WORKING PAPER 8/2024 (ECONOMICS)

VAR Models with Fat Tails and Dynamic Asymmetry

Tamás Kiss, Stepan Mazur, Hoang Nguyen and Pär Österholm

ISSN 1403-0586

Örebro University School of Business
SE-701 82 Örebro, Sweden

VAR Models with Fat Tails and Dynamic Asymmetry

Tamás Kiss[(a)], Stepan Mazur[(a)], Hoang Nguyen[(b)], Pär Österholm [(a,c)]

[(a)] School of Business, Örebro University, Sweden

[(b)] Department of Management and Engineering, Linköping University, Linköping, Sweden

[(c)] National Institute of Economic Research, Stockholm, Sweden

October 9, 2024

Abstract

In this paper, we extend the standard Gaussian stochastic-volatility Bayesian VAR by employing the generalized hyperbolic skew Student's t distribution for the innovations. Allowing the skewness parameter to vary over time, our specification permits flexible modelling of innovations in terms of both fat tails and – potentially dynamic – asymmetry. In an empirical application using US data on industrial production, consumer prices and economic policy uncertainty, we find support – although to a moderate extent – for time-varying skewness. In addition, we find that shocks to economic policy uncertainty have a negative effect on both industrial production growth and CPI inflation.

JEL Classification: C11, C32, C52, E44, E47, G17

Keywords: Bayesian VAR; Generalized hyperbolic skew Students's t distribution; Stochastic volatility; Economic policy uncertainty

1 Introduction

Many macroeconomic time series seem to be associated with properties that are inconsistent with a linear model which has error terms that are Gaussian distributed (and homoscedastic). One such property is that large shocks appear to hit the economy more frequently than what would be expected if the shocks were drawn from a Gaussian distribution. For example, it can be argued that the occurrence of the Global Financial Crisis, the crisis related to the Covid-19 pandemic and the inflation surge of 2022-2023 in a fairly short period of time warrants models with distributions for the error terms that have fat tails. A second property is that the unconditional distribution of the variable can be characterised by skewness. This feature might be driven by disturbances whose distributions have skewness; this could potentially be time-varying, for instance, being different in expansions and contractions. Empirical work documenting and addressing fat tails and skewness include, for example, Acemoglu and Scott (1997), Ni and Sun (2005), Fagiolo et al. (2008), Ascari et al. (2015), Cross and Poon (2016), Liu (2019), Carriero et al. (2020), and Delle Monache et al. (2024).

In this paper, we address these two issues and further the literature by specifying a Bayesian VAR model with stochastic volatility where the disturbances are drawn from a generalized hyperbolic skew Student's t distribution with time-varying skewness. This specification permits flexible modelling of innovations in terms of fat tails and – potentially dynamic – asymmetry. Our model nests specifications with Gaussian and fat-tailed innovations, as well as the specification where asymmetry is constant over time; the last specification was first defined in Karlsson et al. (2023). In setting up the model, we also describe the appropriate priors and the general Markov Chain Monte Carlo procedure used to draw inference.

We illustrate the framework with an application to US data in which we model industrial production, consumer prices and the economic policy uncertainty (EPU) index of Baker et al. (2016). The relation of the EPU index to the macroeconomy has been analysed widely over the last decade and studies tend to show that an increasing EPU index has a dampening effect on the real economy; this includes holding back both investment (e.g. Gulen and Ion, 2016) and consumption (e.g. Nam et al., 2021), an increasing unemployment rate (e.g. Caggiano et al., 2017) and a predictive ability for recessions (e.g. Karnizova and Li, 2014 and Balcilar et al., 2016).

The main contribution of this paper is methodological. The most closely related research in the previous literature includes Karlsson et al. (2023) who introduced a general class of VAR model with generalized hyperbolic skew Student's t innovations and stochastic volatility. Delle Monache et al. (2024) provided evidence of time-varying conditional asymmetry of US GDP using a time-varying location, scale, and asymmetry model. Using the same modelling framework, De Polis et al. (2024) documented largely downside risk in inflation between the mid-1990s and the early 2020s. Iseringhausen (2024) studied the growth-at-risk of advanced economies using a univariate time-varying volatility and skewness model. We contribute to the literature by extending the Karlsson et al. (2023) model allowing the skewness parameters of the generalized hyperbolic skew Student's t innovations to be time-varying. Unlike the observation-driven model of Delle Monache et al. (2024) and De Polis et al. (2024), our proposed model is parameter-driven where the time-varying skewness evolves over time as a stochastic process. Furthermore, we study dynamic skewness of a multivariate distribution, where the skewness parameters follow independent random-walk processes. Therefore, our framework can be seen as an extension of Tsiotas (2012) and Iseringhausen (2020). It also differs from Iseringhausen (2024) who has observable variables in the skewness equation. Our model allows us to account for the interaction between the variables included, as well as innovations that potentially feature dynamic skewness.

In addition to the methodological contribution, we also make a contribution to the empirical literature studying the relationship between the EPU index and the macroeconomy. Our results suggest that there is time-variation in the asymmetry of all three variables, with the evidence being strongest for the EPU index. Furthermore, we confirm the finding of Leduc and Liu (2016) that uncertainty shocks resemble aggregate demand shocks, where an increase in uncertainty is associated with decreasing real activity and lower inflation.

The rest of this paper is organised as follows. We present the methodological framework – including how inference is conducted – in Section 2. The empirical illustration of the framework is provided in Section 3. Finally, Section 4 concludes.

2 Econometric Framework

In this section, we first present a VAR model with a Gaussian distribution and stochastic volatility in the innovations.[1] We then generalize the framework so that we allow for fat tails, dynamic asymmetry and stochastic volatility; this is done by employing a general class of generalized hyperbolic skew Student's t distribution for the innovations. This class of VAR model is an extension of the models developed by Karlsson et al. (2023) and allows us to capture the time-varying effects of disturbances to the conditional mean and the conditional variance of the variables.

