncertainty measure of Jurado et al. (2015), find a strong effect of uncertainty on the lower quantiles of the conditional US output (GDP/IP) growth distribution. Unfortunately, this measure of uncertainty is not readily available for all countries in our sample.

imputed values directly from the dataset made available by these authors. Details on the dataset including the imputation procedure (see also Brownlees and 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 the distortions due to imputed values by only including 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 Equations (1)-(4), where the dependent variable y_{it} is simply shifted by h periods.[6] 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 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.[7]

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 $\gamma_i \sim \mathcal{N}(\mathbf{0}, 10 \times I_k)$, $\phi \sim \mathcal{N}(0, 10)$, 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, Figure 1 shows the estimated model-implied variance of the one-step-ahead predictive distribution for each country, calculated using Equation (6). These plots clearly show the well-known Great Moderation, i.e. a significant decline in output volatility in most industrialised economies with some differences in the timing and magnitude across countries (Stock and Watson, 2003; Summers, 2005; Del Negro and Otrok, 2008). 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

[6] 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. $t + h - 1$.

[7] 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.

the result of the fact that the explanatory variables in the asymmetry equation also affect the conditional variance as explained in the previous section.

Figure 1: Variance of in-sample one-step-ahead predictive distribution

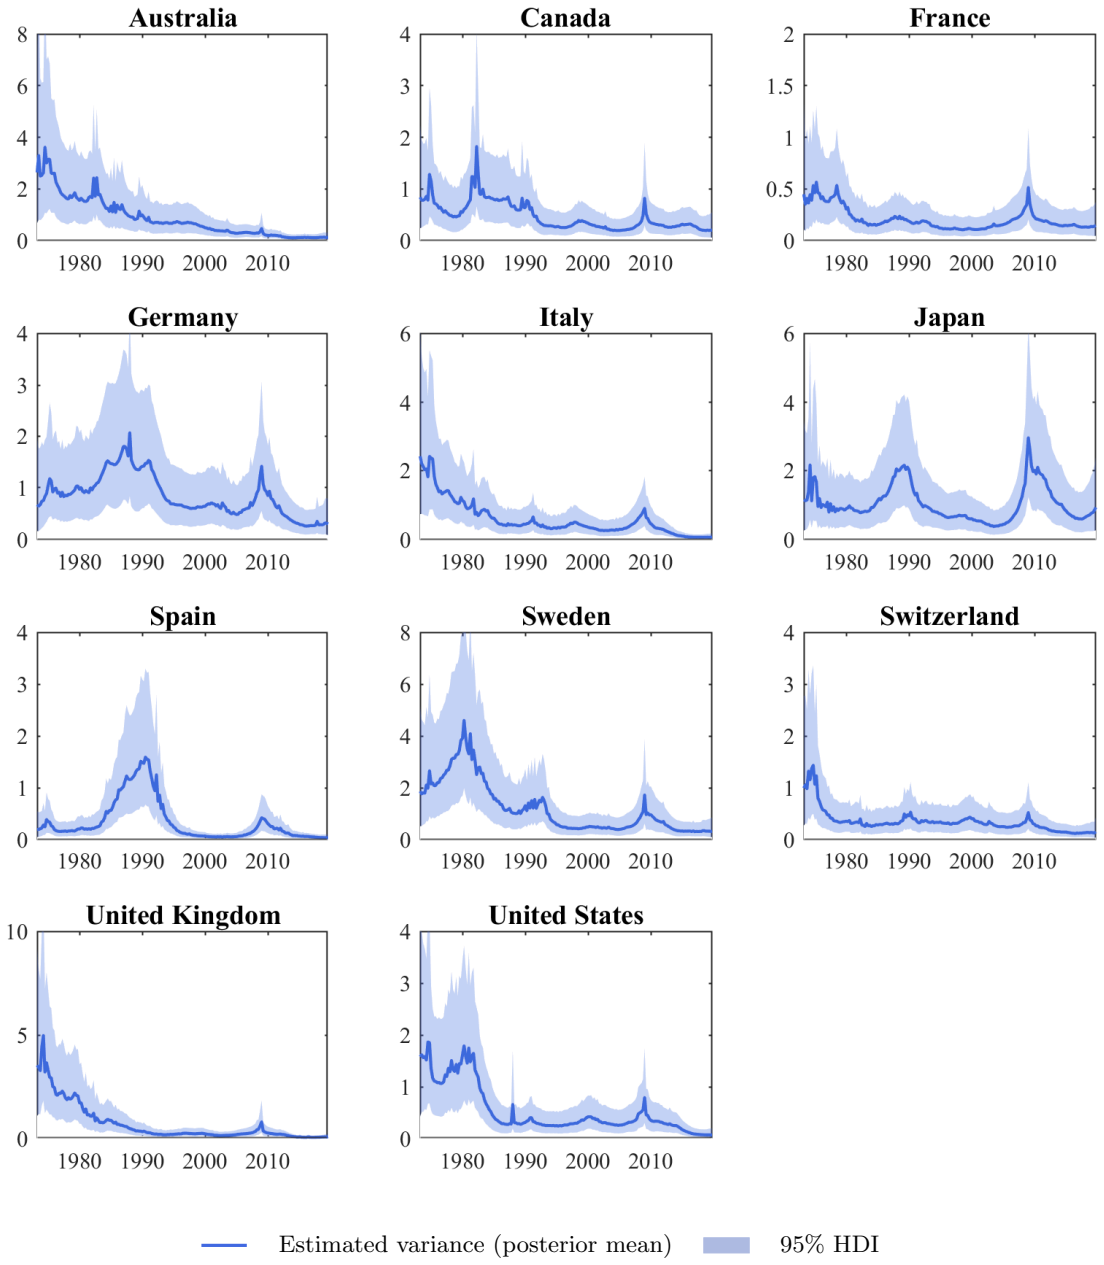

Figure 2 presents one of the key results of the paper, the estimated conditional skewness of the one-step-ahead predictive GDP growth distribution across countries. First, there is signif-

icant 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 skew-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.

Figure 2: Skewness of in-sample one-step-ahead predictive distribution

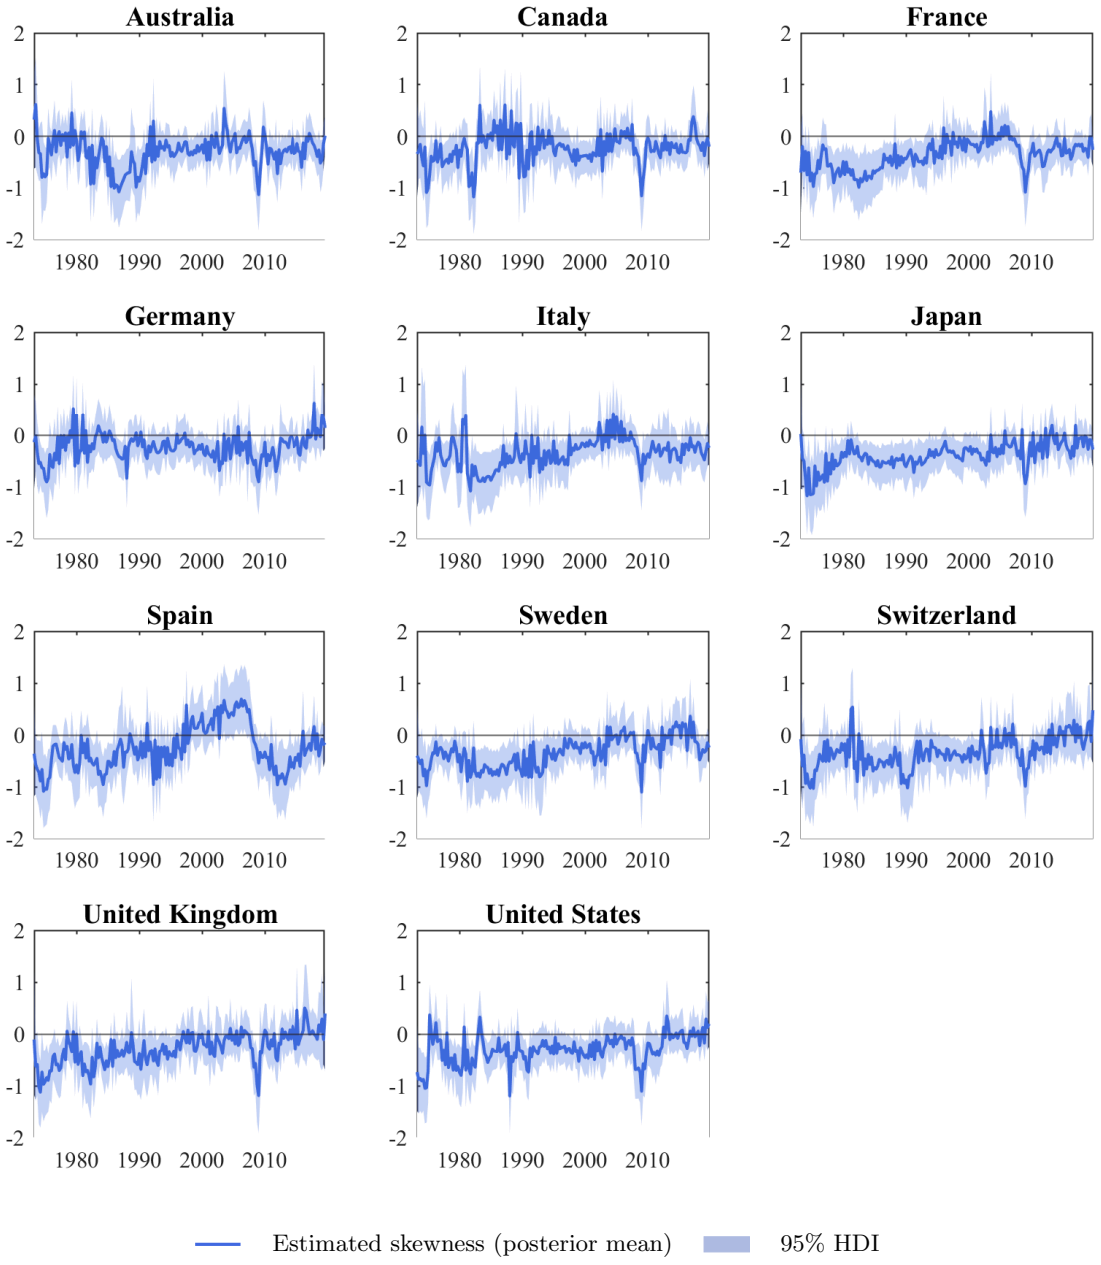

This is a first piece of evidence that suggests the existence of sizable downside risk is a common phenomenon across countries and over time. In particular, during the Great Recession the left tails of conditional growth distributions were pronounced pointing towards a highly vulnerable macroeconomic 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 analyse the role of financial conditions and uncertainty in determining the shape of the predictive growth 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 Equation (4). However, conditional skewness is a non-linear function of δ determined by Equation (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 standardised, 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 Equation (4) into the skewness Equation (7) and obtaining the first derivative of this expression w.r.t. 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 confirm the findings of previous work and suggest an important role of financial conditions in shaping the growth distribution. Worsening financial conditions, as reflected in an increase of 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 generally not support a prominent role of those 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 before, 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 in Hengge (2019) showing that economic policy uncertainty has limited power to explain growth vulnerabilities.

