

Wolf, Elias

Working Paper

Estimating growth at risk with skewed stochastic volatility models

Discussion Paper, No. 2022/2

Provided in Cooperation with:

Free University Berlin, School of Business & Economics

Suggested Citation: Wolf, Elias (2022) : Estimating growth at risk with skewed stochastic volatility models, Discussion Paper, No. 2022/2, Freie Universität Berlin, School of Business & Economics, Berlin,
<https://nbn-resolving.de/urn:nbn:de:kobv:188-refubium-33910-5>

This Version is available at:

<https://hdl.handle.net/10419/249921>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Estimating Growth at Risk with Skewed Stochastic Volatility Models

Elias Wolf

School of Business & Economics

Discussion Paper

Economics

2022/2

Estimating Growth at Risk with Skewed Stochastic Volatility Models*

Elias Wolf[†]

December 18, 2021

Abstract

This paper proposes a Skewed Stochastic Volatility (SSV) model to model time varying, asymmetric forecast distributions to estimate Growth at Risk as introduced in Adrian, Boyarchenko, and Giannone's (2019) seminal paper "Vulnerable Growth". In contrary to their semi-parametric approach, the SSV model enables researchers to capture the evolution of the densities parametrically to conduct statistical tests and compare different models. The SSV-model forms a non-linear, non-gaussian state space model that can be estimated using Particle Filtering and MCMC algorithms. To remedy drawbacks of standard Bootstrap Particle Filters, I modify the Tempered Particle Filter of Herbst and Schorfheide's (2019) to account for stochastic volatility and asymmetric measurement densities. Estimating the model based on US data yields conditional forecast densities that closely resemble the findings by Adrian et al. (2019). Exploiting the advantages of the proposed model, I find that the estimated parameter values for the effect of financial conditions on the variance and skewness of the conditional distributions are statistically significant and in line with the intuition of the results found in the existing literature.

Keywords: Growth at Risk, Macro Finance, Bayesian Econometrics, Particle Filters

JEL classification: C10, E32, E58, G01.

*This paper has benefited from valuable comments and suggestions by Frank Schorfheide, Carlos Montes-Galdon, Joan Paredes, Dieter Nautz, Lars Winkelmann, Helmut Lütkepohl, Yves Schüler, Lea Sieckmann, Max Diegel, Michelle Lenza, Marek Jarochinski, Matteo Cicarelli as well as other participants of the ECB Brownbag Seminar, the 11th RCEA Money Macro Finance Conference in June 2021 and the "Topics in Time Series Econometrics"-Workshop 2020 and 2021 in Tornow. I am grateful for helpful support from the High Performance Cluster at the Freie Universität.

[†]Contact address: Freie Universität Berlin, Chair of Econometrics, Boltzmannstraße 20, 14195 Berlin, Germany. E-mail: e.wolf@fu-berlin.de

1 Introduction

Non-normality and asymmetric distributions have played a role in economic research for a long time (see for example Harvey and Siddique (1999), Engle (2011) or Orlik and Veldkamp (2014)). Yet, while most research was focused on financial markets, the seminal paper "Vulnerable Growth" of Adrian et al. (2019) brought non-normality and asymmetric forecast densities in the spotlight of macro-economists and central banks all around the world. Borrowing from the Value at Risk approach in financial econometrics, Growth at Risk seeks to capture asymmetric tail risks to GDP growth conditioned on national financial conditions and provides policy makers with a risk measure to evaluate economic stability (Prasad, Elekdag, Jeasakul, Lafarguette, Alter, Feng, Wang, and Gust (2019)). However, while their original semi-parametric two-step approach is straight forward, it lacks the possibility to capture the evolution of volatility and skewness in a parametric form to conduct statistical inference and compare the effects of national financial conditions or other exogenous variables. To capture the stylized facts in the existing literature, this paper proposes a Skewed Stochastic Volatility (SSV) model where errors follow a skew Normal distribution (Azzalini (2013)) with the scale and shape parameter modeled as latent states. Building on established estimation methods for non-linear, non-Gaussian state space model as well as discrete Stochastic Volatility models with symmetric densities (see for example Flury and Shephard (2011)), the SSV model can be estimated with particle filtering methods and Monte Carlo Markow Chain (MCMC) Methods. Taking into consideration the well-known weakness of standard Bootstrap Particle Filters to be sensitive to extreme values (Doucet, de Freitas, and Gordon (2001)), I use the Tempered Particle Filter by Herbst and Schorfheide (2019) to obtain robust estimates of the likelihood and latent states. To reduce the runtime of the estimation I modify the tempering schedule to take the asymmetry of the measurement density into account. This results in less tempering iterations. Estimating the model based on US data, I find that the model can capture the stylized facts by Adrian et al. (2019) and other recent studies such as delle Monache, de Polis, and Petrella (2021). Furthermore, the results provide further statistical evidence that national financial conditions have a significant impact on the shape and scale of the one-period ahead forecasting distribution given the 90% credible sets obtained from the posterior distribution of the static model parameters. With a cor-

relation coefficient of -0.41, I find that volatility and skewness have an inverse relationship indicating that downside risks increase in times of high volatility.

The rest of the paper is structured as follows: Section 2 provides an overview of the Growth at Risk concept and the methodology developed by Adrian et al. (2019). The Skewed Stochastic Volatility model and its estimation is introduced in Section 3 and 4 respectively. Section 5 discusses the results for the US. Section 6 concludes.