2.1 A VAR Model with Gaussian Innovations and Stochastic Volatility

We assume that a VAR model with Gaussian innovations and stochastic volatility (Gaussian.SV) is given by

$$\mathbf{y}_t = \mathbf{c} + \mathbf{B}_1\mathbf{y}_{t-1} + \ldots + \mathbf{B}_p\mathbf{y}_{t-p} + \mathbf{A}^{-1}\mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t, \qquad t = 1, \ldots, T, \tag{1}$$

where $\mathbf{y}_t$ is a k-dimensional vector of time series of interest; $\mathbf{c}$ is a k-dimensional vector of constants; $\mathbf{B}_j$ is a $k \times k$ matrix of regression coefficients with $j = 1, \ldots, p$; $\mathbf{A}$ is a $k \times k$ lower triangular matrix with ones on the diagonal that describes the contemporaneous interaction of the endogenous variables; $\mathbf{H}_t = diag(h_{1t}, \ldots, h_{kt})$ is a $k \times k$ diagonal matrix that captures the heteroskedasticity; and $\boldsymbol{\epsilon}_t$ is a k-dimensional vector that follows a multivariate Gaussian distribution with zero mean vector and identity covariance matrix, that is, $\boldsymbol{\epsilon}_t \sim \mathcal{N}_k(\mathbf{0}, \mathbf{I})$. We also assume that $\boldsymbol{\epsilon}_t$ is independent of $\mathbf{H}_t$. Moreover, we assume that the elements of $\mathbf{H}_t = diag(h_{1t}, \ldots, h_{kt})$ follow a random walk process such that

$$\log h_{it} = \log h_{it-1} + \sigma_{h,i}\eta_{it}, \qquad i = 1, \ldots, k, \tag{2}$$

where $\eta_{it} \sim \mathcal{N}(0, 1)$. Let us note that the VAR model with homoskedastic Gaussian innovations can be obtained by restricting the volatility to be constant, i.e. $\boldsymbol{\sigma}_h^2 = (\sigma_{h,1}^2, \ldots, \sigma_{h,k}^2)' = \mathbf{0}$ and $\log h_{it} = \log h_{i0}$ for all $i = 1, \ldots, k$ and $t = 1, \ldots, T$.

[1] We take stochastic volatility as a given feature of the model. There is ample evidence in the literature that many macroeconomic time series benefit from this modelling choice; see, for example, Cogley and Sargent (2005), Sims and Zha (2006), Clark (2011), Carriero et al. (2015), Chan (2017), Akram and Mumtaz (2019), Koop and Korobilis (2019), Karlsson and Österholm (2020a,b, 2023) and Kiss et al. (2023). This modelling choice is also supported in our empirical analysis. As can be seen in Figure 2, estimated volatilities vary substantially over time.

For the sake of notational simplicity, the Gaussian.SV model in (1) can be rewritten as follows

$$\mathbf{y}_t = \mathbf{B}\mathbf{x}_t + \mathbf{u}_t, \tag{3}$$

where $\mathbf{B} = (\mathbf{c}, \mathbf{B}_1, \ldots, \mathbf{B}_p)$ is a $k \times (1+kp)$ matrix, $\mathbf{x}_t = (1, \mathbf{y}'_{t-1}, \ldots, \mathbf{y}'_{t-p})'$ is a $(1+kp)$-dimensional vector, and $\mathbf{u}_t = \mathbf{A}^{-1}\mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t$ is a k-dimensional vector of heteroskedastic innovations associated with the VAR equations. In the next subsection, we extend the innovation term $\mathbf{u}_t$ to a class of generalized hyperbolic skew Student's t distribution with stochastic volatility and dynamic asymmetry.

2.2 A VAR Model with Fat Tails, Dynamic Asymmetry and Stochastic Volatility in the Innovations

We now propose a class of generalized hyperbolic skew Student's t VAR models by assuming that the vector of innovations $\mathbf{u}_t$ in (1) follows a Gaussian variance-mean mixture and is expressed as

$$\mathbf{u}_t = (w_t - \overline{w})\boldsymbol{\gamma}_t + \sqrt{w_t}\mathbf{A}^{-1}\mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t, \tag{4}$$

where the mixing variable w_t follows an inverse gamma distribution with the same shape and rate parameters equal to $\nu/2$, that is, $w_t \sim \mathcal{IG}(\nu/2, \nu/2)$; $\overline{w} = E[w_t] = \nu/(\nu-2)$ is used for centring the innovations vector so that its mean is equal to zero; $\boldsymbol{\gamma}_t = (\gamma_{1t}, \ldots, \gamma_{kt})'$ is the k-dimensional vector of the time-varying skewness parameters which are assumed to follow independent random walk processes such that

$$\gamma_{it} = \gamma_{it-1} + \sigma_{\gamma,i}\eta_{it}, \qquad i = 1, \ldots, k, \tag{5}$$

where $\eta_{it} \sim \mathcal{N}(0,1)$. The matrices $\mathbf{A}$ and $\mathbf{H}_t$, as well as the vector $\boldsymbol{\epsilon}_t$ in (4), are the same as in (1). Moreover, w_t, $\boldsymbol{\gamma}_t$, $\mathbf{H}_t$ and ϵ_t are mutually independent. In what follows, the VAR model with the innovation term defined as in (4) is called *dynSkew-T.SV*.

In general, the dynSkew-T.SV VAR model imposes the fat tails and time-varying asymmetry of the innovations directly in each VAR equation. Even though there is only one parameter to control for the fat tails of all variables, it is still sufficiently flexible to capture the joint dependence of the innovations. Regarding the properties of the model, it can be noted that the following holds:

$$\mathbf{u}_t|w_t, \boldsymbol{\gamma}_t, \mathbf{H}_t \sim \mathcal{N}_k \left(\boldsymbol{\mu}_t = (w_t - \overline{w})\boldsymbol{\gamma}_t,\ \boldsymbol{\Sigma}_t = w_t \mathbf{A}^{-1}\mathbf{H}_t \mathbf{A}^{-1'} \right). \tag{6}$$

It follows that the conditional distribution of $\mathbf{u}_t$ (given $\boldsymbol{\gamma}_t$ and $\mathbf{H}_t$) follows a multivariate generalized hyperbolic skew Student's t distribution (see McNeil et al., 2015). Setting $\boldsymbol{\sigma}^2_\gamma = (\sigma^2_{\gamma,1}, \ldots, \sigma^2_{\gamma,k})' = \mathbf{0}$ and $\gamma_{it} = \gamma_{i0}$ for all $i = 1, \ldots, k$ and $t = 1, \ldots, T$, the dynSkew-T.SV VAR model reduces to the VAR model with generalized hyperbolic skew Student's t innovations and stochastic volatility (Skew-T.SV). Additionally, assuming $\gamma_{it} = \gamma_{i0} = 0$, we arrive to the VAR model with generalized hyperbolic Student's t innovations and stochastic volatility (T.SV). Moreover, considering $\nu \to \infty$, we obtain the VAR model with Gaussian innovations and stochastic volatility (G.SV).

From an economic point of view, the mixing variable w_t is identified from the sudden changes of the conditional volatility of variables while the stochastic volatility term $\exp(h_{it}/2)$ captures smooth changes in the volatility (Chiu et al., 2017). The data will determine whether movements that are large in magnitude come from the mixing variable or a stochastic volatility shift. Is is also worth pointing out that the mixing variable itself is able to generate asymmetry in the unconditional distribution of the variables, even though innovations are drawn from a symmetric distribution. This is due to the fact that the mixing variable shifts the conditional mean of each variable. Since the uncertainty rises during recessions, these effects are different during recessions and expansions, generating asymmetry in the unconditional distribution of the variables (as described in Carriero et al., 2018). Our specification incorporates this feature but also directly allows for skewness in the conditional distribution of the innovations; hence, even the conditional distribution of the variables is allowed to be asymmetric.