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

		$h=1$			$h=4$			$h=8$		
		Mean	$\text{Perc}_{2.5}$	$\text{Perc}_{97.5}$	Mean	$\text{Perc}_{2.5}$	$\text{Perc}_{97.5}$	Mean	$\text{Perc}_{2.5}$	$\text{Perc}_{97.5}$
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
$\text{Growth-at-Risk}_{5\%}$		100	100	0.076	100	91	0.079	91	100	0.080
$\text{Growth-at-Risk}_{95\%}$		100	100	0.067	82	91	0.069	100	100	0.068
		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$\text{Expected shortfall}_{5\%}$		0.171	0.355		0.177	0.382		0.173	0.387	
$\text{Expected longrise}_{95\%}$		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 in Engle and Manganelli (2004). For details see Brownlees and Souza (2021) and also Section 3.3. *TL* is the tick loss, *EKP* refers to the expected shortfall precision measure of Embrechts et al. (2005), and the VaR-ES score is the measure developed in Fissler et al. (2015).

At the one-year-ahead horizon ($h = 4$), tightening financial conditions, as measured by the FCI, do 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 does no longer include zero. Qualitatively, these results support the findings of Adrian et al. (2021), i.e. the existence of a trade-off 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 in 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 Figure 2, the estimated intercept of the asymmetry process is clearly negative.

Table 1 also reports various measures to assess the accuracy of GaR and ES/EL predictions. Throughout the paper we focus on the 5% (95%) quantile to analyse downside (upside) risk. Formally defined, GaR is the p%-quantile of the predictive growth distribution[8],

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

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

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

The reported accuracy measures are also used in the out-of-sample forecasting exercise and will be explained in more detail in the next section. They are reported here to allow for a comparison of both in-sample and out-of-sample results.

To conclude the in-sample analysis, Figure 3 shows the model-implied one-step-ahead Growth-at-Risk and expected shortfall/longrise values over time along with the realised 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 Figures 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).[9]

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 stronger 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

[8] While GaR usually refers to the lower quantile, for simplicity, here we also use the label for the upper quantile.
[9] The average standard deviation across countries for downside (upside) risk, as shown in Figure 3, is 0.59 (0.52) and 0.78 (0.64) for Growth-at-Risk and expected shortfall/longrise, respectively.

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 larger variability of downside risk than of upside risk.

Figure 3: In-sample one-step-ahead Growth-at-Risk and expected shortfall/longrise

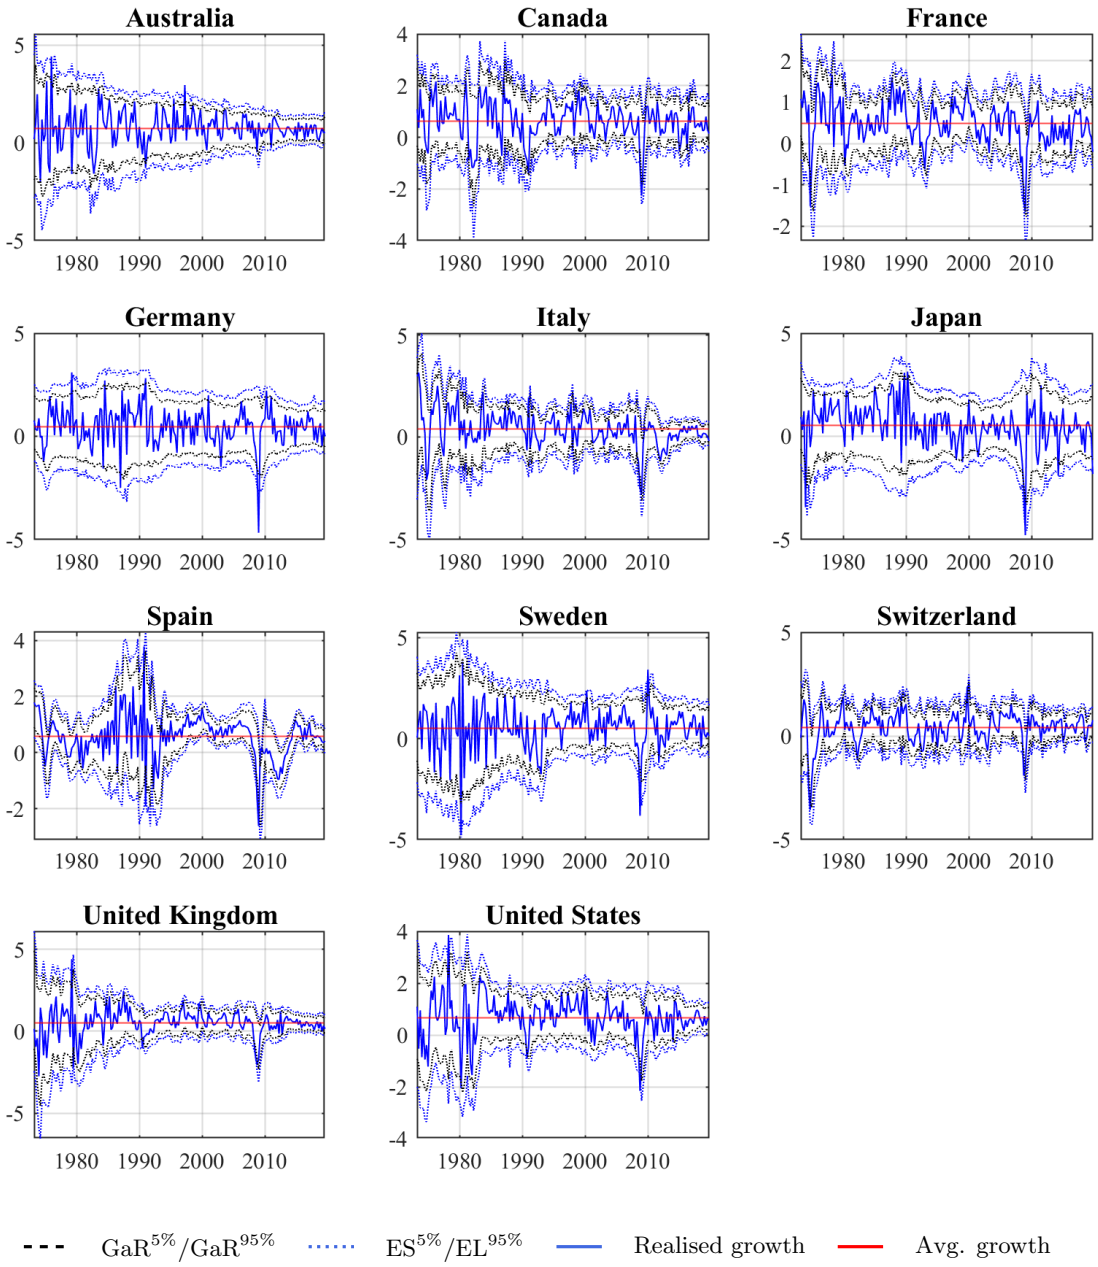

However, the possibility to observe this one particular pattern in both models with sym-

metric 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 will focus 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.

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 1973:Q1–1984:Q4 and then generate the h-step-ahead, $h \in \{1, ..., 4\}$, $GaR^p/ES^p/EL^p$ forecasts.[10] 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 $T^h_{oos} = 140 - (h-1)$.

Competing models and accuracy measures

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 generalised autoregressive conditional heteroscedasticity (GARCH) models applied in Brownlees and Souza (2021).[11] Lastly, we also 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 includes a constant, current GDP growth and three additional lags.

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 forecast for period $t+h$ is simply the empirical quantile and the empirical expected shortfall/longrise calculated using all growth observations of a particular country available up until period t.

Quantile regression

Growth-at-Risk analysis through quantile regression has been popularised by Adrian et al.

[10] 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. The FCIs are also not constructed in real-time and include a look-ahead bias (Brownlees and Souza, 2021).

[11] For these two models, we use the MATLAB routines made available by Adrian et al. (2019) and Brownlees and Souza (2021), respectively.

(2019) and has become a regular monitoring exercise at the IMF and other organisations. Quantile regression was developed by Koenker and Bassett (1978) and we refer to this work for details on the estimation. It can be viewed as a generalisation of a standard (mean) regression where the p%-quantile of the dependent variable, $\mathcal{Q}_p(y)$, is directly modeled as a function of explanatory variables,

$$\mathcal{Q}_p(y_{i,t+h|t}) = \alpha^p_{0i} + \alpha^p_{1i} y_{it} + \alpha^p_{2i} FCI_{it}. \tag{16}$$

The baseline quantile regressions for the forecasting exercise include a constant, current growth and the FCI.[12] The corresponding country-specific quantile regression coefficients α^p_{0i}, α^p_{1i}, and α^p_{2i} 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 in Brownlees and Souza (2021). In essence, their approach is a GARCH(1,1) model with a flexible non-parametric form used to model the standardised growth distribution. The model has the following form,

$$y_{i,t+1} = \mu_{i,t+1|t} + \sqrt{\sigma^2_{i,t+1|t}}\varepsilon_{i,t+1}, \qquad \varepsilon_{i,t+1} \sim \mathcal{D}_{\varepsilon_i}(0,1), \tag{17}$$

$$\sigma^2_{i,t+1|t} = \sigma^2_i(1-\alpha-\beta) + \alpha u^2_{it} + \beta\sigma^2_{it}, \qquad \alpha,\beta > 0,\ (\alpha+\beta) < 1, \tag{18}$$

where $\sigma^2_{i,t+1|t}$ is the conditional (deterministic) one-step-ahead variance, σ^2_i is the unconditional variance, α and β are the GARCH parameters which are pooled across countries, and u_{it} is the non-standardised residual. The model is estimated by so-called composite likelihood methods following Pakel et al. (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

[12] Quantile regressions including all the explanatory variables generally perform worse in predicting downside risk and these results are discussed in the robustness section.

setting. For further methodological details, we refer 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 the previous sections and, motivated by the in-sample analysis, a set that only includes the FCI. Both specifications include a constant and an autoregressive term in the asymmetry equation.[13]

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 Equation (4) by h-quarters, starting by using the draw for the last in-sample period together with the h-period-lagged explanatory variables. This process is continued to obtain the forecast of δ for $h > 1$. Importantly, since X in Equation (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. This model is a SV model with t-distributed growth shocks obtained by imposing the restriction $\delta = 0$ for each country.

Accuracy measures for quantile and expected shortfall forecasts

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 have also been used in Brownlees and Souza (2021) and we refer there for a more detailed explanation.

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

[13] 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.

are the average tick loss across countries computed as

$$TL_p = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{1}{T}\sum_{t=1}^{T}(y_{i,t+h} - GaR^p_{i,t+h|t})(p - \mathbb{I}_{\{y_{i,t+h}<GaR^p_{i,t+h|t}\}})\right). \tag{19}$$

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 *elicitability* property (Fissler and Ziegel, 2016), quantile and shortfall forecasts can be evaluated jointly. Specifically, as Carriero et al. (2020a), we rely on the VaR-ES score of Fissler et al. (2015),

$$\begin{aligned}\text{VaR-ES score}_p = \frac{1}{N}\sum_{i=1}^{N}(\frac{1}{T}\sum_{t=1}^{T}((GaR^p_{i,t+h|t}(\mathbb{I}_{\{y_{i,t+h}<GaR^p_{i,t+h|t}\}} - p) \\ - y_{i,t+h}\mathbb{I}_{\{y_{i,t+h}<GaR^p_{i,t+h|t}\}} + \frac{e^{ES^p_{i,t+h|t}}}{1+e^{ES^p_{i,t+h|t}}}(ES^p_{i,t+h|t} - GaR^p_{i,t+h|t} \\ + p^{-1}(GaR^p_{i,t+h|t} - y_{i,t+h}\mathbb{I}_{\{y_{i,t+h}<GaR^p_{i,t+h|t}\}})) + \ln\frac{2}{1+e^{ES^p_{i,t+h|t}}})),\end{aligned} \tag{20}$$

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 in Embrechts et al. (2005). This measure has been used in Nakajima (2013) and Iseringhausen (2020) and is explained in more detail in these 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.