2 Growth at Risk

The concept of Growth at Risk was introduced with the seminal paper by Adrian et al. (2019) who analyse the variation of the one period ahead forecast distribution of US GDP gdp_{t+1} conditional on the national financial conditions index dubbed $nfcit$.¹ To gain information about the time variation the the one-period ahead forecasting distribution the authors develop a two step semi-parametric procedure. In the first step different quantiles of the distributions are obtained by running quantile regressions of the form

$$gdp_{t+1} = \beta_0 + \beta_1 nfcit_t + \varepsilon_{t+1} \quad (1)$$

where the vector of parameters $\beta = (\beta_0, \beta_1)'$ is chosen to minimize the Koenker Basset loss (Koenker and Bassett (1978)) defined as

$$\mathcal{L}_{KB} = \sum (\tau \cdot \mathbf{1}_{(y_{t+1} > x'_t \beta)} |y_{t+1} + x'_t \beta| + (1 - \tau) \cdot \mathbf{1}_{(y_{t+1} < x'_t \beta)} |y_{t+1} + x'_t \beta|)$$

where $\mathbf{1}_A$ denotes the indicator function, $x_t = (1, nfcit_t)'$ is the vector of explanatory variables including an intercept and $0 \leq \tau \leq 1$ indicates the τ^{th} sample quantile. In a second step, the authors match the predicted 5, 25, 75 and 95 % quantiles from the regressions to the theoretical moments of the skewed T distribution defined by Azzalini (2013). The skew density function takes the form

$$skew T(y|\xi, \omega, \alpha, \nu) = \frac{2}{\omega} \cdot t(z|\nu) \cdot T(\alpha z|\nu + 1) \quad \text{with} \quad z = \frac{y - \xi}{\omega} \quad (2)$$

¹The National Financial Conditions Index is given by the fist principal component of a large number of financial variables and released by the Chicago Fed.

with ξ, ω, α, ν controlling the location, scale, shape and kurtosis of the distribution. Consequently, the resulting densities are quite flexible and not constrained to be symmetric around the mean. Estimating the model based on US data from the 1980s up to 2017, Adrian et al. (2019) document the following properties of the one period ahead forecast distributions:

- (1) Lower quantiles of the conditional forecast distribution vary a lot over over time while the upper quantiles remain relatively stable.
- (2) A deterioration of national financial conditions coincides with increases in the interquartile range and decreases the mean.
- (3) Distributions are more symmetric in normal times and become left skewed in recessionary periods

Since then, this two-step approach has been applied to analyze time-varying forecast distributions of European growth rates (de Santis and van der Veken (2020)) or other macroeconomic variables such as inflation (López-Salido and Loria (2020)). However, while the aforementioned methodology is straight forward to implement, its semi-parametric nature has the drawback that it does not provide a clear representation of the evolution of the time variation in the volatility and asymmetry of the conditional forecast distributions. This prevents to capture the effect of the national financial conditions or any other explanatory variable a researcher might include on the second and third moment parametrically. Consequently, it is also not possible to compare and interpret these effects or conduct statistical inference to determine the relevancy of different variables on different moments of the distributions. To remedy these shortcomings this paper proposes an alternative modeling approach to measure Growth at Risk using a Skewed Stochastic Volatility model that can capture all of the above features of the forecast distributions in a parametric form by providing a law of motion for both volatility and skewness. The model is estimated using Bayesian Methods which yields posterior distributions of the parameters to also conduct statistical inference.

3 Skewed Stochastic Volatility Model

While Adrian et al. (2019) include a simple parametric model that can capture time variation in the first two moments, the resulting forecast distributions remain symmetric and Gaussian. Similarly, the work of Carriero, Clark, and Marcellino (2020) who use a time-varying volatility specification in large Bayesian VAR models can only model time varying symmetric forecasting densities and is thus constraint to capture only skewness in the unconditional density. Montes-Galdon and Ortega (2022) introduce an SVAR with time-varying skew Normal errors, but without including the effect of national financial conditions on the volatility of the structural errors. Other recent contributions in a univariate framework are delle Monache et al. (2021) who use the Generalized Autoregressive Score (GAS) Model framework developed by Creal, Koopman, and Lucas (2013) and Iseringhausen (2021) who develops a panel model of time varying skewness model for a number of different countries. Yet, while both approaches are similar in spirit to the model in this paper, their estimation critically depends on the chosen model specification and distribution. The score driven approach by delle Monache et al. (2021) relies on the gradients of the error distribution which makes further adaptations of the model equations or experiments with other distribution families cumbersome or infeasible while the later study requires conditional conjugacy of the priors. Furthermore, while the model by Iseringhausen (2021) does include national financial conditions, their effect on the volatility of the forecasting distribution only arises indirectly. The paper also relates to an older strain of literature such as Hansen (1994) or Engle and Manganelli (2004) that seeks to model asymmetric densities to obtain estimates of the Value at Risk for financial data. To the best of my knowledge, this is the first approach estimating skewed conditional densities using state of the art Particle Filtering techniques in combination with MCMC-Methods. In contrary to the aforementioned papers, the Skewed Stochastic Volatility model and the estimation strategy proposed in this paper can capture all features documented by Adrian et al. (2019) in a straight forward fashion and is quite flexible with regards to other distributions families and model specifications.

Related to the skew T distribution used by Adrian et al. (2019), the SSV model is based on the skewed Normal distribution of Azzalini (2013) with a probability density function

given by

$$skew \mathcal{N}(y|\mu, \sigma, \alpha) = \frac{2}{\sqrt{(2\pi)}\sigma} e^{-\frac{(y-\mu)^2}{2\sigma^2}} \int_{-\infty}^{\alpha \frac{y-\mu}{\sigma}} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt \quad (3)$$

This distribution has two parameters for location (μ) and scale (σ) plus an additional parameter α that determines its shape. Figure 1 shows how different values for α affect the skewness of the distribution function. While $\alpha < 0$ results in a left skewed distribution $\alpha > 0$ skews the distribution to the right. Setting $\alpha = 0$ recovers the symmetric Normal distribution. Additionally, the skew Normal distribution can capture excess




Figure 1: Skewed Normal distribution for different values of the shape parameter α .

Notes: Negative values of α tilt the distribution to the left while positive values skew the distribution to the right. Setting $\alpha = 0$ yields the standard Normal distribution.