2.3 Bayesian Inference

This section discusses the prior distributions for the parameters in the dynSkew-T.SV models and the general MCMC inference scheme.

2.3.1 Prior Distribution

We denote $\boldsymbol{\theta} = \{vec(\mathbf{B})^{'}, \mathbf{a}^{'}, \boldsymbol{\gamma}_0^{'}, \nu, \mathbf{h}_0^{'}, \boldsymbol{\sigma}_h^{2'}, \boldsymbol{\sigma}_\gamma^{2'}, \boldsymbol{\gamma}_{1:T}^{'}, \boldsymbol{w}_{1:T}^{'}, \mathbf{h}_{1:T}^{'}\}^{'}$ as the set of the dynSkew-T.SV model parameters. We set $\mathbf{a}$ as a vector of the lower triangular elements of the matrix $\mathbf{A}$ such that $\mathbf{a} = (a_{2,1}, a_{3,1}, a_{3,2}, \ldots, a_{k,k-1})^{'}$. Also let $\mathbf{h}_0$ and γ_0 be the vectors of initial values of the stochastic volatility and the dynamic skewness, respectively. Following the prior assumption in Koop and Korobilis (2010) and Karlsson et al. (2023), we assume that the regression coefficients $vec(\mathbf{B})$ follow a Minnesota-type prior,that is $vec(\mathbf{B}) \sim \mathcal{N}_{k(1+kp)}(\mathbf{b}_0, \mathbf{V}_{\mathbf{b}_0})$ such that the overall shrinkage is $l_1 = 0.2$ and the cross-variable shrinkage is $l_2 = 0.5$. For other parameters, we consider vague but proper prior distributions. For example, the prior of $\mathbf{a}$ is also Gaussian, $a_{ij} \sim \mathcal{N}(0, 10I_{0.5k(k-1)})$, which implies a weak assumption of no interaction among endogenous variables. The prior for the degrees of freedom parameter is $\nu \sim \mathcal{G}(2, 0.1)$ such that the prior mean of the degrees of freedom is 20. The initial skewness has zero prior mean and unit variance, while the initial log volatility is obtained from the OLS estimation of the AR(p) model for each time series. In the case of the VAR model without stochastic volatility, we consider $\log h_{i0} \sim \mathcal{N}\left(\log \hat{\Sigma}_{i,OLS}, 4\right)$, where $\hat{\Sigma}_{i,OLS}$ are variance estimators of AR(p) models using the ordinary least square method; see Clark and Ravazzolo (2015). The prior for the variance of a shock to the volatility and the time-varying skewness are $\sigma^2_{\gamma,i} \sim \mathcal{G}(\frac{1}{2}, \frac{1}{2V_\sigma})$ which is equivalent to $\pm\sqrt{\sigma^2_{\gamma,i}} \sim \mathcal{N}(0, V_\sigma)$; see Kastner and Frühwirth-Schnatter (2014). We choose $V_\sigma = 1$, noting that this prior is less influential compared to the conjugated inverse gamma prior especially in the case of a small true value. Finally, the prior for the mixing variable w_t is based on the model assumption $w_t|\nu \sim \mathcal{IG}(\nu/2, \nu/2)$.

2.3.2 Estimation Procedure

The Bayesian inference using a Gibbs sampler is extended to make inference on model parameters. To simplify the notation, let $\mathbf{\Psi}$ be a set of conditioning parameters except the one that we sample from.

1. To sample $\pi(\mathbf{b}|\mathbf{\Psi})$ where $\mathbf{b} = vec(\mathbf{B})$, we consider Equation (3) as

$$\begin{aligned}
\mathbf{y}_t - (w_t - \overline{w})\boldsymbol{\gamma}_t &= \mathbf{B}\mathbf{x}_t + w_t^{1/2}\mathbf{A}^{-1}\mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t, \\
w_t^{-1/2}\mathbf{A}\left(\mathbf{y}_t - (w_t - \overline{w})\boldsymbol{\gamma}_t\right) &= w_t^{-1/2}\mathbf{A}\mathbf{B}\mathbf{x}_t + \mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t \\
&= w_t^{-1/2}\mathbf{x}_t' \otimes \mathbf{A}\,\mathrm{vec}(\mathbf{B}) + \mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t, \\
\widetilde{\mathbf{y}}_t &= \widetilde{\mathbf{X}}_t\mathbf{b} + \mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t,
\end{aligned}$$

   where $\widetilde{\mathbf{y}}_t = w_t^{-1/2}\mathbf{A}\left(\mathbf{y}_t - (w_t - \overline{w})\boldsymbol{\gamma}_t\right)$ and $\widetilde{\mathbf{X}}_t = w_t^{-1/2}\mathbf{x}_t' \otimes \mathbf{A}$. Hence, the conditional posterior of $\mathbf{b}$ is a conjugate normal distribution that

$$\pi(\mathbf{b}|\mathbf{\Psi}) \sim \mathcal{N}(\mathbf{b}^*, \mathbf{V}_{\mathbf{b}}^*),$$

   where

$$\begin{aligned}
\mathbf{V}_{\mathbf{b}}^{*-1} =& \mathbf{V}_{\mathbf{b}_0}^{-1} + \sum_{t=1}^{T} \widetilde{\mathbf{X}}_t'\mathbf{H}_t^{-1}\widetilde{\mathbf{X}}_t, \\
\mathbf{b}^* =& \mathbf{V}_{\mathbf{b}}^* \left[\mathbf{V}_{\mathbf{b}_0}^{-1}\mathbf{b}_0 + \sum_{t=1}^{T} \widetilde{\mathbf{X}}_t'\mathbf{H}_t^{-1}\widetilde{\mathbf{y}}_t\right].
\end{aligned}$$

2. To sample $\pi(\boldsymbol{\gamma}_{0:T}|\mathbf{\Psi})$, we consider the time-varying coefficient γ_t on the representation of the dynSkew-T.SV model,

$$\begin{aligned}
\mathbf{y}_t - \mathbf{B}\mathbf{x}_t &= (w_t - \overline{w})\boldsymbol{\gamma}_t + w_t^{1/2}\mathbf{A}^{-1}\mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t, \\
\boldsymbol{\gamma}_t &= \boldsymbol{\gamma}_{t-1} + \mathrm{diag}(\boldsymbol{\sigma}_\gamma)\boldsymbol{\varsigma}_t.
\end{aligned}$$

   We apply the forward filter backward smoothing algorithm as in Carter and Kohn (1994) to filter out the last state variable of the time-varying skewness. Then, $\boldsymbol{\gamma}_t$ is sampled from its smoothing Gaussian distribution for $t = T, \ldots, 0$; see Primiceri (2005). As σ_γ^2 has a non-conjugate prior distribution, we sample $\pi(\boldsymbol{\sigma}_\gamma^2|\mathbf{\Psi})$ using an independent Metropolis-Hastings step; see Kastner and Frühwirth-Schnatter (2014).