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. In contrast, the GARCH model generally outperforms both the historical benchmark and quantile regression for the lower and the upper quantile. Lastly, all three stochastic volatility models perform, on average, similarly compared to the GARCH approach with only small differences among them.

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

	$h=1$			$h=2$			$h=3$			$h=4$		
Historical benchmark												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	73	18	0.101	73	82	0.102	73	82	0.103	73	82	0.103
$GaR_{95\%}$	45	64	0.081	45	73	0.081	45	64	0.082	45	**82**	0.082
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.452	0.575		0.455	0.588		0.470	**0.593**		0.473	**0.596**	
$EL_{95\%}$	0.446	0.690		0.450	0.691		0.454	0.691		0.457	0.692	
Country-specific quantile regressions (FCI only)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	73	45	0.090	82	55	0.096	64	73	0.102	64	64	0.109
$GaR_{95\%}$	91	82	0.079	64	73	0.080	64	**91**	0.082	64	**82**	0.082
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.276	0.497		0.387	0.538		0.589	0.622		0.519	0.701	
$EL_{95\%}$	0.272	0.676		0.334	0.698		0.240	0.697		0.250	0.693	
Panel-GARCH(1,1)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	91	**82**	0.085	82	**100**	0.094	91	**100**	0.101	**91**	91	0.104
$GaR_{95\%}$	**100**	91	0.070	82	**100**	0.072	64	**91**	0.076	55	64	0.078
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.189	0.447		0.423	0.526		0.470	0.612		0.478	0.652	
$EL_{95\%}$	0.294	0.650		0.343	0.654		0.342	0.666		0.422	0.679	
SV model with symmetric shocks												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	**100**	64	0.083	**100**	73	**0.092**	91	82	0.100	**91**	**100**	0.103
$GaR_{95\%}$	**100**	**100**	**0.068**	**100**	82	0.070	**100**	82	0.074	91	73	**0.077**
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.203	0.427		0.301	0.519		0.440	0.615		0.440	0.665	
$EL_{95\%}$	0.175	0.628		0.201	0.639		0.181	0.656		0.204	0.666	
SV model with time-varying skewness (FCI only)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	**100**	73	**0.082**	**100**	82	**0.092**	**100**	91	0.100	**91**	**100**	0.103
$GaR_{95\%}$	**100**	**100**	**0.068**	**100**	82	**0.069**	**100**	82	**0.073**	**100**	73	**0.077**
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	**0.155**	**0.415**		**0.276**	**0.513**		**0.411**	0.610		**0.429**	0.666	
$EL_{95\%}$	0.153	0.628		**0.175**	**0.637**		**0.162**	**0.653**		**0.194**	**0.664**	
SV model with time-varying skewness (FCI + TS + HP + WUI)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	**100**	64	**0.082**	**100**	73	0.093	**100**	73	**0.099**	**91**	**100**	**0.102**
$GaR_{95\%}$	**100**	**100**	**0.068**	91	91	0.071	**100**	82	0.075	82	**82**	0.078
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.183	0.418		0.328	0.527		0.437	0.611		0.455	0.663	
$EL_{95\%}$	**0.148**	**0.627**		0.191	0.642		0.178	0.657		0.200	0.670	

Note: This table contains the results of the out-of-sample forecasting exercise. 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 in Engle and Manganelli (2004). For details see also Brownlees and Souza (2021). *TL* is the tick loss, *EKP* refers to the expected shortfall precision measure of Embrechts et al. (2005), and the VaR-ES score is the measure developed in Fissler et al. (2015). Bold numbers indicate the model with the highest average accuracy for each measure.

Second, in terms of the tick loss (TL), the stochastic volatility models are the best performing models across quantiles and horizons. Specifically, the model only including the FCI as a driver of the shape of the predictive distribution has, on average, some small advantages in predicting downside risk for $h = 1$ and $h = 2$, and for upside risk across horizons. The other variables next to the FCI only marginally improve forecasts of downside risk at the remaining horizons. Generally, 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 model with time-varying skewness driven by the FCI proves 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 the difficulty of econometric models to precisely forecast downside risk at longer horizons.

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 only including the FCI is clearly smaller than those of the strongest competitors when measuring downside risk. 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 and Mariano, 1995). Specifically, as in Brownlees and Souza (2021), we conduct pair-wise tests with the null hypothesis being equal predictive ability against an 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), respectively. 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 and 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 3-5 (out of 11) countries, and for 1-2 countries when evaluated against the symmetric SV model. In turn, our proposed model is outperformed by other models for at most 1-2 countries. For $h = 2$ the 'FCI only' SV model continues to show, on average, the most cases of outperformance. Overall, and in line with Brownlees and Souza (2021), the evidence becomes weaker as h increases and for $h > 2$ most models have difficulties to outperform 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 including only the FCI is performing particularly well.

Table 3: Tests for superior forecasting performance

Downside risk (5%)

		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 (73)	64 (73)	64 (73)	55 (73)	-	36 (45)	45 (55)	55 (55)	55 (64)	55 (55)
	QR	0 (0)	-	27 (27)	27 (27)	27 (36)	27 (27)	0 (0)	-	36 (36)	36 (45)	36 (45)	27 (45)
	GARCH	0 (0)	9 (9)	-	36 (36)	27 (27)	27 (36)	0 (0)	9 (9)	-	36 (45)	36 (36)	18 (27)
	SV (sym.)	0 (0)	9 (18)	0 (0)	-	0 (18)	0 (0)	0 (0)	9 (18)	0 (0)	-	0 (9)	0 (0)
	SV (FCI)	0 (0)	9 (18)	0 (0)	0 (9)	-	0 (9)	0 (0)	0 (9)	0 (0)	0 (0)	-	0 (0)
	SV (full)	0 (0)	9 (18)	0 (0)	9 (9)	9 (36)	-	0 (0)	0 (18)	0 (0)	9 (18)	9 (27)	-
2	Hist.	-	27 (36)	36 (55)	18 (36)	36 (45)	36 (45)	-	18 (18)	27 (36)	9 (18)	9 (36)	9 (18)
	QR	0 (0)	-	27 (27)	18 (27)	18 (27)	18 (27)	0 (0)	-	18 (18)	9 (18)	18 (18)	18 (18)
	GARCH	0 (9)	0 (0)	-	18 (27)	18 (18)	9 (27)	0 (0)	0 (0)	-	0 (9)	9 (9)	0 (9)
	SV (sym.)	9 (9)	0 (0)	9 (9)	-	9 (18)	9 (18)	0 (0)	0 (0)	0 (9)	-	0 (9)	0 (18)
	SV (FCI)	9 (9)	0 (9)	0 (9)	9 (9)	-	0 (9)	0 (0)	0 (0)	0 (0)	9 (9)	-	0 (0)
	SV (full)	9 (9)	0 (0)	9 (9)	0 (9)	0 (0)	-	0 (0)	0 (0)	0 (0)	0 (9)	0 (0)	-
3	Hist.	-	27 (36)	9 (9)	9 (9)	9 (9)	9 (9)	-	27 (36)	9 (9)	9 (9)	9 (9)	9 (9)
	QR	9 (18)	-	18 (45)	9 (18)	9 (9)	9 (9)	0 (27)	-	9 (27)	0 (9)	9 (9)	9 (9)
	GARCH	9 (9)	0 (0)	-	18 (18)	0 (9)	0 (18)	0 (9)	0 (0)	-	0 (9)	0 (0)	0 (9)
	SV (sym.)	9 (18)	0 (0)	0 (9)	-	0 (0)	0 (0)	0 (9)	0 (0)	0 (0)	-	0 (9)	9 (9)
	SV (FCI)	9 (18)	0 (0)	9 (9)	18 (27)	-	18 (27)	0 (9)	0 (0)	0 (0)	0 (0)	-	18 (18)
	SV (full)	0 (18)	0 (0)	9 (18)	0 (27)	0 (0)	-	0 (9)	0 (0)	0 (0)	9 (18)	0 (9)	-
4	Hist.	-	9 (9)	9 (9)	18 (18)	18 (18)	18 (18)	-	0 (9)	9 (9)	9 (9)	18 (18)	18 (18)
	QR	9 (18)	-	27 (27)	18 (27)	18 (36)	18 (36)	0 (18)	-	18 (18)	9 (18)	9 (18)	9 (18)
	GARCH	18 (36)	0 (0)	-	18 (27)	18 (18)	18 (18)	18 (18)	0 (0)	-	18 (27)	18 (27)	18 (27)
	SV (sym.)	18 (18)	0 (0)	0 (0)	-	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	-	0 (0)	9 (9)
	SV (FCI)	18 (18)	0 (0)	0 (0)	36 (45)	-	36 (36)	0 (18)	0 (0)	0 (0)	27 (27)	-	27 (27)
	SV (full)	18 (18)	0 (0)	0 (0)	9 (18)	0 (9)	-	0 (18)	0 (0)	0 (0)	9 (18)	0 (9)	-