Kurtosis which is a function of the scale and shape parameters as discussed in Azzalini (2013). Thus, the skewed Normal distribution allows modelling all moments of interest while maintaining a parsimonious modelling approach. The resulting Skewed Stochastic Volatility Model (SSV) takes the form

$$gdp_{t+1} = \gamma_0 + \gamma_1 nfcit + \varepsilon_{t+1} \quad \text{with} \quad \varepsilon_t \sim skew \mathcal{N}(0, \sigma_t, \alpha_t) \quad (4)$$

and

$$\ln(\sigma_t^2) = \delta_{1,0} + \delta_{1,1} nfcit + \delta_{1,2} \ln(\sigma_{t-1}^2) + \nu_{1,t} \quad (5)$$

$$\alpha_t = \delta_{2,0} + \delta_{2,1} nfcit + \delta_{2,2} \alpha_{t-1} + \nu_{2,t} \quad (6)$$

where $\nu_{1,t}$ and $\nu_{2,t}$ are assumed to be uncorrelated Gaussian White Noise innovations. Equations (4) to (6) form a non-linear, non-Gaussian state space model with $\ln(\sigma_t)$ and α_t as latent states, that can capture all characteristics found by Adrian et al. (2019). Including $nfcit$ in both state equations 5 and 6 introduce time variation of the scale and shape of the forecast distribution that depends on the national financial conditions. Furthermore, the lagged states allow for persistence in the evolution of the state variables. Due to the increased complexity of the model arising from non-linearity and non-Gaussianity, the model is estimated with a combination of Sequential and Markow Chain Monte Carlo methods that are introduced in the next section.

4 Model Estimation

Due to the skew Normal distribution of the forecasting errors as well as the non-linearity in the state equations, the SSV model forms a non-linear state space model. Therefore, the states and static model parameters cannot be estimated using Kalman Filtering and EM-Algorithms that usually provide feasible estimation methods. Yet, non-linear state space models can feasibly be estimated in a Bayesian setting using a combination of SMC-Methods such as Particle filters and MCMC-Algorithms such as the Metropolis-Hastings Sampler (Schön, Lindsten, Dahlin, Wagberg, Naesseth, Svensson, and Dai (2015)). In particular, Kim, Shephard, and Chib (1998) show that estimation of stochastic volatility models using a Particle Metropolis Hastings algorithm is straight forward. Given a non-linear, non-Gaussian state space system consisting of measurements y_t and latent states variables s_t that evolve according to the densities

$$y_t \sim p(y_t | s_t) \tag{7}$$

$$s_t \sim p(s_t | s_{t-1}). \tag{8}$$

the estimation of the model is divided in two parts: Posterior distributions of the static model parameters

$$\theta = (\delta_0, \delta_1, \gamma_{1,0}, \gamma_{1,1}, \gamma_{1,2}, \gamma_{2,0}, \gamma_{2,1}, \gamma_{2,2}, \sigma_{\nu,1}, \sigma_{\nu,2})$$

are obtained with a Metropolis Hastings sampler that generates draws from the posterior

$$p(\theta|y_{1:T}, s_{1:T}) = \frac{p(y_{1:T}|s_{1:T}, \theta)p(s_{1:T}|\theta)p(\theta)}{p(y_{1:T})} \quad (9)$$

by constructing a Markov Chain $\{\theta_i\}_{i=1}^N$ with stationary distribution equal to the posterior. Posterior distributions of the time varying model parameters $s_t = (\ln \sigma_t, \alpha_t)$ are given by the filtering distribution

$$p(s_t|y_{1:t}) = \frac{p(y_t|s_t)p(s_t|y_{1:t-1})}{\int p(y_t|s_t)p(s_t|y_{1:t-1})ds_t} \quad (10)$$

$$\text{with } p(s_t|y_{1:t-1}) = p(s_t|s_{t-1})p(s_{t-1}|y_{1:t-1}) \quad (11)$$

using a particle filter that sequentially approximates $p(s_t|y_{1:t})$ using importance sampling. Therefore, at each point in time t , M particles consisting of the tuples $\{s_t^i, W_{i=1}^M\}$ are drawn from a proposal $q(s_t|y_{1:t})$. In principle, the proposal can be chosen freely, yet a convenient choice generates draws based on the mixture density

$$q(s_t|y_{1:t}) = \sum_{i=1}^M W_{t-1,i} p(s_t|s_{t-1}^i) \quad \text{with} \quad \sum_{i=1}^M W_{t-1,i} = 1 \quad (12)$$

This choice of proposal yields the standard Bootstrap Particle filter that assigns weights $W_{t,i}$ given by

$$W_{t,i} = \frac{w_{t,i}}{\sum_{i=1}^M w_{t,i}} \quad \text{with} \quad w_{t,i} = p(y_t|s_{t,i}). \quad (13)$$

Hence, the particles are resampled proportional to the likelihood of the measurement y_t given the proposed states $s_{t,i}$ (Doucet et al. (2001)). Additionally, conditional on some choice of parameters θ the Particle filter generates an approximation of the likelihood function $\mathcal{L}(\theta)$ given by

$$\mathcal{L}(\theta) = p(y_{1:T}|s_{1:T}, \theta) \approx \prod_{j=1}^T \frac{1}{M} \sum_{i=1}^M W_{t,i}. \quad (14)$$

that can be embedded in the Metropolis Hastings Algorithm for the static model parameters. As shown by Andrieu, Doucet, and Holenstein (2010), these Particle MCMC-Algorithms remain valid such that the distribution of the chain $\{\theta_i\}_{i=1}^N$ converges to the exact posterior $p(\theta|y_{1:T}, s_{1:T})$, if the likelihood function is approximated using a particle filter.

Yet, it has often been pointed out that the proposal distribution given by (12) is not optimal since it ignores information about the states s_t contained in y_t (see for example Herbst and Schorfheide (2019)). As argued by Pitt and Shephard (1999), this leads to a poor approximation of the tails of the filtering distribution. Especially in case of outliers, most particles that are proposed based on $s_{t-1,i}$ will only have a low likelihood of y_t . This results in a poor approximation of the filtering density as only a few particles are resampled and consequently in a high variance of the weights $W_{t,i}$. This phenomenon is commonly referred to as weight degeneracy (Pitt and Shephard (1999)). Conversely, a more uniform the distribution of the weights, will lead to a smaller variance and a the better approximation of the filtering density. Therefore, the accuracy of the approximation at time t is gauged by the Inefficiency Ratio given as

$$Ineff_t = \frac{1}{M} \sum_{i=1}^M \left(\frac{w_{t,i}}{\frac{1}{M} \sum_{i=1}^M w_{t,i}} \right)^2 \quad (15)$$

where $w_{t,i}$ are the unnormalized weights.² A high inefficiency ratio indicates very uneven weights, while an inefficiency ratio close to 1 indicates evenly distributed weights.