3. To sample $\pi(\mathbf{a}|\mathbf{\Psi})$, we follow Cogley and Sargent (2005) and rewrite Equation (3) as a system

of linear equations with equation,

$$\mathbf{A}\widetilde{\mathbf{u}}_t = \mathbf{H}_t^{1/2}\boldsymbol{\epsilon}_t,$$

where $\tilde{\mathbf{u}}_t = w_t^{-1/2}(\mathbf{y}_t - \mathbf{B}\mathbf{x}_t - (w_t - \overline{w})\boldsymbol{\gamma}_t)$. Then, the conditional posterior of elements in $\mathbf{a}$ can be drawn equation by equation using the conditionally Gaussian posterior distribution; see Cogley and Sargent (2005).

4. To sample $\pi(\mathbf{H}_{1:T}|\mathbf{\Psi})$, we follow Kim et al. (1998), Primiceri (2005) and Del Negro and Primiceri (2015). Let $\tilde{\tilde{\mathbf{u}}}_t = \mathbf{A}\tilde{\mathbf{u}}_t$, for each series $i = 1, \ldots, k$, we have that $\log \tilde{\tilde{u}}_{it}^2 = \log h_{it} + \log \epsilon_t^2$. We follow Kim et al. (1998) who approximate the distribution of $\log(\epsilon_t^2)$ using a mixture of 7 normal components. Then, we sample $\log h_t$ from its smoothing Gaussian distribution using the forward filter backward smoothing algorithm in Carter and Kohn (1994). Finally, we sample $\pi(\boldsymbol{\sigma}_h^2|\mathbf{\Psi})$ from an independent Metropolis-Hastings; see Kastner and Frühwirth-Schnatter (2014).

5. To sample $\pi(\nu_i|\mathbf{\Psi}) \propto \mathcal{G}(\nu_i; 2, 0.1)\prod_{t=1}^{T}\mathcal{IG}\left(w_t; \frac{\nu_i}{2}, \frac{\nu_i}{2}\right)$ for $i = 1, \ldots, k$, we use an adaptive random walk Metropolis-Hastings algorithm to accept/reject the draw $\nu_i^{(*)} = \nu_i + \eta_i \exp(c_i)$, where $\eta_i \sim \mathcal{N}(0, 1)$ and the adaptive variance c_i is adjusted automatically such that the acceptance rate is around 0.25 (Roberts and Rosenthal, 2009).

6. To sample $\pi(w_t|\mathbf{\Psi})$ for $t = 1, \ldots, T$, we write

$$\mathbf{u}_t = \mathbf{y}_t - \mathbf{B}\mathbf{x}_t + \bar{w}\boldsymbol{\gamma}_t = w_t\boldsymbol{\gamma}_t + w_t^{1/2}\mathbf{A}^{-1}\mathbf{H}_t^{1/2}\epsilon_t.$$

The full conditional posterior is thus generalized inverse Gaussian following Hörmann and Leydold (2014), $GIG(\lambda, \psi, \chi)$ with $\lambda = -(\nu + k)/2$, $\chi = q_t^2 + \nu$ and $\psi = p_t^2$ for $q_t^2 =$

$\mathbf{u}_t'\mathbf{A}'\mathbf{H}_t^{-1}\mathbf{A}\mathbf{u}_t$ and $p_t^2 = \boldsymbol{\gamma}_t'\mathbf{A}'\mathbf{H}_t^{-1}\mathbf{A}\boldsymbol{\gamma}_t$ according to

$$\begin{aligned}
\pi(w_t|\boldsymbol{\Psi}) &\propto w_t^{-k/2}\exp\left(-\frac{1}{2}(\mathbf{y}_t - \mathbf{B}\mathbf{x}_t - (w_t-\bar{w})\boldsymbol{\gamma}_t)'\boldsymbol{\Sigma}_t^{-1}(\mathbf{y}_t - \mathbf{B}\mathbf{x}_t - (w_t-\bar{w})\boldsymbol{\gamma}_t)\right)\mathcal{IG}\left(w_t; \frac{\nu_i}{2}, \frac{\nu_i}{2}\right) \\
&\propto w_t^{-k/2}\exp\left(-\frac{1}{2}(\mathbf{u}_t - w_t\boldsymbol{\gamma}_t)'(w_t^{-1}\mathbf{A}'\mathbf{H}_t^{-1}\mathbf{A})(\mathbf{u}_t - w_t\boldsymbol{\gamma}_t)\right)\mathcal{IG}\left(w_t; \frac{\nu_i}{2}, \frac{\nu_i}{2}\right) \\
&\propto w_t^{-(k/2+\nu/2+1)}\exp\left(-\frac{1}{2}\frac{\mathbf{u}_t'\mathbf{A}'\mathbf{H}_t^{-1}\mathbf{A}\mathbf{u}_t + \nu}{w_t} - \frac{1}{2}w_t\boldsymbol{\gamma}_t'\mathbf{A}'\mathbf{H}_t^{-1}\mathbf{A}\boldsymbol{\gamma}_t\right).
\end{aligned}$$

3 Empirical Illustration

We next estimate models using monthly US data on industrial production, consumer prices and the the economic policy uncertainty (EPU) index of Baker et al. (2016). Analysis is conducted both within- and out-of-sample.

The data – which range from 1985M1 to 2024M3 – were sourced from the FRED-MD database of Federal Reserve Bank of St. Louis.[2] We calculate the year-on-year growth rate of industrial production and year-on-year inflation using the twelve-month log-difference of the industrial production index and the consumer price index. Denoting industrial production at time t by IP_t, the former is defined as $i_t = 100[\log(IP_t) - \log(IP_{t-12})]$; denoting the consumer price index at time t by CPI_t, the latter is given as $p_t = 100[\log(CPI_t) - \log(CPI_{t-12})]$. For the numerical stability of the estimation procedure, we divide the EPU by 10; we accordingly define $e_t = EPU_t/10$. The data series are shown in the Figure 1.

From a methodological point of view, our primary interest is in the skewness estimates. But it is also relevant to see how our empirical findings relate to earlier research that has studied the relation of the EPU index to the macroeconomy. As was pointed out in the introduction, an increasing EPU index tends to be associated with a slowing real economy, where a weaker development of investment and consumption is in line with the more general notion that uncertainty should hold back the real economy; see, for example, Bernanke (1983) and Bloom (2009). Previous findings regarding the real economy are also in line with the claim of Leduc and Liu (2016) who argue that uncertainty shocks resemble aggregate demand shocks. However, Jones and Olson (2013) find that the relation between the EPU index and inflation appears ambiguous. Adding more data and new

[2] See McCracken and Ng (2016) for details.

methodology to study this issue is hence relevant.