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)	64 (73)	82 (91)	82 (82)	73 (91)	-	36 (36)	73 (82)	91 (91)	91 (91)	91 (91)
	QR	27 (27)	-	64 (73)	64 (73)	73 (73)	64 (73)	18 (27)	-	45 (45)	64 (73)	64 (73)	64 (73)
	GARCH	0 (0)	0 (0)	-	27 (27)	36 (45)	36 (36)	0 (0)	0 (0)	-	27 (45)	36 (45)	36 (45)
	SV (sym.)	0 (0)	0 (0)	9 (9)	-	18 (18)	0 (9)	0 (0)	0 (0)	18 (18)	-	9 (18)	0 (9)
	SV (FCI)	0 (0)	0 (0)	9 (9)	0 (0)	-	0 (9)	0 (0)	0 (0)	9 (18)	0 (9)	-	0 (9)
	SV (full)	0 (0)	0 (0)	9 (9)	27 (36)	27 (55)	-	0 (0)	0 (0)	9 (9)	18 (27)	36 (45)	-
2	Hist.	-	18 (27)	73 (73)	55 (73)	64 (73)	55 (64)	-	45 (45)	73 (82)	73 (82)	82 (82)	82 (82)
	QR	0 (9)	-	27 (64)	36 (55)	36 (73)	36 (73)	0 (0)	-	36 (64)	36 (73)	45 (73)	45 (73)
	GARCH	0 (0)	0 (0)	-	18 (18)	27 (36)	27 (27)	0 (0)	0 (0)	-	27 (36)	27 (45)	27 (36)
	SV (sym.)	0 (0)	0 (0)	9 (27)	-	45 (55)	18 (27)	0 (0)	0 (0)	9 (18)	-	45 (64)	18 (27)
	SV (FCI)	0 (0)	0 (0)	0 (18)	0 (0)	-	0 (0)	0 (0)	0 (0)	0 (9)	0 (0)	-	0 (0)
	SV (full)	0 (0)	0 (0)	9 (27)	0 (0)	55 (55)	-	0 (0)	0 (0)	0 (27)	0 (0)	55 (55)	-
3	Hist.	-	36 (36)	55 (55)	36 (55)	36 (55)	36 (55)	-	36 (36)	55 (73)	55 (55)	55 (55)	55 (55)
	QR	0 (0)	-	18 (27)	9 (27)	9 (36)	9 (27)	0 (9)	-	9 (18)	9 (36)	9 (36)	9 (27)
	GARCH	9 (9)	9 (9)	-	18 (36)	36 (36)	36 (36)	9 (9)	9 (9)	-	36 (36)	27 (45)	27 (36)
	SV (sym.)	0 (0)	0 (18)	9 (18)	-	27 (36)	9 (9)	0 (0)	0 (9)	0 (9)	-	18 (27)	9 (9)
	SV (FCI)	0 (0)	0 (9)	0 (9)	0 (0)	-	0 (0)	0 (0)	0 (9)	0 (0)	0 (0)	-	0 (0)
	SV (full)	0 (0)	0 (18)	0 (18)	0 (0)	45 (55)	-	0 (0)	0 (9)	0 (9)	0 (0)	45 (45)	-
4	Hist.	-	27 (27)	36 (45)	27 (27)	27 (45)	27 (36)	-	45 (45)	45 (55)	36 (45)	45 (45)	36 (45)
	QR	0 (9)	-	27 (36)	27 (27)	27 (27)	27 (27)	0 (0)	-	18 (27)	18 (36)	18 (27)	18 (27)
	GARCH	9 (9)	18 (18)	-	18 (27)	18 (36)	18 (27)	9 (9)	27 (27)	-	27 (36)	27 (27)	27 (27)
	SV (sym.)	0 (0)	0 (9)	9 (9)	-	36 (45)	18 (36)	0 (0)	0 (9)	9 (9)	-	18 (55)	18 (27)
	SV (FCI)	0 (0)	0 (9)	9 (9)	0 (0)	-	0 (0)	0 (0)	0 (0)	0 (9)	0 (0)	-	0 (0)
	SV (full)	0 (0)	0 (9)	9 (9)	0 (9)	36 (55)	-	0 (0)	0 (9)	0 (9)	0 (9)	55 (55)	-

Note: This table contains the results of Diebold and Mariano (1995) tests 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).

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.

Figure 4: Out-of-sample one-step-ahead Growth-at-Risk

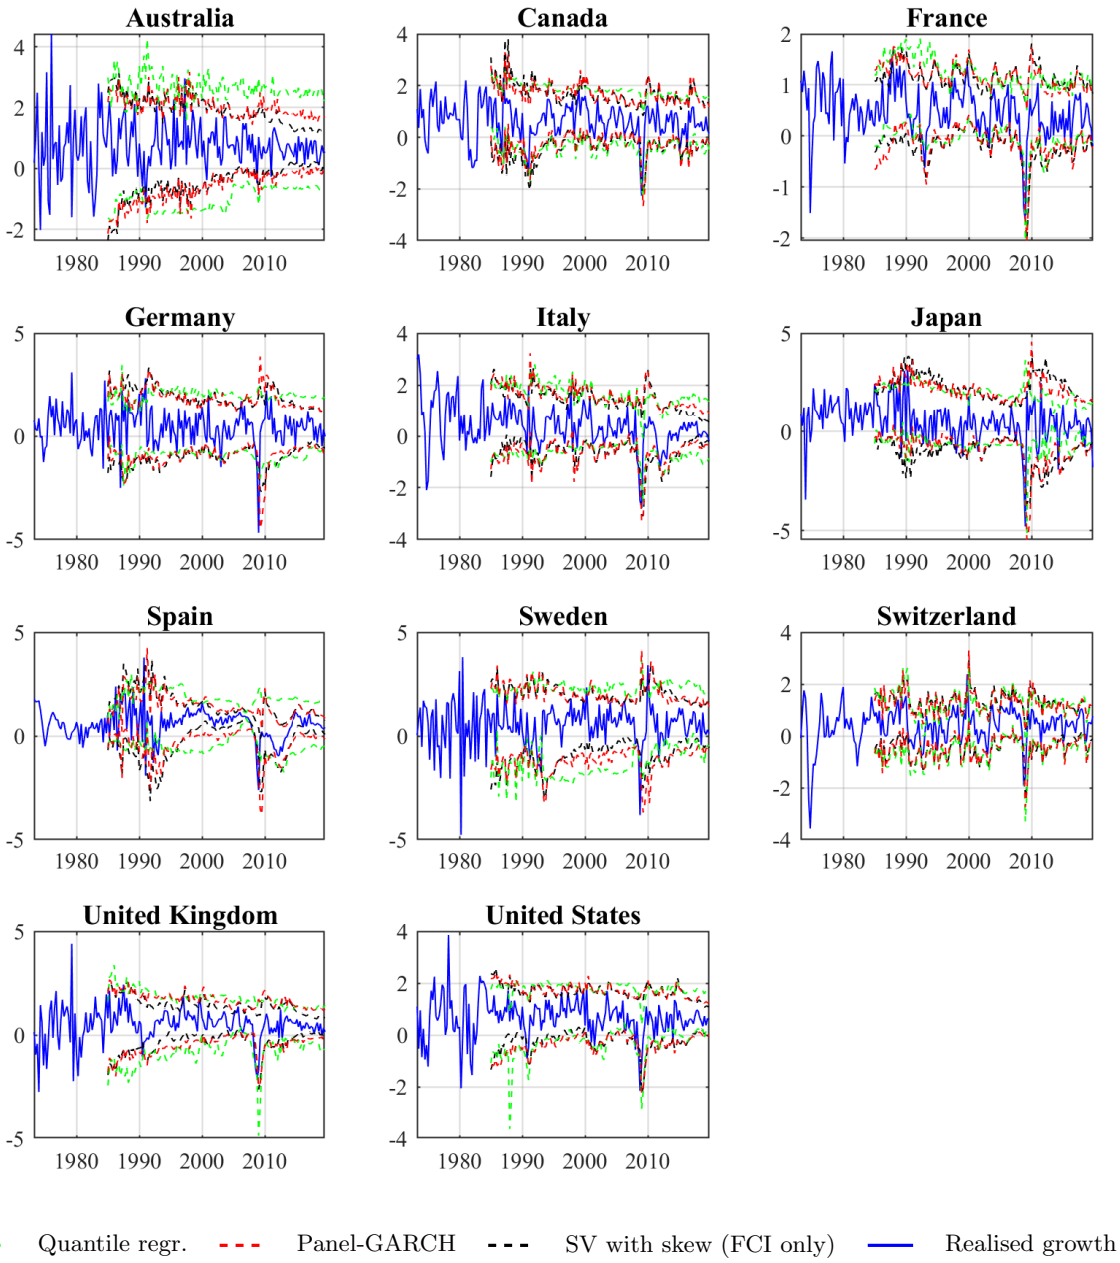

In summary, the results confirm the previous literature in some ways while adding new

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 et al. (2020b) also find the former to perform significantly better in out-of-sample forecasting.[14] Importantly from a policy perspective, we find that a time-varying skewness model including the FCI can help to improve average forecasts of Growth-at-Risk and expected shortfall. However, the gains compared to the symmetric SV model are often small, especially for the quantile forecasts. A larger model including in addition 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.

Figure 4 shows the out-of-sample one-step-ahead GaR forecasts obtained from quantile regressions, panel-GARCH and the SV model with time-varying skewness that only includes the FCI. The corresponding chart showing the one-step-ahead ES/EL forecasts derived from these models can be found in Appendix E. 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 case of Australia) or narrower (e.g. in case of Japan), and are often more erratic. The latter confirms the findings in 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, Figure 5 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. Alongside, we plot the cross-country average quantile regression coefficients for the FCI and the 5% and 95% quantiles, respectively. 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-

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

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 compared to the lower quantile. Future work could take this apparent time variation in the effect of financial conditions on the asymmetry of predictive growth distributions more explicitly into account.

Figure 5: Recursive estimates of the impact of FCI on asymmetry (one-quarter-ahead)

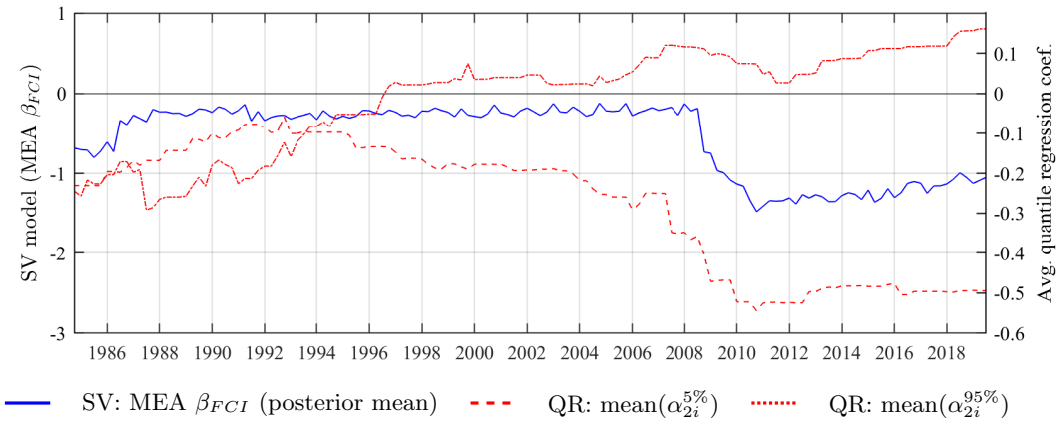

Note: The dates refer to the last observation of the respective estimation (in-sample) period.

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 as well as to discuss the impact of alternative modeling choices. The detailed results are presented in Table E-1 in Appendix E.

First, 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 including the full set of explanatory variables. However, the results indicate that in most cases this leads to a deterioration of the forecasting performance compared to the more parsimonious specification. Second, 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. Third, to ensure that skewness, which is driven by macro-financial conditions, is a useful feature for forecasting, we consider a specification where the FCI is also included in the conditional mean (with country-specific coefficients). Allowing financial conditions to also impact the conditional mean overall improves the average forecasting per-

formance further for $h = 1$ in terms of the loss measures. At the remaining horizons, the accuracy gains mostly occur when predicting downside risk while precision in the right tail of the distribution on average decreases. Importantly, the SV model with time-varying skewness and the FCI in the mean equation still maintains, on average, small advantages compared to the symmetric SV model that also includes the FCI in the mean equation. This is similar to the baseline comparison, i.e. when both models only include AR-dynamics in the mean equation. Finally, 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 this helps in some cases 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. 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.

4 Conclusion

Economic policy-makers 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 scrutinised various aspects of the Growth-at-Risk framework. In particular, the question to what extent financial variables can help to inform this analysis and improve Growth-at-Risk forecasts, has been extensively discussed. Adding to this discussion, this paper 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.

For a panel of 11 OECD countries over the period 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 across horizons when forecasting upside risk. Including, in addition, a measure of economic and policy uncertainty, or some of the prominent components of the FCI separately, does generally not improve or only adds very little to the forecasting performance of the model.