Since risk analysis is most relevant to accurately capture extreme values, this paper uses the Tempered Particle Filter introduced by Herbst and Schorfheide (2019) as a more complex but also more accurate algorithm to improve the estimation of the states and to obtain a better approximation of the likelihood function for the Metropolis Hastings step. A comparison of the performance of the Bootstrap and the Tempered Particle Filter based on simulated data from the SSV model is given in Appendix A.3

4.1 Tempered Particle Filter

In theory, the optimal proposal distribution at time t is given by $p(s_t|y_t, s_{t-1})$, which in general is not available in closed form.³ Therefore, the Tempered Particle Filter proposed by Herbst and Schorfheide (2019) adjusts the proposal distribution to the observation y_t using annealed Importance Sampling first proposed by Neal (2001). After updating the particles $\{s_{t-1,i}, w_{t,i}\}_{i=1}^M$ from $t-1$ to t , the proposed states are evaluated on a sequence

²It can be shown that if one can draw particles from the optimal proposal density $p(s_t|y_t, s_{t-1})$ the weights become $w_{t,i} = \frac{1}{M} \forall i$ which gives $Ineff_t = 1$

³To be more precise, $p(s_t|y_t, s_{t-1})$ is only available in closed form if the distributions of the measurement and states are conjugate to each other.

of N_ϕ bridge distributions defined as the "tempered" likelihood function $p_n(y_t|s_{t,i})$ with an inflated measurement variance

$$\sigma^2/\phi_n \quad \text{with} \quad 0 < \phi_n < 1 \quad \text{and} \quad \lim_{n \rightarrow N_\phi} \phi_n = 1.$$

Thus, initially "flattening" the bell shape of the normal distribution results in more equal weights of the particles and decreases the Inefficiency Ratio towards a value r^* set by the researcher. Reducing the measurement error variance to its actual level while targeting a user-defined inefficiency ratio r^* the proposed states $s_{t,i}$ are then sequentially adapted to the more optimal proposal distribution $p(s_{t,i}, s_{t-1,i}|y_{1:t})$ ⁴. In contrary to Neal (2001), the tempering schedule is adaptive because at each step ϕ_n is chosen such that

$$Ineff(\phi_n) - r^* = 0. \tag{16}$$

Targeting a lower r^* will result in a better approximation of the latent states, but comes at the price of longer run-times and higher computational costs. The adaptive tempering procedure of Herbst and Schorfheide (2019) is designed with a focus on DSGE Models with non-linear dynamics but symmetric Gaussian measurement densities.

I modify the adaptive tempering schedule such that the asymmetry of $p_n(y_t|s_{t,i})$ is taken into account by additionally varying the asymmetry of the measurement density. Hence, starting from a symmetric and flat distribution, each tempering step brings the scale and shape of the measurement distribution closer to the final level. Following the reasoning in Herbst and Schorfheide (2019) as the skew Normal distribution converges to the symmetric Normal distribution for $\alpha \rightarrow 0$ this reduces the number of tempering iterations and prevents that particles close to the mode are overweighted. More formally, the unnormalized weights $\tilde{w}_{t,i}(\phi_0)$ are given by

$$p_0(y_t|s_{t,i}) = skew \mathcal{N}(y_t|\mu_t, \sigma_{t,i}/\phi_0, \phi_0\alpha_{t,i})$$

⁴This updating requires an additional MCMC-step that mutates states $s_{t,i}$ using a Transition Kernel that is invariant to the distribution $p_n(s_t|y_t, s_{t-1})$

Given expression (3) and Algorithm 2 in Herbst and Schorfheide (2019), this gives the following expression for the unnormalized weights for at the n^{th} tempering step

$$\tilde{w}_{t,i}(\phi_n) = \left(\frac{\phi_n}{\phi_{n-1}} \right)^2 \exp \left(\frac{-(\phi_n - \phi_{n-1})(y_t - \mu_t)}{2\sigma_{t,i}} \right)^2 \tilde{\Lambda}_{t,i}(\phi_n) \quad (17)$$

with

$$\tilde{\Lambda}_{t,i}(\phi_n) = \frac{\int_{-\infty}^{\alpha_{t,i}\phi_n^{2/3} \frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp \left(\frac{-t^2}{2} \right) dt}{\int_{-\infty}^{\alpha_{t,i}\phi_{n-1}^{2/3} \frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp \left(\frac{-t^2}{2} \right) dt} \quad (18)$$

Expression (17) and (18) show that the weights of the skew Normal distribution differ from a symmetric Normal density by a factor $\tilde{\Lambda}$ that is greater or smaller than 1 depending on the sign of $\alpha_{t,i}$ (see Appendix 1 for more details). Figure 2 and 3 illustrate the effect based on US GDP data. Figure 2 shows how the distribution of the proposed states $\ln(\sigma_{i,t}^2)$ is adjusted to the measurement in Q2 of 1978. The distribution is mutated from values of 2.04 (left) or 1.41 (right) to a final value of 3.80 to capture the higher risks of an unusually large upward spike in US GDP observed in the data. While exclusively tempering the scale of the measurement density takes a total of 10 iterations, also tempering the symmetry of the requires only 8 tempering steps. Given that the Particle Metropolis Hastings Algorithm requires $N \times M$ sweeps of the Tempered Particle filter with $T \times N_\phi$ tempering iterations per sweep, this adjustment decreases the runtime of the algorithm and enables researchers to improve the accuracy by targeting a lower Inefficiency Ratio r^* . Figure 3

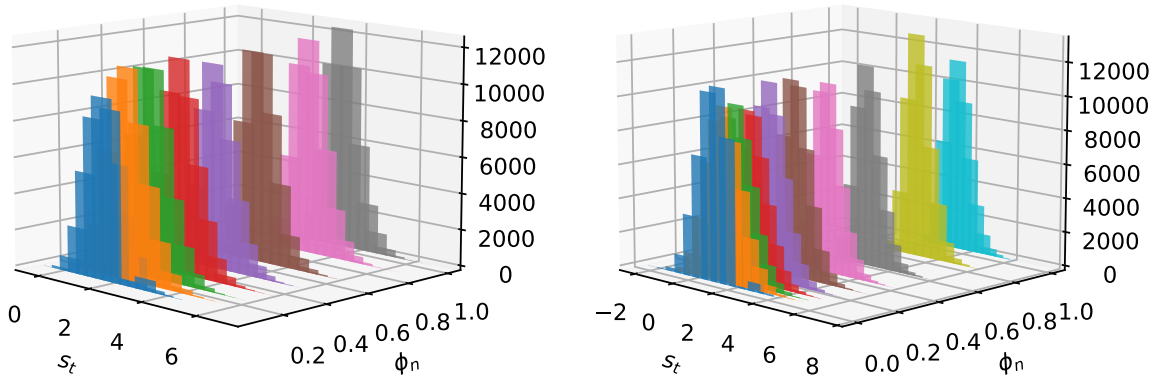


Figure 2: Tempering of the state distributions of $\ln(\sigma_{t,i}^2)$ in 1978 (Q2)