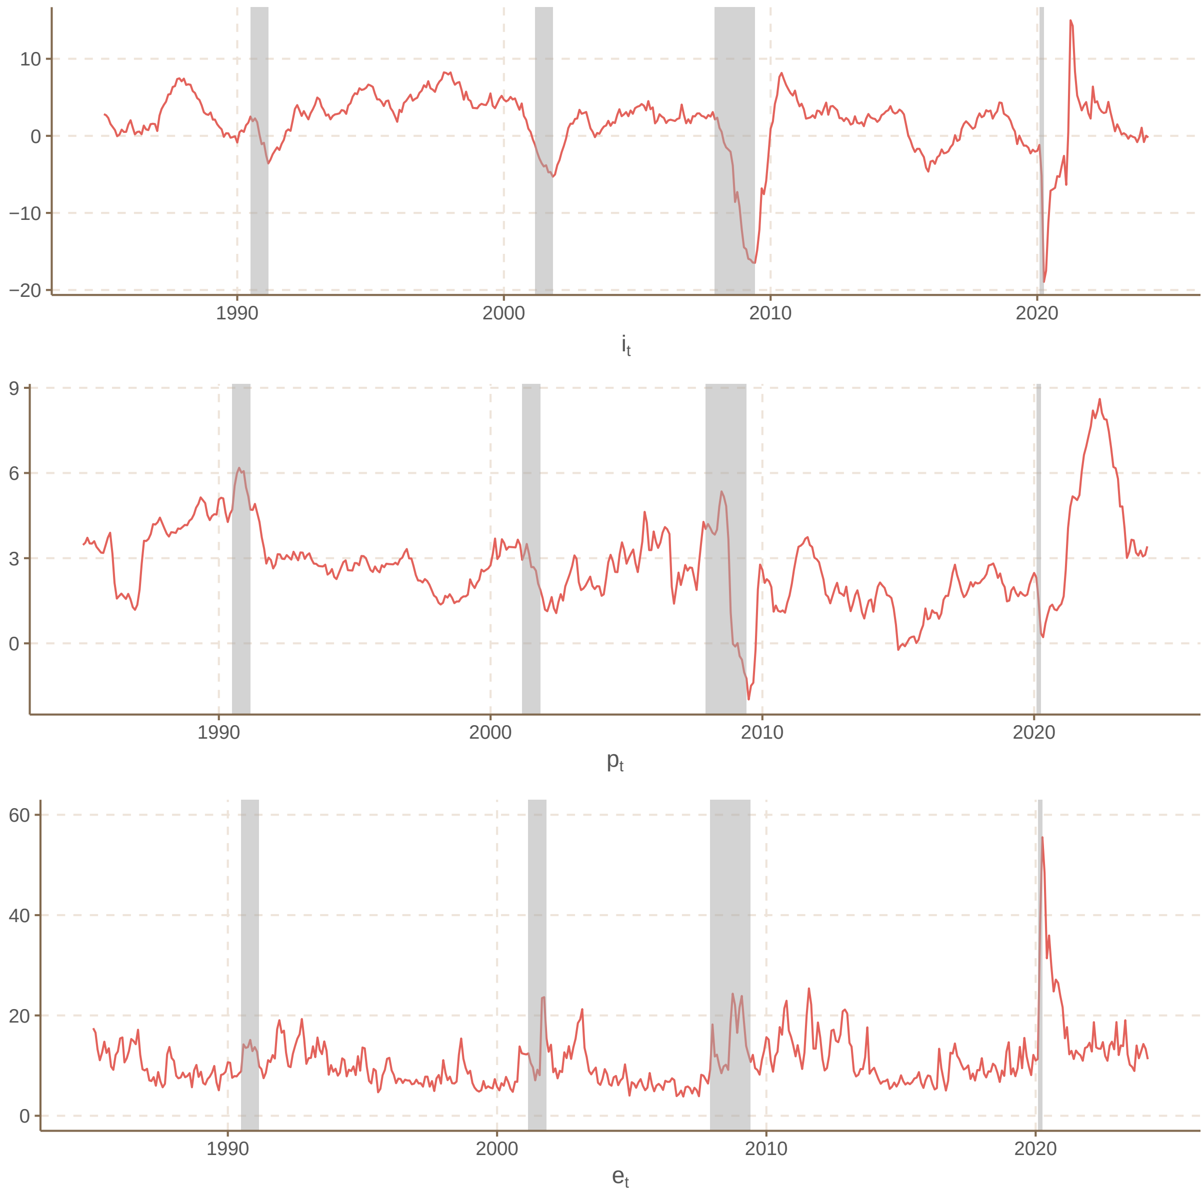

Figure 1: Data.

Industrial production growth rate (i_t) and CPI inflation (p_t) are given as 100 times the twelve-month log-difference. The plotted EPU (e_t) is the original series divided by 10. The shaded areas highlight the recession periods based on the NBER indicators.

3.1 Within-Sample Analysis

Our within-sample analysis focuses on the properties of a trivariate VAR(3) model[3] with dynamic skewness, where the vector of dependent variables is defined as $\mathbf{y}_t = (i_t, p_t, e_t)'$.[4]

Initially, we note the importance of stochastic volatility for all variables; see Figure 2. Generally, volatility varies substantially over time and, in particular, it rises around crisis times. This finding gives support to our choice to take stochastic volatility as a given feature. Interestingly, the estimated stochastic volatility is barely affected by allowing for non-Gaussian innovations; the estimates are slightly higher for the Gaussian model, but differences are often within the credible intervals (larger deviations are only present for some periods for the EPU index). This suggests that the dispersion of the innovation distribution is mostly captured well by stochastic volatility; potentially fat tails play a secondary role in capturing tail observations.[5]

Figure 3 shows the evolution of the skewness parameters for the trivariate VAR model with dynamic skewness. Despite the relatively wide credible intervals, we detect time variation in the asymmetry parameter for the variables. The EPU index – which typically has a positive skewness – exhibits the strongest time variation; it tends to increase in association with crises and decrease afterwards. For inflation, we observe a gradual decline in skewness, from an apparently symmetric distribution (no skewness) at the beginning of the sample to some evidence for negative skewness later on. This negative skewness is in line with the downside risk of inflation typically found between the mid-1990s and the early 2020s by De Polis et al. (2024), and for a somewhat shorter period by Le Bihan et al. (2023). The evidence for (time-varying) skewness is the weakest for industrial production, where credible intervals basically always include zero for the skewness parameter (implying a symmetric distribution), except for a short period in the mid 2010s, when skewness turns slightly positive. While this should not be over-interpreted, we note that industrial production growth can be seen as representing the evolution of the real economy and from

[3]The number of lags is set based on the Schwarz (1978) criterion applied to VARs assumed to have Gaussian and homoskedastic disturbances which have been estimated with maximum likelihood.

[4]Chan and Eisenstat (2018) and Karlsson et al. (2023) compute the marginal likelihood for model comparisons using an importance sampling method. Since the VAR dynSkew-T.SV model contains high-dimensional latent state space parameters, integrating the latent variables in the importance sampling steps and coming up with a good proposal sampling distribution is challenging. Therefore, formal model selection based on marginal likelihoods is beyond the scope of this paper. It is, however, an interesting topic for future research.