The results derived from the model proposed in this paper can provide valuable insights for policy-makers, especially during the current COVID-19 crisis. Both fiscal and monetary policy have provided large-scale support that has helped to keep financial conditions favorable during the pandemic. While this has undoubtedly been crucial to stabilise the economic environment during this unprecedented crisis, at some point the support will be phased out and financial conditions could be adversely affected. Our results provide additional evidence that such a tightening of financial conditions could create short-term vulnerabilities that require close monitoring. In this regard, the model can provide quantitative evidence to support economic policy decisions. Finally, from a forecasting perspective the suggested approach competes well in predicting risks to the economic outlook, and could thus be a tool that benefits regular macroeconomic risk analysis at policy institutions.

References

Adams, P. A., Adrian, T., Boyarchenko, N., and Giannone, D. (2021). Forecasting macroeconomic risks. *International Journal of Forecasting*, forthcoming.

Adrian, T., Boyarchenko, N., and Giannone, D. (2019). Vulnerable Growth. *American Economic Review*, 109(4):1263–89.

Adrian, T., Grinberg, F., Liang, N., Malik, S., and Yu, J. (2021). The Term Structure of Growth-at-Risk. *American Economic Journal: Macroeconomics*, forthcoming.

Ahir, H., Bloom, N., and Furceri, D. (2018). The World Uncertainty Index. *mimeo*.

Alessandri, P. and Mumtaz, H. (2017). Financial conditions and density forecasts for us output and inflation. *Review of Economic Dynamics*, 24:66–78.

Azzalini, A. and 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.

Box, M. (1971). Bias in Nonlinear Estimation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 33(2):171–190.

Brownlees, C. T. and Souza, A. (2021). Backtesting Global Growth-at-Risk. *Journal of Monetary Economics*, 118:312–330.

Caldara, D., Cascaldi-Garcia, D., Cuba-Borda, P., and Loria, F. (2020). Understanding Growth-at-Risk: A Markov-Switching Approach. *mimeo*.

Caldera Sánchez, A. and Röhn, O. (2016). How do policies influence GDP tail risks? *OECD Economics Department Working Papers*, No. 1339.

Carriero, A., Clark, T. E., and Marcellino, M. G. (2020a). Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions. *Federal Reserve Bank of Cleveland Working Paper*, No. 20-02R.

Carriero, A., Clark, T. E., and Marcellino, M. G. (2020b). Nowcasting tail risks to economic activity with many indicators. *Federal Reserve Bank of Cleveland Working Paper*, No. 20-13R.

Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Chan, J. C. and Hsiao, C. Y. (2014). Estimation of Stochastic Volatility Models with Heavy Tails and Serial Dependence. In Jeliazkov, I. and Yang, X.-S., editors, *Bayesian Inference in the Social Sciences*, pages 155–176. Wiley-Blackwell.

Chan, J. C. and 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.

Chib, S. (2001). Markov Chain Monte Carlo Methods: Computation and Inference. *Handbook of Econometrics*, 5:3569–3649.

D'Agostino, A., Gambetti, L., and Giannone, D. (2013). Macroeconomic Forecasting and Structural Change. *Journal of Applied Econometrics*, 28(1):82–101.

Del Negro, M. and 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., and Petrella, I. (2020). Modeling and Forecasting Macroeconomic Downside Risk. *CEPR Discussion Papers*, DP15109, Centre for Economic Policy Research.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods: Second Edition*. Oxford University Press.

Embrechts, P., Kaufmann, R., and Patie, P. (2005). Strategic Long-Term Financial Risks: Single Risk Factors. *Computational Optimization and Applications*, 32(1-2):61–90.

Engle, R. F. and Manganelli, S. (2004). CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles. *Journal of Business & Economic Statistics*, 22(4):367–381.

Estrella, A. and Hardouvelis, G. A. (1991). The Term Structure as a Predictor of Real Economic Activity. *The Journal of Finance*, 46(2):555–576.

Fissler, T. and Ziegel, J. (2016). Higher order elicitability and Osband's principle. *Annals of Statistics*, 44(4):1680–1707.

Fissler, T., Ziegel, J., and Gneiting, T. (2015). Expected Shortfall is Jointly Elicitable with Value at Risk – Implications for Backtesting. *Risk*.

Gelman, A., Shirley, K., et al. (2011). Inference from Simulations and Monitoring Convergence. *Handbook of Markov Chain Monte Carlo*, 6:163–174.

Geweke, J. (1992). Evaluating the Accuracy of Sampling-Based Approaches to Calculating Posterior Moments. In Bernardo, J. M., Berger, J., Dawid, A. P., and Smith, J. F. M., editors, *Bayesian Statistics 4*, pages 169–193. Oxford University Press.

Giglio, S., Kelly, B., and 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. and 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., and 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., and Santoro, E. (2020). Leverage and Deepening Business-Cycle Skewness. *American Economic Journal: Macroeconomics*, 12(1):245–81.

Johnson, N. L., Kotz, S., and Balakrishnan, N. (1995). Continuous Univariate Distributions, Vol. 2, *Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics*.

Jovanovic, B. and Ma, S. (2020). Uncertainty and Growth Disasters. *NBER Working Papers*, No. 28024, National Bureau of Economic Research.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–393.

Koenker, R. and Bassett, G. (1978). Regression Quantiles. *Econometrica*, 46(1):33–50.

Koop, G. and Korobilis, D. (2014). A new index of financial conditions. *European Economic Review*, 71:101–116.

Koop, G. M. (2003). *Bayesian Econometrics*. John Wiley & Sons Inc.

McCausland, W. J., Miller, S., and Pelletier, D. (2011). Simulation smoothing for state–space models: A computational efficiency analysis. *Computational Statistics & Data Analysis*, 55(1):199–212.

Morley, J. and Piger, J. (2012). The Asymmetric Business Cycle. *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.

Omori, Y., Chib, S., Shephard, N., and Nakajima, J. (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140(2):425–449.

Orlik, A. and 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., and Sheppard, K. (2011). Nuisance parameters, composite likelihoods and a panel of GARCH models. *Statistica Sinica*, pages 307–329.

Plagborg-Møller, M., Reichlin, L., Ricco, G., and Hasenzagl, T. (2020). When is Growth at Risk? *Brookings Papers on Economic Activity*, Spring 2020.

Prasad, M. A., Elekdag, S., Jeasakul, M. P., Lafarguette, R., Alter, M. A., Feng, A. X., and Wang, C. (2019). Growth at Risk: Concept and Application in IMF Country Surveillance. *IMF Working Papers*, No. 19/36 , International Monetary Fund.

Reichlin, L., Ricco, G., and Hasenzagl, T. (2020). Financial Variables as Predictors of Real Growth Vulnerability. *CEPR Discussion Papers*, DP14322, Centre for Economic Policy Research.

Salgado, S., Guvenen, F., and Bloom, N. (2019). Skewed business cycles. *NBER Working Papers*, No. 26565, National Bureau of Economic Research.

Stock, J. H. and Watson, M. W. (2003). Has the Business Cycle Changed and Why? In *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pages 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. and 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.

Appendix A Moments of the noncentral t-distribution

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 (Hogben et al., 1961). The mean, variance, and third and fourth central moment of X are given by:

$$
\begin{aligned}
&\mathbb{E}[X] = c_{11}(\nu)\delta, && \text{if } \nu > 1,\\
&\mathbb{E}\left[(X - \mathbb{E}[X])^2\right] = c_{22}(\nu)\delta^2 + c_{20}(\nu), && \text{if } \nu > 2,\\
&\mathbb{E}\left[(X - \mathbb{E}[X])^3\right] = c_{33}(\nu)\delta^3 + c_{31}(\nu)\delta, && \text{if } \nu > 3,\\
&\mathbb{E}\left[(X - \mathbb{E}[X])^4\right] = c_{44}(\nu)\delta^4 + c_{42}(\nu)\delta^2 + c_{40}, && \text{if } \nu > 4.
\end{aligned}
$$

The coefficients have the following functional forms:

$$
\begin{aligned}
&c_{11}(\nu) = \sqrt{\frac{1}{2}\nu}\frac{\Gamma\left[\frac{1}{2}(\nu-1)\right]}{\Gamma(\frac{1}{2}\nu)}, \qquad c_{22}(\nu) = \frac{\nu}{\nu-2} - c_{11}(\nu)^2, \qquad c_{20}(\nu) = \frac{\nu}{\nu-2},\\
&c_{33}(\nu) = c_{11}(\nu)\left[\frac{\nu(7-2\nu)}{(\nu-2)(\nu-3)} + 2c_{11}(\nu)^2\right], \qquad c_{31}(\nu) = \frac{3\nu}{(\nu-2)(\nu-3)}c_{11}(\nu),\\
&c_{44}(\nu) = \frac{\nu^2}{(\nu-2)(\nu-4)} - \frac{2\nu(5-\nu)c_{11}(\nu)^2}{(\nu-2)(\nu-3)} - 3c_{11}(\nu)^4,\\
&c_{42}(\nu) = \frac{6\nu}{\nu-2}\left[\frac{\nu}{\nu-4} - \frac{(\nu-1)c_{11}(\nu)^2}{\nu-3}\right], \qquad c_{40}(\nu) = \frac{3\nu^2}{(\nu-2)(\nu-4)}.
\end{aligned}
$$

Appendix B Details on the MCMC algorithm

This appendix provides details on the MCMC algorithm and the conditional posterior distributions from which the unobserved components and parameters of the panel stochastic volatility model with time-varying skewness are drawn. The presentation closely follows Iseringhausen (2020).

Block 1: Sample the conditional mean coefficients γ from $p(\gamma|y, X_\mu, h, \delta, \lambda, \nu)$

The country-specific k-dimensional vector of regression coefficients γ_i of the conditional mean specification μ_i can be sampled as outlined in Tsionas (2002). Sampling is done country-by-country, for $i = 1, ..., N$, from the following conditional posterior distribution

$$
\gamma_i|y_i, X_{\mu_i}, h_i, \delta_i, \lambda_i, \nu \sim \mathcal{N}\left([X'_{\mu_i}\Lambda_i^{-1}X_{\mu_i}]^{-1}X'_i\Lambda_i^{-1}(\widetilde{y}_i - \delta_i \odot e^{h_i/2} \odot \lambda_i^{1/2}), e_i^h[X'_{\mu_i}\Lambda_i^{-1}X_{\mu_i}]^{-1}\right), \tag{B-1}
$$

where X_{μ_i} is a $T \times k$ matrix containing a constant, current growth, and three additional lags. Furthermore, $\Lambda_i = \text{diag}(\lambda_{i1}, ..., \lambda_{iT})$, $\widetilde{y}_i = y_i + e^{h_i/2} c_{11}(\nu)\delta_i$ and $\odot$ denotes the element-wise (Hadamard) product of two vectors. Different from Tsionas (2002), the second summand of the transformed dependent variable $\widetilde{y}_i$ reflects the fact that we consider the de-meaned version of the noncentral t-distribution.