Note: In both cases, the mean of the distribution moves from 2.04 (left) or 1.41 (right) to about 3.80. However, while tempering only the scale of the distribution (right) results in 10 iterations, tempering the shape parameter $\alpha_{t,i}$ as well reduces the number of iterations to 8 only steps (left).

shows the total number of tempering steps for both tempering variants. In line with the intuition, the number of tempering steps increases during times of higher volatility with values of subsequent measurements being further apart. The plot shows that starting from a more symmetric density the filter requires fewer tempering steps with differences becoming more notable in times of higher volatility.⁵




Figure 3: Total number of tempering steps for both tempering variants

Note: In line with intuition tempering steps increase during times of higher volatility. The plot shows that additionally tempering the shape of the measurement density requires fewer tempering steps.

4.2 Data and Priors

The proposed model is estimated on the same data set as used by Adrian et al. (2019) with a mixture of uninformative and informative priors on the static parameters. On the one hand, priors on the parameters of the mean equation are centered at the OLS estimates of a linear regression of the variable $nfcit$ on gdp_{t+1} . Since the OLS estimator is equal to the ML-Estimator under the assumption of normally distributed errors, this reflects the prior believe that neither time-varying skewness nor volatility are present. On the other hand, priors on the parameters of the state equations are diffuse to be identified by the data Table 2 in A.2 gives a comprehensive overview of the prior specification of

⁵A simulation exercise based on 500 runs of both also yielded that including the symmetry of the measurement distribution in reduces the number of tempering steps by about 25%.

the static parameters as well as the data. The Tempered Particle Filter is run using $M = 40000$ particles with a targeted Inefficiency ratio $r^* = 1.2$ and 2 mutation steps in each tempering iteration. Draws for the static model parameters are generated using a standard Random Walk proposal with four chains ran in parallel on the HPC-Cluster at the Freie Universität. To increase the efficiency of the Metropolis Hastings algorithm the constrained model parameters $\delta_{1,2}$ and σ_{ν_i} are mapped to the real line using the following transformations

$$\delta_{1,2} = \tanh(\psi) \in [-1, 1] \quad (19)$$

$$\sigma_{\nu_i} = \exp(\xi_i) \in \mathbb{R}^+ \quad (20)$$

As described in Schön et al. (2015) this requires to correct the original priors of $\delta_{1,2}$ and σ_{ν_i} for the Jacobians of the inverse functions based on the change of variables rule.

$$\frac{d \tanh^{-1}(\psi)}{d\psi} = \frac{1}{1 - \psi^2} \quad \text{and} \quad \frac{d \log(\xi_i)}{d\xi_i} = \frac{1}{\xi_i}$$

to obtain samples from the transformed target distribution of

$$\tilde{\theta} = (\gamma_0, \gamma_1, \delta_{1,0}, \delta_{1,2}, \delta_{2,0}, \delta_{2,1}, \psi, \xi_1, \xi_2) \in \mathbb{R}.$$

The posterior distributions of the original model parameters can then be recovered using Equation (19) and (20). To further improve the mixing properties of the chains an initial estimate of $Var(\tilde{\theta}) = \Omega$ is obtained based on a pre-run of 5000 draws. The proposal variance is scaled to target an acceptance ratio between 20% and 30%.

5 Results

Table 1 presents the estimates for the static model parameters based on the Particle Metropolis Hastings approach described in the previous section. With an impact of about -0.7 on the mean (γ_0), 0.39 on the scale ($\delta_{1,1}$) and -0.42 on the shape ($\delta_{2,1}$) of the forecasting distribution of the one-period ahead US-GDP growth rate, the impact of national financial conditions is in line with the findings of Adrian et al. (2019). As financial conditions deteriorate, the expected growth rate decreases while the interquartile range as well

| Model Parameter | Mean | SD | q05 | q95 |
|------------------|--------|-------|--------|--------|
| γ_0 | 2.217 | 0.335 | 1.69 | 2.799 |
| γ_1 | -0.695 | 0.236 | -1.125 | -0.351 |
| $\delta_{1,0}$ | 1.295 | 0.385 | 0.784 | 2.06 |
| $\delta_{1,1}$ | 0.388 | 0.139 | 0.197 | 0.647 |
| $\delta_{1,2}$ | 0.292 | 0.198 | -0.106 | 0.556 |
| $\delta_{2,0}$ | 0.401 | 0.305 | 0.904 | 0.649 |
| $\delta_{2,1}$ | -0.429 | 0.226 | -0.81 | -0.038 |
| σ_{ν_1} | 0.451 | 0.119 | 0.298 | 0.688 |
| σ_{ν_2} | 0.533 | 0.182 | 0.319 | 0.877 |

Table 1: Posterior Means, Standard deviations (SD) and 90% credible sets of the static model parameters

Notes: The model was estimated using $N = 20000$ draws of the Particle Metropolis Hastings Algorithm. The first half of the sample was discarded as burn in. The model specification containing a lagged value of α_t in the state equation for the shape parameter was strongly rejected against a model specification without an AR(1) term based on a Bayes Ratio of 200.004. Marginal Data Densities were estimated using the modified harmonic mean estimator (Geweke (1999))

as downside risks to GDP-growth increase. Furthermore, all three coefficients are significantly different from zero based on the 90 % credible set constructed from the posterior draws. Additionally, Figure 4 shows the sample approximations and prior distribution of the parameters of the mean equation γ_0 and γ_1 as well as the parameters that capture the effect of the $nfcit$ variable on the shape and scale of the forecasting distribution. All posteriors are well-behaved, uni-modal and clearly centered away from zero. Additionally, the upper two panels show that the mode of the posterior distributions is different from the OLS estimates that the priors are centered on. The effect of the national financial conditions index is dampened compared to the estimated values under symmetric, normally distributed errors hinging to the fact that some variation is captured by the contemporaneous effect of national financial conditions on the skewness of the distribution. The same holds for the intercept, which is also smaller than its OLS counterpart. With a value of approximately -0.4 the absolute effect of $nfcit$ on the shape is only marginally larger than on the volatility. The estimates for the variances of the innovations of the state equations indicate that there also seems to be more uncertainty in the evolution of the asymmetry compared to the volatility of the forecasting distribution.