[5]The posterior mean of the degrees of freedom parameter is approximately 15, suggesting somewhat heavier-than-Gaussian tails.

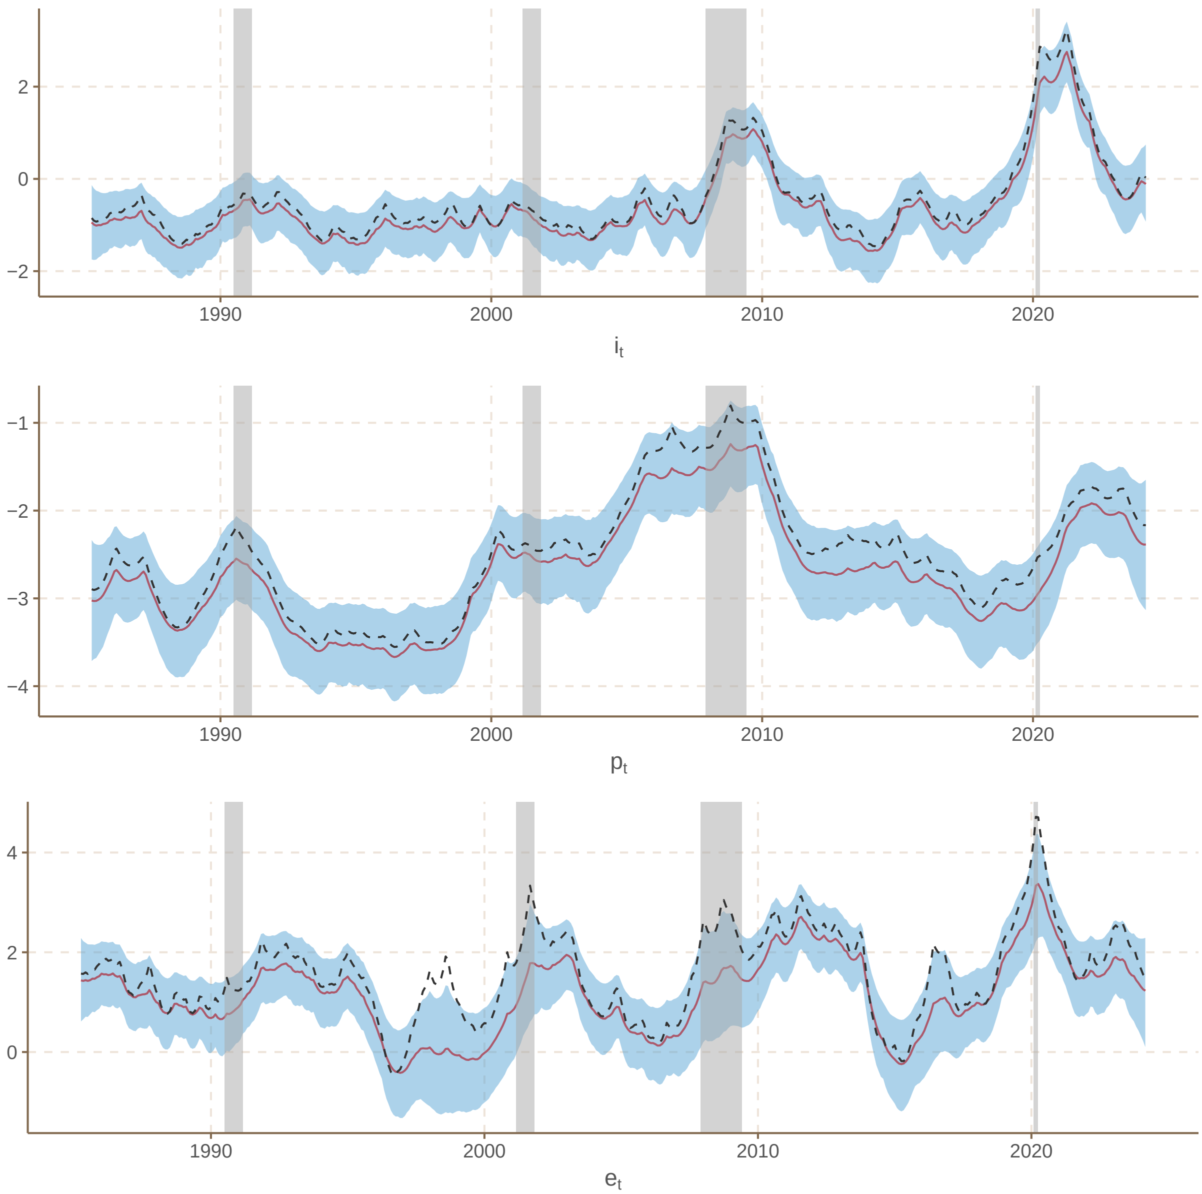

Figure 2: Estimated log stochastic volatility.

i_t is industrial production growth rate, p_t is CPI inflation and e_t is the EPU index (divided by ten). Estimated log volatility of the VAR model with fat tails and dynamic skewness is given by the red solid line; the coloured band gives the 80% credible interval. The dashed black line shows the estimated mean log volatility from the VAR model with Gaussian error terms. The shaded areas highlight the recession periods based on the NBER indicators.

that perspective, the positive skewness found stands in contrast with the negative skewness that Iseringhausen (2024) tended to find for GDP growth; it is, however, in line with the finding of Karlsson et al. (2023). In addition, due to the dynamics of the model and the contemporaneous correlation between the (reduced form) shocks, the model can still capture asymmetry in industrial production growth, such as larger negative movements during recessions (see Figure 1).[6]

Looking at the data series in Figure 1 together with the estimates of the stochastic volatility and the conditional skewness in Figures 2 and 3, we see that large movements in the variables around crisis times is mostly captured by an increase in the volatility and less so by an increase in the asymmetry. That is, the model attributes large movements to draws from an innovation distribution that is more spread-out, rather than draws from a distribution with increased asymmetry. The only exception is the EPU index, where not only dynamic skewness is strongest, but the difference in the stochastic volatility estimates for models with and without skewness are the most salient (last panel of Figure 2).

Next, we present the VAR dynamics using the impulse-response functions based on unit shocks. In this case, the impulse-response functions only depend on $\mathbf{A}$ and the dynamic parameters of the model ($\mathbf{B}_j$, $j = 1, \ldots, p$), which are constant over time, so responses to unit shocks are also time invariant. These impulse-response functions, together with the 80 percent credible intervals, are shown in Figure 4.[7]

The impulse-response functions are in line with what one would expect. In particular, the shock to the EPU index (third column in Figure 4) is in line with the findings of the previous literature in that it holds back real activity – here given by industrial production growth (first row). In addition, the effect of the shock to the EPU index resembles a (negative) aggregate demand shock in that it also has a negative effect on inflation (second row). Regarding the other shocks, it can be noted that the point estimate of a shock to industrial production growth (first column) also resembles

[6]As a general caveat, it should be noted that model misspecification could affect our results. Since the model features dynamic skewness but constant parameters in the conditional mean equation, time variation in the conditional mean of one or more of the variables in the model – such as the time-varying "trend inflation" in, for example, Chan et al. (2018) – might partially be absorbed by the skewness parameter. To investigate this further, a model with time variation both in the skewness equation and the conditional mean equation would have to be estimated. However, such an extension is beyond the scope of the current paper.