Block 2: Sample the mixture indicators s from $p(s|y, X_\mu, h, \delta, \nu, \gamma)$

The country-by-country sampling of the mixture indicators s_i builds on the approach of Kim et al. (1998) but includes an extension to account for the fact that the specific mixture components depend on ν (which changes over MCMC iterations) and δ_{it} (which changes over MCMC iterations and time). Specifically, s_{it} is a discrete random variable with $M = 10$ possible realisations, and where each s_{it} has the following probability mass function

$$p(s_{it} = j|y_{it}, X_{\mu_{it}}, h_{it}, \delta_{it}, \nu, \gamma_i) = \frac{1}{k_{it}} q_j(\nu, \delta_{it}) p_{\mathcal{N}}\left(\widetilde{y}_{it}; h_{it} + m_j(\nu, \delta_{it}), v_j^2(\nu, \delta_{it})\right). \tag{B-2}$$

In terms of notation, $\widetilde{y}_{it} = \log((y_{it} - X_{\mu_{it}}\gamma_i)^2 + c)$, $c = 10^{-6}$ is an offset constant and $k_{it} = \sum_{j=1}^{10} q_j(\nu, \delta_{it}) p_{\mathcal{N}}\left(\widetilde{y}_{it}; h_{it} + m_j(\nu, \delta_{it}), v_j^2(\nu, \delta_{it})\right)$ is a normalising constant. The indicator sampling is operationalised by using the inverse-transform method as in Chan and Hsiao (2014).

Block 3: Sample the (log-)volatilities h from $p(h|y, X_\mu, s, \delta, \nu, \gamma, \sigma_h^2)$

First, let us specify a general state space model of the following form (see Durbin and Koopman, 2012)

$$w_t = Z_t \kappa_t + e_t, \qquad e_t \sim \mathcal{N}(0, H_t), \tag{B-3}$$

$$\kappa_{t+1} = d_t + T_t \kappa_t + R_t \eta_t, \qquad \eta_t \sim \mathcal{N}(0, Q_t), \tag{B-4}$$

where w_t is an observed realisation of a dependent variable and κ_t the unobserved state. The matrices Z_t, T_t, H_t, Q_t, R_t, and d_t are assumed to be given and thus conditioned upon. The error terms of the observation and state equation, e_t and η_t, are assumed serially uncorrelated and independent at all leads and lags. Given this general form, the specific state space representation employed in this block to sample the latent (log-)volatility series

h_i, country-by-country, is then

$$\underbrace{\widetilde{y}_{it} - m_{s_{it}}(\nu, \delta_{it})}_{w_{it}} = \underbrace{\begin{bmatrix}1\end{bmatrix}}_{Z_{it}} \underbrace{h_{it}}_{\kappa_{it}} + \underbrace{\epsilon_{it}}_{e_{it}}, \tag{B-5}$$

$$\underbrace{h_{i,t+1}}_{\kappa_{i,t+1}} = \underbrace{\begin{bmatrix}1\end{bmatrix}}_{T_{it}} \underbrace{h_{it}}_{\kappa_{it}} + \underbrace{\begin{bmatrix}1\end{bmatrix}}_{R_{it}} \underbrace{\eta_{it}}_{\eta_{it}}, \tag{B-6}$$

where $\widetilde{y}_{it} = \log((y_{it} - X_{\mu_{it}}\gamma_i)^2 + c)$, $H_{it} = v^2_{s_{it}}(\nu, \delta_{it})$ and $Q_{it} = \sigma^2_h$. We filter the unknown state variable h_{it} from the linear Gaussian state space model given by Equations (B-5) and (B-6) using recently developed sparse matrix algorithms (Chan and Jeliazkov, 2009; McCausland et al., 2011). In particular, we follow Chan and Hsiao (2014) who show how to efficiently sample the (log-)volatilities h_{it} based on these algorithms. A detailed explanation of the so-called precision sampler can be found on pp. 5-8 in Chan and Hsiao (2014).

Block 4: Sample the latent state λ from $p(\lambda|y, X_\mu, h, \delta, \nu, \gamma)$

The sampling approach for the latent state variable λ_i, country-by-country, follows Tsionas (2002). The conditional distribution of each λ_{it} is

$$p(\lambda_{it}|y_{it}, X_{\mu_{it}}, h_{it}, \delta_{it}, \nu, \gamma_i) \propto \lambda_{it}^{-(\nu+3)/2} \exp\left[-\frac{u_{it}^2/e^{h_{it}} + \nu}{2\lambda_{it}} + \delta_{it}(u_{it}/e^{h_{it}/2})\lambda_{it}^{-1/2}\right], \tag{B-7}$$

where $u_{it} = y_{it} - X_{\mu_{it}}\gamma_i + e^{h_{it}/2}c_{11}(\nu)\delta_{it}$. If $\delta_{it} = \delta = 0$, the shocks are Student t-distributed and sampling λ_{it} from its inverse-gamma conditional posterior distribution is straightforward as in e.g. Chan and Hsiao (2014). In the general noncentral case the conditional distribution is non-standard and we require acceptance sampling. Tsionas (2002) notes that the conditional distribution of $w_{it} = \lambda_{it}^{-1/2}$ is log-concave and each w_{it} follows the distribution

$$p(w_{it}|y_{it}, X_{\mu_{it}}, h_{it}, \delta_{it}, \nu, \gamma_i) \propto w_{it}^{\nu} \exp\left(-\frac{u_{it}^2/e^{h_{it}} + \nu}{2} w_{it}^2 + \frac{\delta_{it} u_{it}}{e^{h_{it}/2}} w_{it}\right). \tag{B-8}$$

This distribution is part of a more general family of distributions with kernel function

$$f(x) \propto x^{N-1} \exp(-(A/2)x^2 + Bx), \tag{B-9}$$

where $N_1 = \nu + 1$, $A = u^2/e^h + \nu$ and $B = \delta u/e^{h/2}$.

As a proposal density for the acceptance sampling $g(x) \sim \text{Gamma}(N_1, \theta^*)$ is used, where

$\theta^* = N_1/x^*$ and x^* is the positive root that is the solution to

$$A_{it}x^2 - B_{it}x - N_1 = 0. \tag{B-10}$$

The candidate draw w_{it}^* is accepted with probability

$$R = \exp(r^* - r), \tag{B-11}$$

where $r^* = \log(f(x)/g(x))$ evaluated at w_{it}^* and $r = \log(f(x)/g(x))$ evaluated at x^*. Specifically,

$$r^* = -(A_{it}/2){w_{it}^*}^2 + (B_{it} + \theta^*)w_{it}^* - N_1 \log(\theta^*), \tag{B-12}$$

$$r = -(A_{it}/2){x^*}^2 + (B_{it} + \theta^*)x^* - N_1 \log(\theta^*). \tag{B-13}$$

Once a candidate draw w_{it}^* has been accepted, the original state variable is simply recovered as $\lambda_{it} = {w_{it}^*}^{-2}$.

Block 5: Sample the degrees of freedom ν from $p(\nu|\lambda)$

The description of the sampling approach for the degrees of freedom ν, which is pooled across countries, closely follows Chan and Hsiao (2014). In particular, the (log-)conditional posterior distribution of ν remains identical compared to the (symmetric) Student-t case (Tsionas, 2002). Using the $NT \times 1$ vector $\widetilde{\lambda}$, which stacks the N country-specific column vectors λ_i, the (log-)conditional posterior distribution of interest is

$$\log p(\nu|\widetilde{\lambda}) = \frac{NT\nu}{2}\log(\nu/2) - NT\log\Gamma(\nu/2) - (\nu/2+1)\sum_{j=1}^{NT}\log\widetilde{\lambda}_j - \frac{\nu}{2}\sum_{j=1}^{NT}\widetilde{\lambda}_j^{-1} + k, \tag{B-14}$$

for $0 < \nu < \bar{\nu}$ and k is a normalisation constant. The first and second derivative of this distribution with respect to ν are given by

$$\frac{d\log p(\nu|\widetilde{\lambda})}{d\nu} = \frac{NT}{2}\log(\nu/2) + \frac{NT}{2} - \frac{NT}{2}\Psi(\nu/2) - \frac{1}{2}\sum_{j=1}^{NT}\log\widetilde{\lambda}_j - \frac{1}{2}\sum_{j=1}^{NT}\widetilde{\lambda}_j^{-1}, \tag{B-15}$$

$$\frac{d^2\log p(\nu|\widetilde{\lambda})}{d\nu^2} = \frac{NT}{2\nu} - \frac{NT}{4}\Psi'(\nu/2), \tag{B-16}$$

where $\Psi(x) = \frac{d}{dx}\log\Gamma(x)$ and $\Psi'(x) = \frac{d}{dx}\Psi(x)$ denote the digamma and trigamma function, respectively. The first and second derivatives can be evaluated easily, and thus $\log p(\nu|\widetilde{\lambda})$ can

be maximised by well-known algorithms (e.g. the Newton-Raphson method). In addition, the mode and the negative Hessian evaluated at the mode, $\hat{\nu}$ and K_ν, are obtained. Lastly, to sample a draw of ν, a Metropolis-Hastings step is implemented with proposal density $\mathcal{N}(\hat{\nu}, K_\nu^{-1})$.

Block 6: Sample the latent noncentrality parameter δ from $p(\delta|y, X_\mu, h, \lambda, \nu, \gamma, X_\delta, \phi, \beta, \sigma_\delta^2)$

Similarly to Iseringhausen (2020), for the sampling of the time-varying noncentrality parameter δ_i, country-by-country, we explore the following state space model

$$\underbrace{\widetilde{y}_{it}}_{w_{it}} = \underbrace{\left[\lambda_{it}^{1/2} - c_{11}(\nu)\right]}_{Z_{it}} \underbrace{\delta_{it}}_{\kappa_{it}} + \underbrace{\epsilon_{it}}_{e_{it}}, \tag{B-17}$$

$$\underbrace{\delta_{i,t+1}}_{\kappa_{i,t+1}} = \underbrace{[X_{\delta_{it}}\beta]}_{d_{it}} + \underbrace{\left[\phi\right]}_{T_{it}} \underbrace{\delta_{it}}_{\kappa_{it}} + \underbrace{\left[1\right]}_{R_{it}} \underbrace{\omega_{it}}_{\eta_{it}}, \tag{B-18}$$

where $\widetilde{y}_{it} = (y_{it} - X_{\mu_{it}}\gamma_i)e^{-h_{it}/2}$, $H_{it} = \lambda_{it}$, and $Q_{it} = \sigma_\delta^2$. The observation Equation (B-17) is derived by rewriting the model in Equation (9). Again, we use the routine developed in Chan and Jeliazkov (2009) to obtain a sample of $\delta_i = (\delta_{i1}, ..., \delta_{iT})$ for $i = 1, ..., N$.