Figure 9 shows the filtered states given the mean estimates of the static model parameters. The sharp increase in volatility and downside risk in the 1980s as well as during the financial crises in 2009 is well captured by the evolution of the two latent states. These

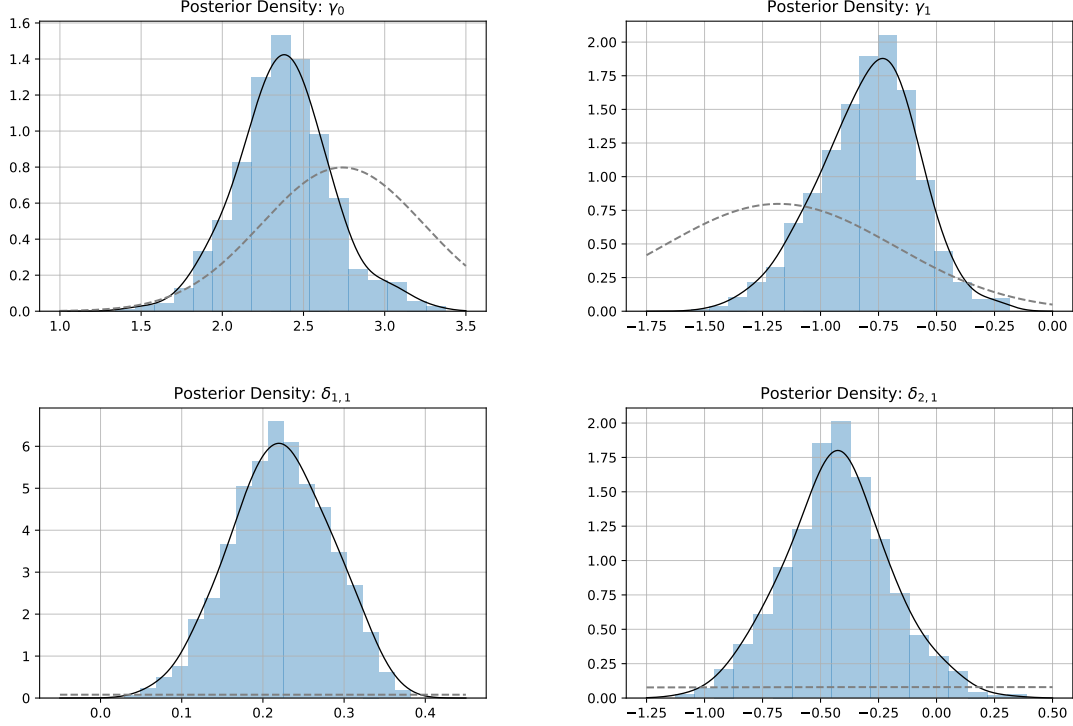


Figure 4: Posterior Distributions obtained using the Particle MCMC Algorithm

Posterior distributions for the parameters of the mean equation γ_0 and γ_1 as well as the parameters that capture the effect of the $nfcit_t$ variable on the scale ($\delta_{1,1}$) and shape ($\delta_{2,1}$). All posteriors are well-behaved, uni-modal and clearly away from zero with coefficients significantly different from zero based on the 90 % credible set constructed from the posterior draws. The grey dashed lines indicate the prior distributions.

findings confirm the results by Adrian et al. (2019) and show that the proposed model is able to replicate the stylized facts given in section 2. However, even though the distribution is more symmetric in normal times, the estimated state of α_t even exhibits positive skewness in times of moderation. This is in line the results of delle Monache et al. (2021) who find evidence for a cyclical behavior of the shape of the one-step ahead conditional forecasting distribution of US GDP growth with positive skewness in expansionary periods and negative skewness during recessions. Furthermore, the filtered states exhibit a strong inverse behavior with periods of high volatility coinciding with an increase in downside risk. This complementary behavior of the two states is further illustrated by ?? that shows the negative correlation between the two filtered series with an estimated correlation coefficient of -0.41. Eventually, Figure 7 shows the resulting forecasting distributions of the estimated model with the lower and upper 5% and 25% quantiles. The effect of the strong increase in the scale and shape parameters during the Great Recession as well as during




Figure 5: Filtered states obtained with the Tempered Particle Filter

The filter is tuned to target an Inefficiency Ratio of $r^* = 1.2$, 2 Mutation steps and $M = 40000$ particles. The posterior means in Table 1 are used for the static model parameters.

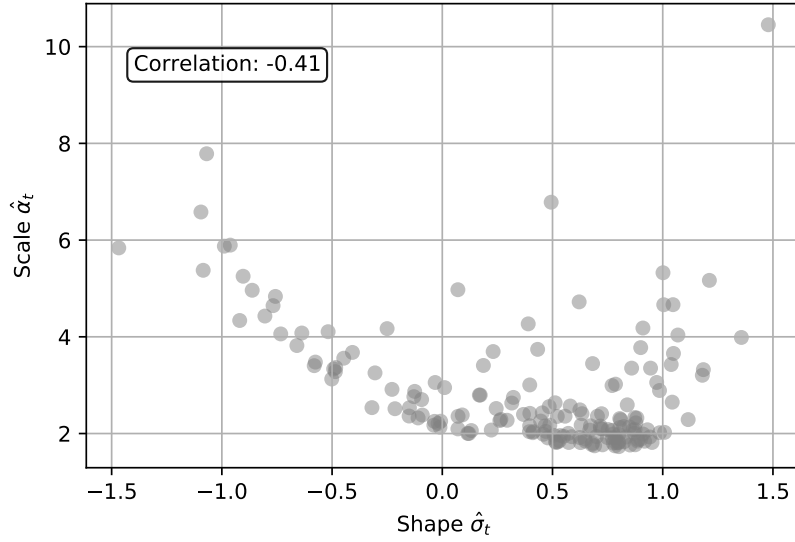


Figure 6: Scatterplot of the time-varying scale and shape parameter

Notes: The negative correlation of -0.41 is clearly visible. This implies that increases in volatility occurs with an increase in downside risks implied by left-skewed distributions.

the Oil Crises in the 1970s and 80s is clearly visible in the behavior of the lower quantiles. The distributions show a very similar behavior with stable upper quantiles while the lower quantiles vary significantly over time. Hence, the Skewed Stochastic Volatility Model is not only well suited to appropriately capture all the features described by Adrian et al. (2019), additionally providing statistical insights and parameter estimates that are in line with economic intuition as well as with the existing literature.