[7]Our model does feature stochastic volatility, so the impulse-response functions of a one standard deviation shock are time-varying. In this case, time variation solely comes from the time-varying standard deviation, presented in Figure 2. By presenting responses to unit shocks, we isolate the dynamics of the conditional mean from the stochastic volatility, and hence make interpretation of the impulse-response functions easier.

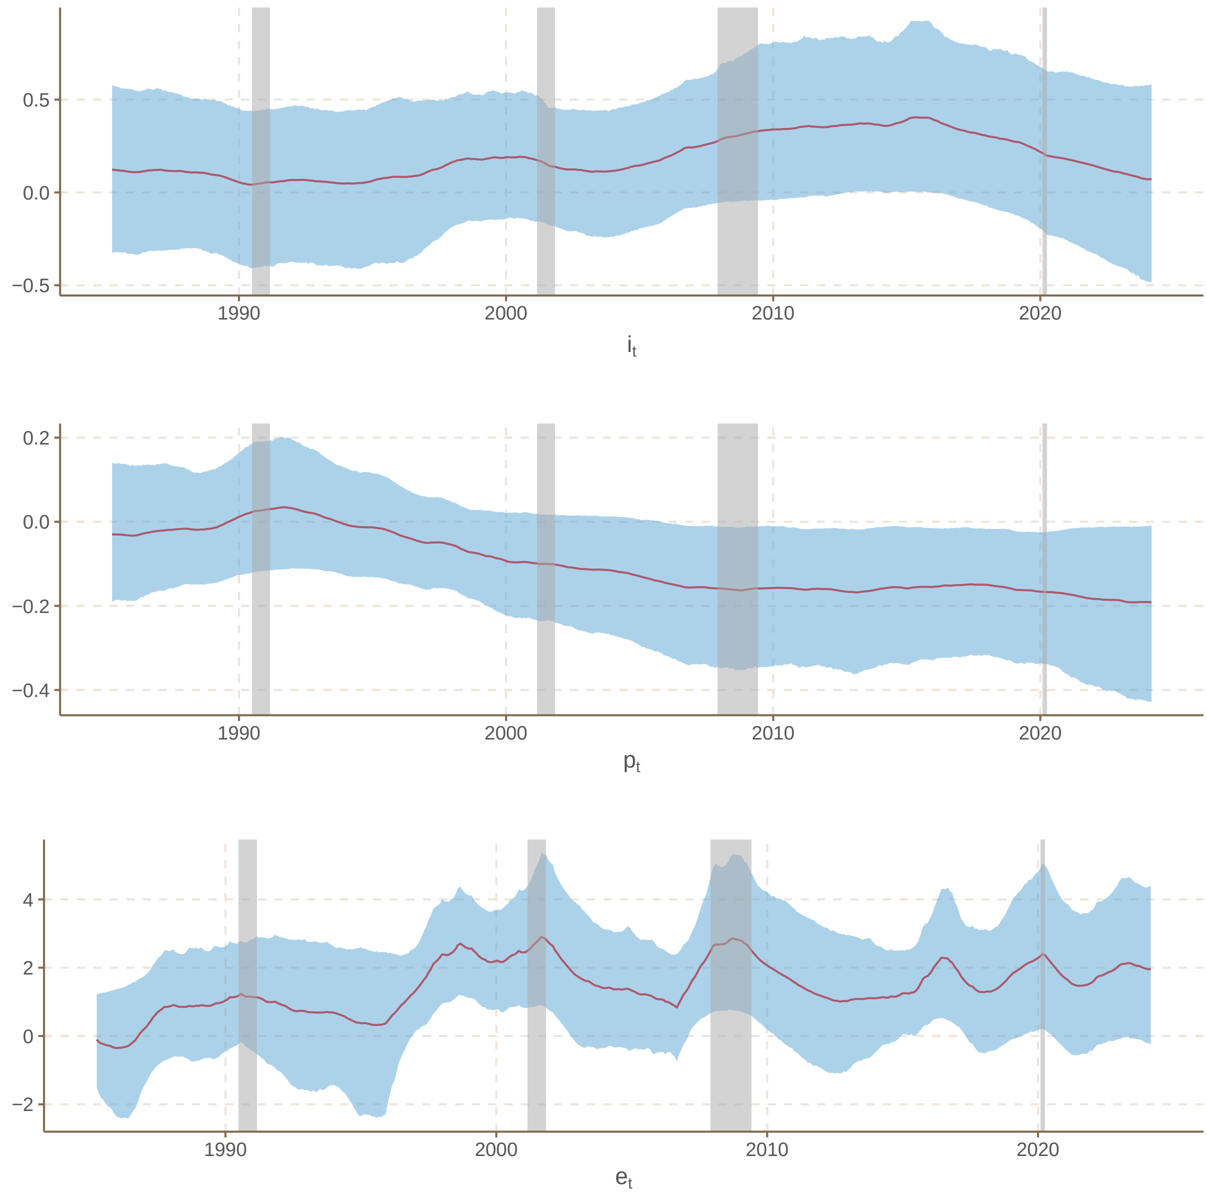

Figure 3: Estimated time-varying skewness.

i_t is industrial production growth rate, p_t is CPI inflation and e_t is the EPU index (divided by ten). Estimated skewness (posterior mean) of the VAR model with fat tails and dynamic skewness is given by the red solid line; the coloured band gives the 80% credible interval. The shaded areas highlight the recession periods based on the NBER indicators.

an aggregate demand shock as it is associated with an increase in inflation and a decrease in the EPU index. It should be noted though that the increase in inflation is not significant; the credible interval always includes zero. Finally, a shock to inflation (second column) looks very much like a supply-side (cost-push-type) shock; the effect on industrial production growth is negative (after a short, slightly positive effect) and the EPU increases.

Overall, the within-sample evidence shows that there is some support for time-varying skewness. And the dynamics of the model – illustrated by the impulse-response functions – are in line with mainstream findings from the previous literature on the relation between the EPU index and the macroeconomy.

3.2 Forecasting Performance

We next consider how different model features – in particular dynamic skewness – affect forecasting performance; both point and density forecasts are considered. Five different forecast horizons are assessed: one month, two months, three months, six months and one year. Our benchmark model is the univariate AR(3) model with normally distributed error terms. This is first compared to univariate specifications with error terms that allow for fat tails and asymmetry, where we consider both constant and time-varying skewness. We then also compare the benchmark model to trivariate VAR specifications, where three different trivariate specifications are used: Gaussian innovations, fat-tailed innovations with constant skewness, and fat-tailed innovations with time-varying skewness. All models are estimated allowing for stochastic volatility.