Block 7: Sample the coefficients of the asymmetry process δ from $p(\phi, \beta|\delta, X_\delta, \sigma_\delta^2)$

Conditional on the latent asymmetry parameter δ, the task of sampling the pooled coefficients on the explanatory variables that determine skewness, simplifies to the linear Bayesian regression problem that is treated in standard textbooks (see e.g. Koop, 2003) with an additional acceptance-rejection step to ensure that ϕ remains in the stationary region. Defining $\widetilde{\beta} = (\phi, \beta')'$, $\widetilde{\beta}_0 = (\phi_0, \beta_0')'$, $\sigma^2_{\widetilde{\beta}_0} = \sigma^2_{\phi 0} = \sigma^2_{\beta 0}$, the conditional posterior is

$$\widetilde{\beta}|\delta, X_\delta, \sigma_\delta^2 \sim \mathcal{N}(\widehat{\widetilde{\beta}}, \Sigma_{\widetilde{\beta}}), \tag{B-19}$$

$$\widehat{\widetilde{\beta}} = \left(\sigma^2_{\widetilde{\beta}_0} I^{-1}_{K+1} + \frac{1}{\sigma_\delta^2}\widetilde{X_\delta}'\widetilde{X_\delta}\right)^{-1} \left(\sigma^2_{\widetilde{\beta}_0} I^{-1}_{K+1}\widetilde{\beta}_0 + \frac{1}{\sigma_\delta^2}\widetilde{X_\delta}'\widetilde{X_\delta}\right)^{-1}, \tag{B-20}$$

$$\Sigma_{\widetilde{\beta}} = \left(\sigma^2_{\widetilde{\beta}_0} I^{-1}_{K+1} + \frac{1}{\sigma_\delta^2}\widetilde{X_\delta}'\widetilde{\delta}\right)^{-1}, \tag{B-21}$$

where $\widetilde{X_\delta} = [\delta_{t-1}', X_{\delta_1}', ..., X_{\delta_N}']'$, $\widetilde{\delta} = [\delta_1', ..., \delta_N']'$, and I_{K+1} is the identity matrix of dimension $K+1$. After obtaining a draw of $\widetilde{\beta}$, the stationarity condition of ϕ is checked and if not fulfilled, sampling of $\widetilde{\beta}$ is repeated until a draw of ϕ lies within the unit circle.

Block 8: Sample the shock variances σ_h^2 from $p(\sigma_h^2|h)$ and σ_δ^2 from $p(\sigma_\delta^2|\delta, X_\delta, \phi, \beta)$

The innovation variances of the (log-)volatility h_{it} and the noncentrality parameter δ_{it} are pooled across countries and have inverse-gamma conditional posterior distributions (Kim et al., 1998). The conditional posterior distribution of σ_h^2 is

$$\sigma_h^2|h \sim \mathcal{IG}(c_{h0} + N(T-1)/2, C_h), \tag{B-22}$$

where notation follows Chan and Hsiao (2014) and where C_h is defined as,

$$C_h = C_{h0} + \left[\sum_{i=1}^{N}\sum_{t=2}^{T}(h_{it} - h_{i,t-1})^2\right]/2. \tag{B-23}$$

The conditional posterior distribution of σ_δ^2 is

$$\sigma_\delta^2|\delta, X_\delta, \phi, \beta \sim \mathcal{IG}(c_{\delta 0} + N(T-1)/2, C_\delta), \tag{B-24}$$

where C_δ is defined as,

$$C_\delta = C_{\delta 0} + \left[\sum_{i=1}^{N}\sum_{t=2}^{T}(\delta_{it} - \phi\delta_{i,t-1} - X_{\delta_{it}}\beta)^2\right]/2. \tag{B-25}$$

Appendix C Monte Carlo simulation

This section uses a small simulation exercise to study the performance of the MCMC algorithm presented in Section 2.2 and to show its ability to recover the true data generating parameters. We generate $1,000$ datasets from the following data generating process (DGP), which constitutes a simplified version of the main model that, for the sake of simplicity, has no predictive dimension ($h = 0$), excludes the conditional mean specification, and only considers one exogenous variable in the asymmetry equation:

$$y_{it} = e^{h_{it}/2}\varepsilon_{it}, \qquad i = 1, ..., N, \quad t = 1, ..., T, \tag{C-1}$$

$$\varepsilon_{it} = u_{it} - \mathbb{E}[u_{it}], \qquad u_{it} \sim \mathcal{NCT}(\nu, \delta_{it}), \tag{C-2}$$

$$h_{it} = h_{i,t-1} + \eta_{it}, \qquad \eta_{it} \sim \mathcal{N}(0, \sigma_h^2), \tag{C-3}$$

$$\delta_{it} = \phi\delta_{i,t-1} + \beta_0 + \beta_1 x_{it} + \omega_{it}, \qquad \omega_{it} \sim \mathcal{N}(0, \sigma_\delta^2), \quad |\phi| < 1, \tag{C-4}$$

$$x_{it} = \rho x_{i,t-1} + e_{it}, \qquad e_{it} \sim \mathcal{N}(0, \sigma_x^2), \quad |\rho| < 1. \tag{C-5}$$

The true model parameters are set as follows: $\phi = 0.7$, $\beta_0 = -0.2$, $\beta_1 = 1$, $\sigma_h^2 = 0.1^2$, $\sigma_\delta^2 = 0.1^2$, $\nu = 6$, $\rho = 0.7$, and $\sigma_x^2 = 0.1$. To avoid unrealistic realisations of the (unrestricted)

random walk process for h_{it}, when generating the data we restrict h_{it} to lie in the interval $[-1, 1]$. Moreover, to reduce the impact of initial conditions on the processes for h, δ, and x, we run the DGP for 50 periods before generating each dataset.

After obtaining a dataset, we estimate the model excluding the Gibbs block for the conditional mean, and using as input only y and x. The prior configurations are identical to those used in the main analysis. Table C-1 shows the results of the Monte Carlo simulation. In general, the algorithm performs well in recovering the underlying parameters of the DGP. For a relatively small sample size comparable to the one of our macroeconomic dataset ($N = 10$ and $T = 200$), the average posterior means deviate somewhat from the true values. This is not surprising since the priors are not centered around the true values thus inducing a distortion and, more generally, Bayesian estimators are not unbiased (Box, 1971). However, the Monte Carlo distribution covers the true value with large probability mass in all cases.

Table C-1: Results Monte Carlo simulation

	$N = 10$ / $T = 200$						$N = 10$ / $T = 1,000$					
	ϕ	β_0	β_1	σ_h^2	σ_δ^2	ν	ϕ	β_0	β_1	σ_h^2	σ_δ^2	ν
True	0.70	-0.20	1.00	0.01	0.01	6.00	0.70	-0.20	1.00	0.01	0.01	6.00
Mean	0.60	-0.26	1.16	0.01	0.02	6.67	0.68	-0.20	0.96	0.01	0.02	5.94
Std.	0.11	0.11	0.32	0.00	0.01	1.89	0.03	0.03	0.11	0.00	0.01	0.45
CD	0.48	0.48	0.46	0.47	0.47	0.46	0.48	0.48	0.48	0.49	0.43	0.46
IF	275.5	237.3	296.9	18.2	123.3	61.3	159.4	137.9	216.0	24.0	505.9	25.2

Note: The table shows the average posterior means across all Monte Carlo samples and the standard deviations of the estimated posterior means. *CD* is the average p-value of the convergence diagnostic of Geweke (1992) and *IF* is the average inefficiency factor described in Chib (2001).

To show that these minor inaccuracies are purely due to limited information in a smaller dataset rather than actual problems with the MCMC algorithm, we also consider a simulation exercise with a longer time series dimension ($T = 1,000$). The results indicate that estimation precision quickly increases with rising T. In this case, the Monte Carlo averages are even closer to the true values with only small dispersion.

The convergence diagnostic of Geweke (1992) shows that the null hypothesis of a converged Markov chain can, on average, not be rejected for all six model parameters in both simulations. Finally, the algorithm is somewhat inefficient in sampling uncorrelated draws from the conditional posterior distributions of the parameters in the asymmetry equation. Similar results have already been reported in Iseringhausen (2020) for the original model, but we note that an even larger number of Gibbs iterations does not alter the results meaningfully.

Appendix D Data

Table D-1: Data description, sources, and information on imputation

Variable	Description	Source	Imputed obs.
y	Quarter-over-quarter real GDP growth (in percent, seasonally adjusted)	OECD	—
FCI	Financial Conditions Index	IMF (2017, 2018)	CA: 1973Q1–1980Q4 IT: 1973Q1–1980Q4 ES: 1973Q1–1980Q4 SE: 1973Q1–1980Q4 CH: 1973Q1–1980Q4
TS	Term spread: 10-year interest rate minus 3-month interest rate	OECD	IT: 1973:Q1–1990:Q4 JP: 1973:Q1–2002Q1 ES: 1973:Q1–1979Q4 SE: 1973:Q1–1986Q3 CH: 1973:Q1–1973Q4 GB: 1973:Q1–1985Q4
HP	Quarter-over-quarter percentage change of seasonally adjusted real house price index	OECD	—
WUI	World Uncertainty Index	WorldUncertaintyIndex.com (Ahir et al., 2018)	—

Note: This table contains details on the variables used in the estimations, their definitions, and sources. The last column mentions the countries and periods for which missing values have been imputed. The imputed values are taken directly from the available dataset of Brownlees and Souza (2021). These authors impute missing observations with the cross-sectional average of the other countries in each period, which are re-scaled such that the standard deviation of the original series remains unchanged. While for the FCI missing values are imputed based on the six countries with a complete FCI series that are also used in this paper, for the term spread some cross-sectional averages are based on more countries that are part of the larger 24-country sample of Brownlees and Souza (2021). For further details, the reader is referred to Brownlees and Souza (2021).