Figure 7: The one-step ahead forecasting distribution for one-quarter ahead GDP growth

Notes: Lower and upper 5% and 25% percent quantiles display the same characteristics as found by Adrian et al. (2019). While the upper quantiles remain relatively stable, the lower quantiles vary strongly over time indicating increased downside risk to GDP growth in times of financial distress.

6 Conclusion

This paper investigates the possibilities of Skewed Stochastic Volatility models as an alternative method to perform Growth at Risk Analysis as introduced by Adrian et al. (2019). To capture the features of the one-period ahead conditional forecast distributions, I propose a Skewed Stochastic Volatility model that allows researchers to estimate and conduct statistical inference on the estimated parameters to determine and to test the effects of financial conditions or other exogenous driving variables on the first three moments of the conditional forecast distribution of future GDP growth. The resulting state space model is non-linear with non-Gaussian errors and can be estimated with SMC and MCMC Methods. I obtain robust estimates of the model likelihood and the evolution of the latent states, using the Tempered Particle Filter introduced by Herbst and Schorfheide (2019). I modify the adaptive tempering schedule proposed by the authors to account for the asymmetry in the error distribution of the measurement equation. This reduces the number tempering steps to save computational time while achieving the same accuracy in the estimation of the states and likelihood function. Estimating the model based on US data yields conditional forecast densities that closely resemble the findings

by Adrian et al. (2019). Exploiting the advantages of the proposed model, I find that the estimated parameter values for the effect of financial conditions on the variance and skewness of the conditional distributions are significant and in line with the intuition of results found in the existing literature.

References

- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review* 109(4), 1263–89.
- Andrieu, C., A. Doucet, and R. Holenstein (2010). Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(3), 269–342.
- Azzalini, A. (2013). *The Skew-Normal and Related Families*. Institute of Mathematical Statistics Monographs. Cambridge University Press.
- Blum, A., J. Hopcroft, and R. Kannan (2020). *Foundations of Data Science*. Cambridge University Press.
- Carriero, A., T. E. Clark, and M. Marcellino (2020, January). Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions. Working Papers 202002R, Federal Reserve Bank of Cleveland.
- Creal, D., S. J. Koopman, and A. Lucas (2013). Generalized Autoregressive Score Models With Applications. *Journal of Applied Econometrics* 28(5), 777–795.
- de Santis, R. A. and W. van der Veken (2020). Forecasting macroeconomic risk in real time: Great and Covid-19 Recessions. Working Papers 2436, European Central Bank.
- delle Monache, D., A. de Polis, and I. Petrella (2021). Modeling and forecasting macroeconomic downside risk. Working Papers 1324, Banca d’Italia.
- Doucet, A., N. de Freitas, and N. J. Gordon (Eds.) (2001). *Sequential Monte Carlo Methods in Practice*. Statistics for Engineering and Information Science. Springer.
- Engle, R. and S. Manganelli (2004). Caviar: Conditional autoregressive value at risk by regression quantiles. *Journal of Business Economic Statistics* 22, 367–381.
- Engle, R. F. (2011). Long-Term Skewness and Systemic Risk. *Journal of Financial Econometrics* 9(3), 437–468.
- Flury, T. and N. Shephard (2011). Bayesian Inference based only on Simulated Likelihood: Particle Filter Analysis of Dynamic Economic Models. *Econometric Theory* 27(5), 933–956.
- Geweke, J. (1999). Using simulation methods for bayesian econometric models: inference, development, and communication. *Econometric Reviews* 18(1), 1–73.

- Hansen, B. (1994). Autoregressive conditional density estimation. *International Economic Review* 35(3), 705–30.
- Harvey, C. R. and A. Siddique (1999). Autoregressive conditional skewness. *The Journal of Financial and Quantitative Analysis* 34(4), 465–487.
- Herbst, E. and F. Schorfheide (2019). Tempered particle filtering. *Journal of Econometrics* 210(1), 26–44.
- Iseringhausen, M. (2021). A Time-varying Skewness Model for Growth-at-Risk. Working Papers 49, European Stability Mechanism.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies* 65(3), 361–393.
- Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica* 46(1), 33–50.
- López-Salido, D. and F. Loria (2020, February). Inflation at Risk. Finance and Economics Discussion Series 2020-013, Board of Governors of the Federal Reserve System (U.S.).
- Montes-Galdon, C. and E. Ortega (2022). Skewed SVARS: Tracking the structural sources of macroeconomic tail risks. *Advances in Econometrics* (forthcoming).
- Neal, R. M. (2001). Annealed importance sampling. *Statistics and Computing* 11(2), 125–139.
- Orlik, A. and L. Veldkamp (2014, August). Understanding Uncertainty Shocks and the Role of Black Swans. Working Paper 20445, National Bureau of Economic Research.
- Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* 94(446), 590–599.
- Prasad, A., S. Elekdag, P. Jeasakul, R. Lafarguette, A. Alter, A. X. Feng, C. Wang, and C. A. Gust (2019). Growth at Risk: Concept and Application in IMF Country Surveillance. Working Paper 036, International Monetary Fund.
- Schön, T. B., F. Lindsten, J. Dahlin, J. Wagberg, C. A. Naesseth, A. Svensson, and L. Dai (2015). Sequential monte carlo methods for system identification. *IFAC-PapersOnLine* 48(28), 775–786. 17th IFAC Symposium on System Identification SYSID 2015.

A Appendix

A.1 Inefficiency Ratio given a skew Normal

From Herbst and Schorfheide (2019), the weights $w_t^i(\phi_0)$ for $0 < \phi_0 \leq 1$ are given by the tempered likelihood function evaluated at the states $s_{t,i}$

$$p_0(y_t|s_{t,i}) = \frac{2\phi_0^{1/2}}{\sqrt{2\pi}\sigma_{t,i}} \exp\left(\frac{-\phi_0(y_t - \mu_t)^2}{2\sigma_{t,i}^2}\right) \int_{-\infty}^{\alpha_{t,i}\phi_0^{1/2}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt. \quad (21)$$

Given the annealed importance sampling method described in Neal (2001) and Algorithm 2 by in Herbst and Schorfheide (2019), the expression of the unnormalized weights $w_t^i(\phi_n)$ are defined as the ratio of the bridge distributions

$$w_{t,i}(\phi_n) = \frac{p_n(y_t|s_{t,i}^i)}{p_{n-1}(y_t|s_{t,i}^i)} \quad (22)$$

Using expression (3) for the density of the skew normal distribution yields

$$w_{t,i}(\phi_n) = \left(\frac{\phi_n}{\phi_{n-1}}\right)^{1/2} \exp\left(\frac{-(\phi_n - \phi_{n-1})(y_t - \mu_t)^2}{2\sigma_{t,i}^2}\right) \frac{\int_{-\infty}^{\alpha_{t,i}\phi_n^{1/2}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt}{\int_{-\infty}^{\alpha_{t,i}\phi_{n-1}^{1/2}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt} \quad (23)$$

Expression (17) shows that in comparison to normally distributed measurement errors, the weights of the skew normal errors are scaled by a factor

$$\Lambda_{t,i}(\phi_n) = \frac{\int_{-\infty}^{\alpha_{t,i}\phi_n^{1/2}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt}{\int_{-\infty}^{\alpha_{t,i}\phi_{n-1}^{1/2}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt} \quad (24)$$

Additionally tempering the symmetry of the skew normal distribution modifies this factor to

$$\tilde{\Lambda}_{t,i}(\phi_n) = \frac{\int_{-\infty}^{\alpha_{t,i}\phi_n^{2/3}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt}{\int_{-\infty}^{\alpha_{t,i}\phi_{n-1}^{2/3}\frac{(y_t - \mu_t)}{\sigma_{t,i}}} \exp\left(\frac{-t^2}{2}\right) dt} \begin{cases} < 1 \text{ if } \alpha_{t,i} < 0 \\ > 1 \text{ if } \alpha_{t,i} > 0 \end{cases} \quad \forall 0 > \phi_n > 1. \quad (25)$$

Inequality (25) holds since

$$\int_{-\infty}^x \exp\left(\frac{-t^2}{2}\right) dt$$

is monotonically increasing in x on a range from 0 to $\sqrt{2\pi}$ (for a proof see for example (Blum, Hopcroft, and Kannan (2020))) and the condition that $\phi_n > \phi_{n-1}$.

The tempering schedule is adaptive since in each iteration, solving the equation

$$\sum_{i=1}^M \left[\frac{\exp\left(\frac{-(\phi_n - \phi_{n-1})(y_t - \mu_t)}{2\sigma_{t,i}}\right)^2 \tilde{\Lambda}_{t,i}(\phi_n)}{\sum_{i=1}^M \exp\left(\frac{-(\phi_n - \phi_{n-1})(y_t - \mu_t)}{2\sigma_{t,i}}\right)^2 \tilde{\Lambda}_{t,i}(\phi_n)} \right]^2 - r^* = 0. \quad (26)$$

provides a new value for ϕ_n .⁶

A.2 US Data and Priors




Figure 8: US GDP and National Financial Conditions Index.

The sampling frequency for both series is quarterly and the sample ranges from 1973 Q1 to 2016 Q2.

⁶Using (21), it can be shown that for a model with stochastic volatility, the lower bound on r^* for $\phi_0 \rightarrow 0$ is given as

$$\lim_{\phi_0 \rightarrow 0} Ineff(\phi_0) = \frac{\frac{1}{M} \sum_{i=0}^M \left(\frac{1}{\sigma_{i,t}}\right)^2}{\left(\frac{1}{M} \sum_{i=0}^M \frac{1}{\sigma_{i,t}}\right)^2} > 1$$

| Model Parameter | Distribution | Param 1 | Param 2 |
|------------------|--------------|---------|---------|
| γ_0 | N | 2.1 | 0.5 |
| γ_1 | N | -1.18 | 0.5 |
| $\delta_{1,0}$ | N | 0 | 5 |
| $\delta_{1,1}$ | N | 0 | 5 |
| $\delta_{1,2}$ | U | -1 | 1 |
| $\delta_{2,0}$ | N | 0 | 5 |
| $\delta_{2,1}$ | N | 0 | 5 |
| σ_{ν_1} | IG | 3 | 6 |
| σ_{ν_2} | IG | 3 | 6 |

Table 2: Priors for the static model parameters in the Metropolis Hastings Algorithm.

N denotes normal priors with Param 1 and Param 2 giving mean and variances, U denotes uniform priors with Param 1 and Param 2 for the upper and lower bound. IG denotes the inverse Gamma distribution with Param 1 and Param 2 for α and β .

A.3 Bootstrap Particle Filter vs. Tempered Particle Filter




Figure 9: Mean Squared Errors of the filtered states

Mean Squared Errors are calculated based on 500 simulations obtained with the Bootstrap Particle Filter and the Tempered Particle Filter. Tuning parameters of the Tempered Particle Filter were set to target an Inefficiency Ratio of $r^* = 1.2$, 2 Mutation steps and $M = 40000$ particles. The superior performance of the Tempered Particle Filter is clear from the mean and standard errors of the distributions.

Diskussionsbeiträge - Fachbereich Wirtschaftswissenschaft - Freie Universität Berlin
Discussion Paper - School of Business & Economics - Freie Universität Berlin

2022 erschienen:

2022/1 AHRENS, Steffen; Ciril BOSCH-ROSA und Thomas MEISSNER:
Intertemporal Consumption and Debt Aversion: A Replication and Extension
Economics