Figure D-1: Time series plots of (non-standardised) explanatory variables

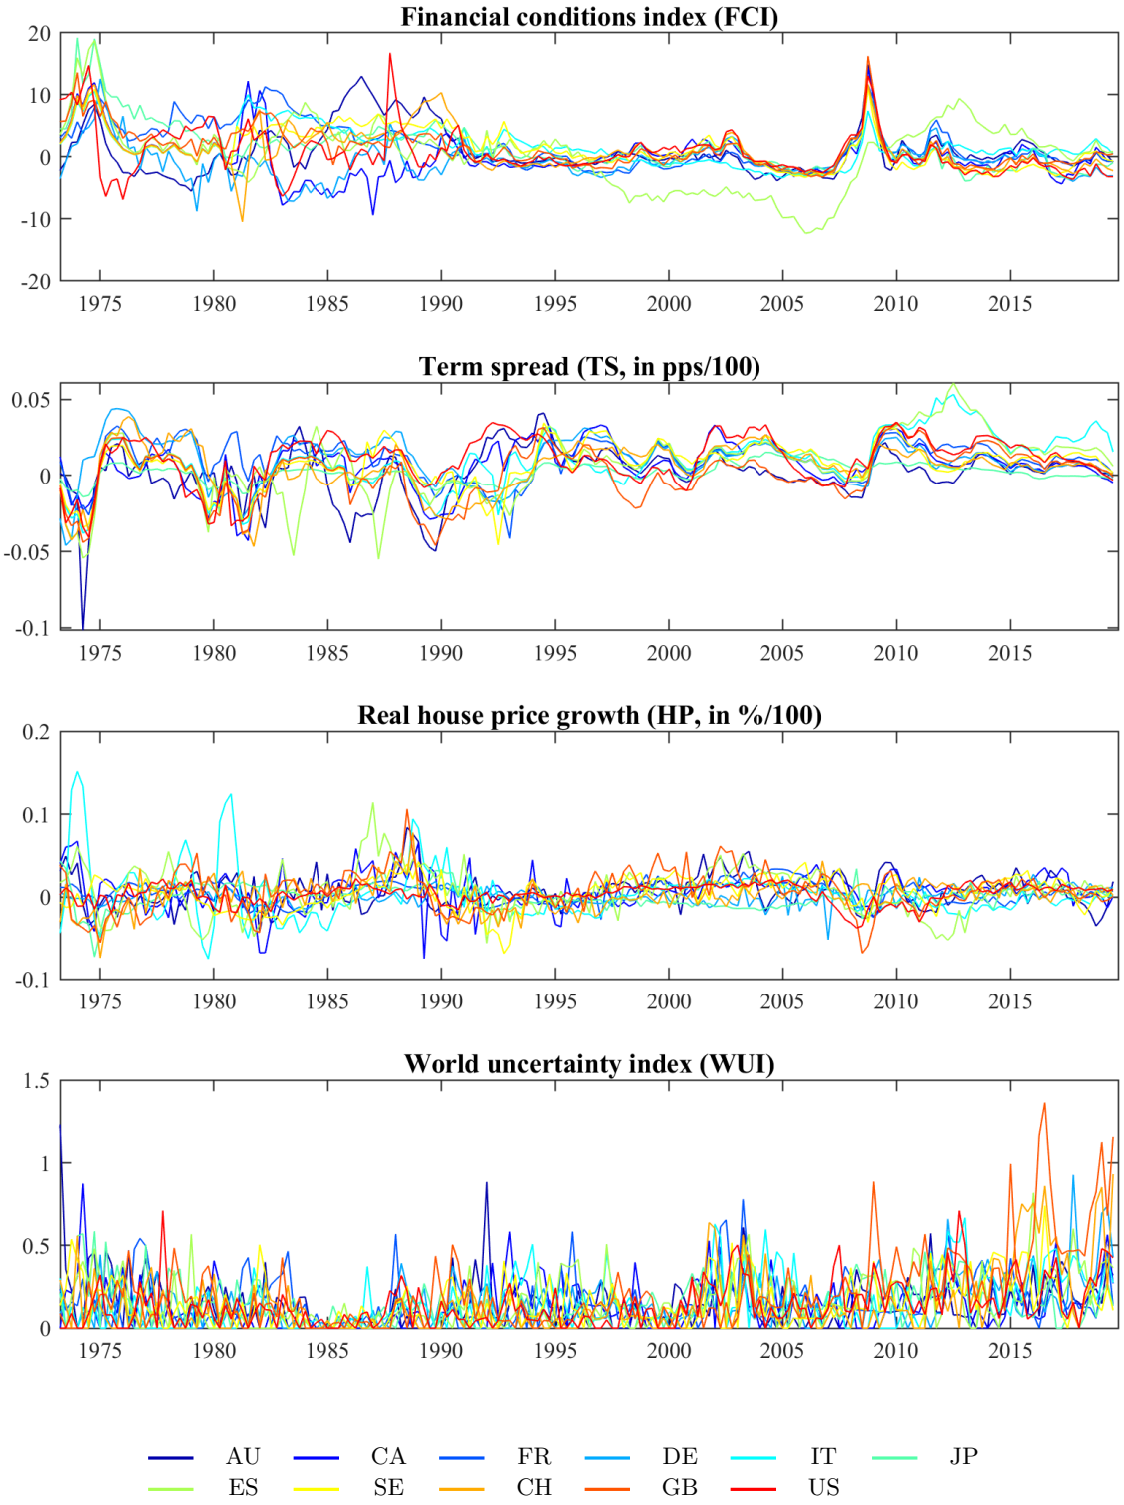

Appendix E Additional results

Figure E-1: Out-of-sample one-step-ahead expected shortfall/longrise

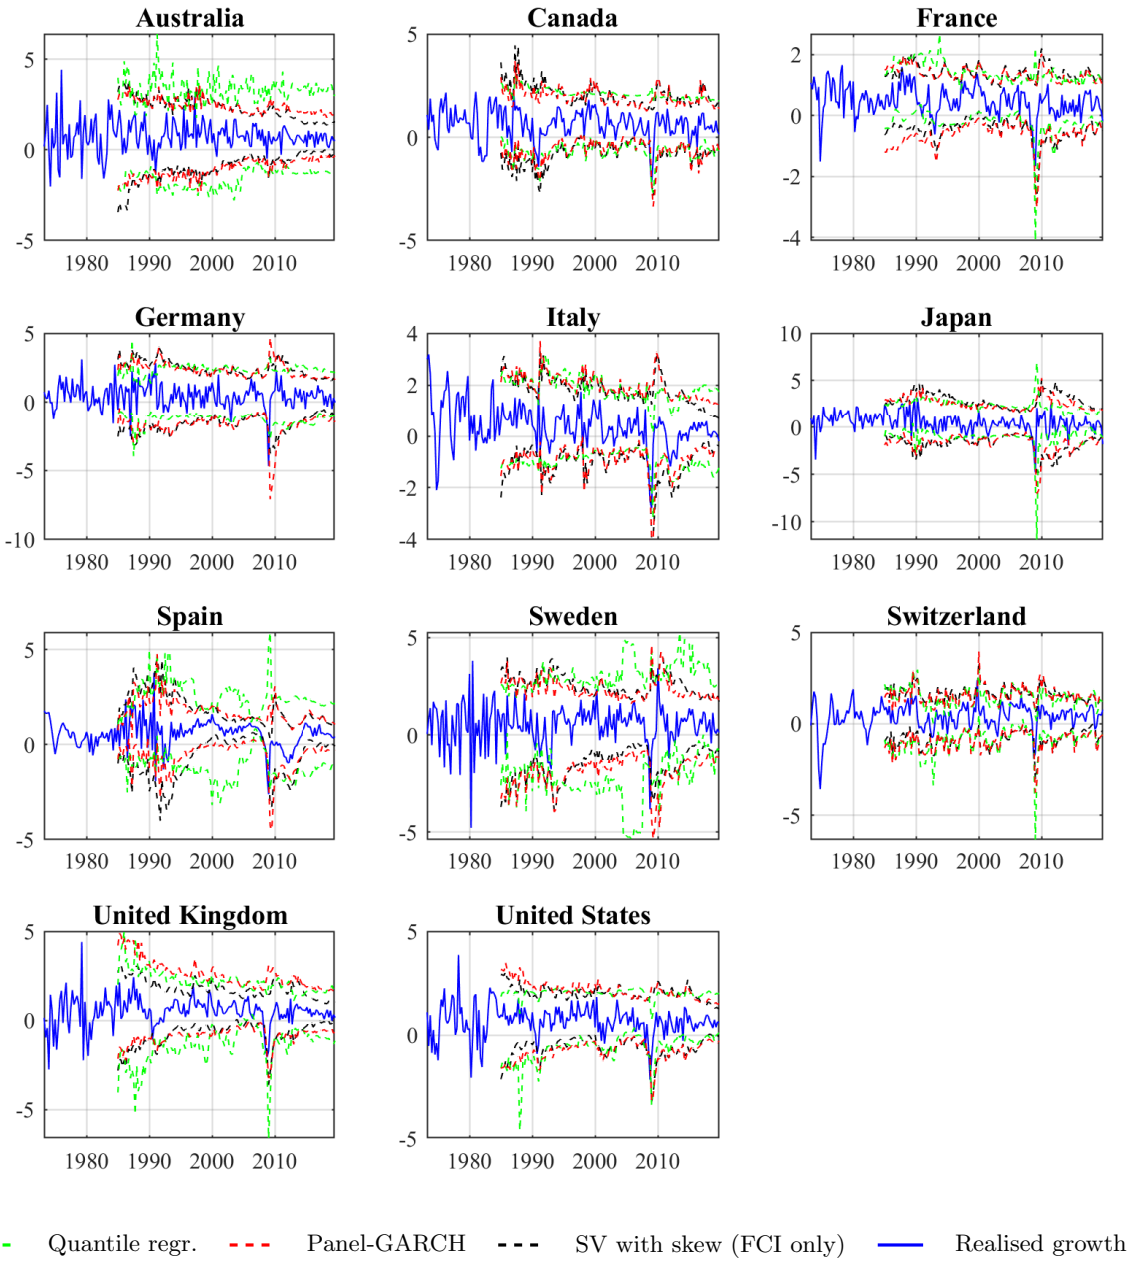

Table E-1: Robustness checks and alternative specifications for forecasting exercise

	$h=1$			$h=2$			$h=3$			$h=4$		
Country-specific quantile regressions (FCI + TS + HP + WUI)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	64	**55**	0.097	73	55	**0.095**	64	**82**	0.109	55	**91**	0.114
$GaR_{95\%}$	82	82	0.081	**91**	**100**	0.092	**100**	91	0.088	**82**	**100**	0.095
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.400	0.628		**0.338**	0.558		**0.429**	0.736		0.520	0.793	
$EL_{95\%}$	0.367	0.717		**0.284**	0.790		0.289	0.742		0.337	0.805	
SV model with time-varying skewness (FCI + y_t)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	100	64	0.082	100	82	0.092	100	**100**	0.100	91	100	0.103
$GaR_{95\%}$	100	100	0.068	100	82	0.069	100	82	0.074	100	73	0.077
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.177	0.417		**0.275**	**0.509**		0.417	0.610		0.434	0.669	
$EL_{95\%}$	0.164	0.628		**0.171**	**0.635**		0.162	0.653		**0.189**	0.664	
SV model with symmetric shocks (FCI in cond. mean)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	91	55	**0.082**	91	73	0.092	**100**	64	0.100	91	73	0.104
$GaR_{95\%}$	100	100	0.068	91	82	0.071	100	**100**	0.075	91	73	0.078
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	**0.188**	**0.412**		**0.295**	**0.512**		**0.415**	**0.612**		**0.422**	0.682	
$EL_{95\%}$	0.184	0.629		**0.187**	0.648		0.183	0.658		**0.184**	0.667	
SV model with time-varying skewness (FCI in asym. eq. + FCI in cond. mean)												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	91	64	**0.080**	91	64	0.092	100	91	0.100	**100**	64	0.104
$GaR_{95\%}$	100	100	**0.067**	100	82	0.070	100	**100**	0.074	91	73	0.077
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	**0.146**	**0.398**		0.285	**0.508**		**0.383**	**0.608**		**0.417**	0.683	
$EL_{95\%}$	0.172	**0.626**		0.178	0.645		0.169	0.655		**0.177**	0.665	
SV model with symmetric shocks and heterogenous σ_h^2												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	100	36	0.083	100	**91**	**0.091**	**100**	**91**	**0.099**	**100**	91	**0.101**
$GaR_{95\%}$	100	100	0.068	100	82	0.070	100	**91**	0.074	91	73	0.077
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.205	**0.426**		0.308	**0.515**		0.441	**0.608**		0.447	**0.650**	
$EL_{95\%}$	0.181	0.630		0.203	0.643		0.190	0.658		0.217	0.669	
SV model with time-varying skewness (FCI) and heterogenous σ_h^2												
	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL	DQ_{uc}	DQ_{hits}	TL
$GaR_{5\%}$	100	45	**0.081**	100	73	**0.091**	100	91	**0.099**	**100**	91	**0.101**
$GaR_{95\%}$	100	100	0.068	100	**91**	0.069	100	**91**	0.073	91	73	**0.076**
	EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score		EKP	VaR-ES score	
$ES_{5\%}$	0.163	**0.412**		0.287	**0.509**		**0.410**	**0.603**		0.443	**0.652**	
$EL_{95\%}$	0.153	0.628		**0.167**	0.639		0.165	0.653		0.196	0.664	

Note: See Table 2. Bold numbers indicate an improvement in the measure compared to the respective baseline model shown in Table 2, i.e. the 'FCI only' quantile regression, the baseline SV model with symmetric shocks, or the baseline SV model with time-varying skewness ('FCI only').

European Stability Mechanism

6a Circuit de la Foire Internationale
L-1347 Luxembourg

Tel: +352 260 292 0

www.esm.europa.eu

info@esm.europa.eu