The forecast evaluation is performed on an expanding sample. We use the mean squared forecast error (MSFE) for point forecast evaluation. At forecast horizon h, the MSFE is calculated as

$$\mathrm{MSFE}_{i,h} = \frac{1}{T - T_0 - h + 1} \sum_{t=T_0}^{T-h} \left[\left(\bar{y}_{i,t+h|t} - y^o_{i,t+h}\right)^2 \right],$$

where $\bar{y}_{i,t+h|t}$ is the point forecast – the posterior mean – using all the information up to time t and $y^o_{i,t+h}$ is the realization of the variable i, h-steps ahead. A lower MSFE indicates better model performance.

For density forecasts, we use the log predictive score (LPS). The LPS is defined such that a

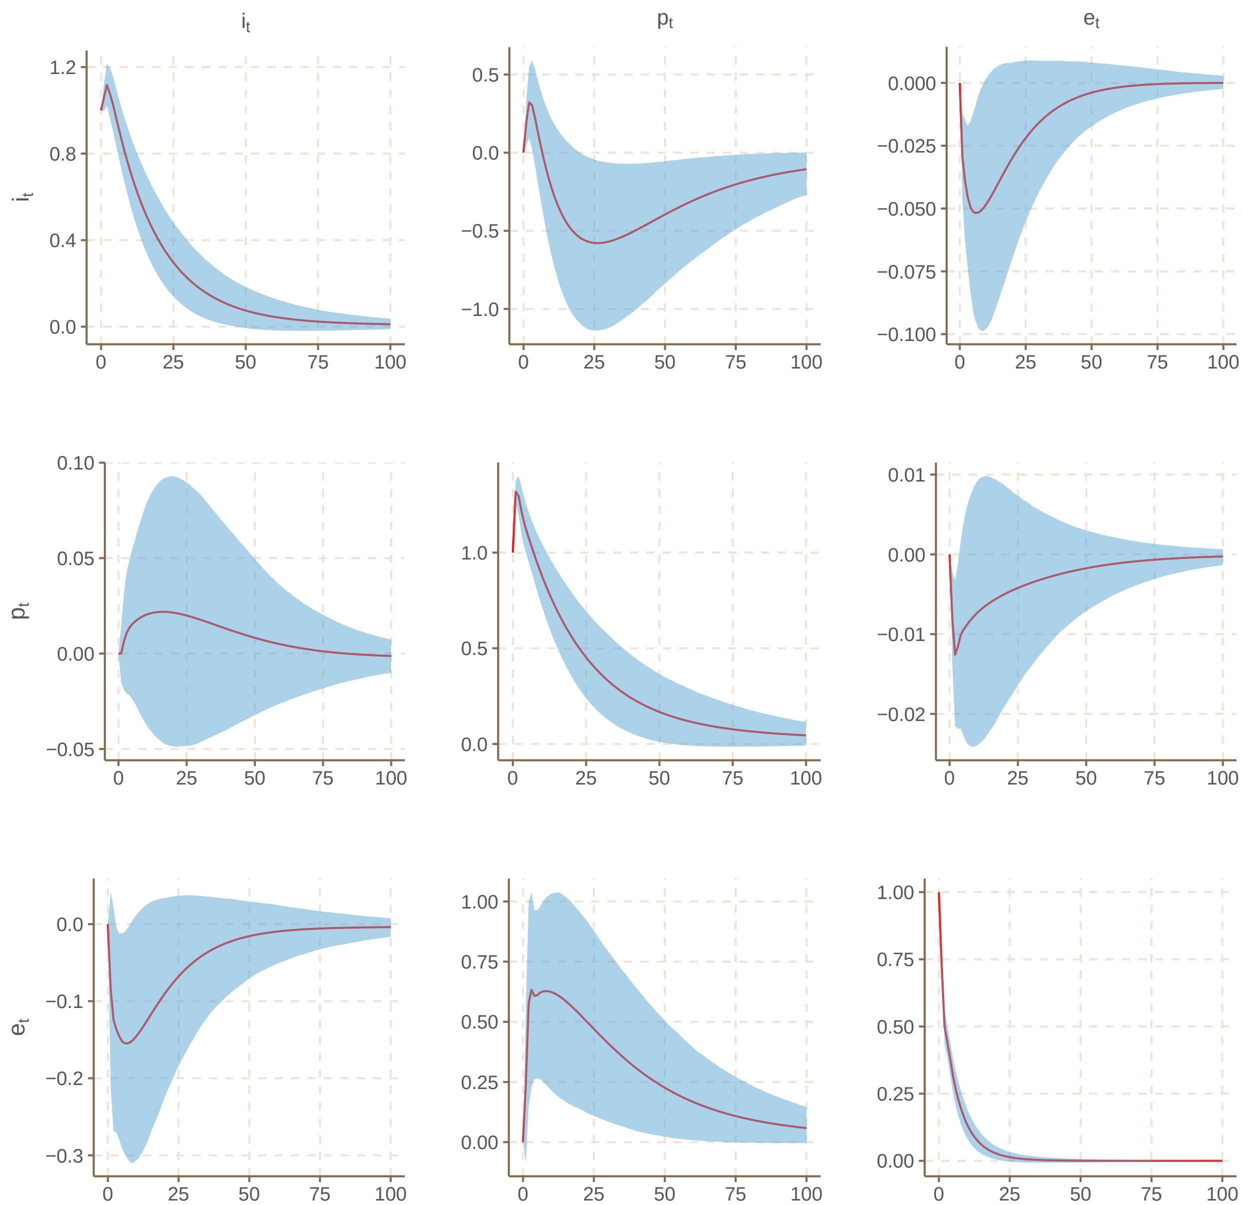

Figure 4: The impulse response functions of the model with time-varying skewness and fat tails.

i_t is industrial production growth rate, p_t is CPI inflation and e_t is the EPU index (divided by ten). Impulses are shown in the column headers, response variables are given by the rows. The size of the impulse is one unit shock. The red solid line gives the point estimate and the coloured bands show the 80% credible interval. The size of the effect in percentage points is represented on the vertical axis. The horizontal axis shows the horizon in months.

larger value indicates a better forecast performance of the model and is given by

$$\begin{aligned}
\mathrm{LPS}_{i,h} =& \frac{1}{T_1 - T_0 - h + 1} \sum_{t=T_0}^{T_1-h} \left[\log \int p(y^o_{i,t+h} | \boldsymbol{\theta}, \mathbf{y}_{1:t}) p(\boldsymbol{\theta} | \mathbf{y}_{1:t}) d\boldsymbol{\theta} \right] \\
=& \frac{1}{T_1 - T_0 - h + 1} \sum_{t=T_0}^{T_1-h} \left[\log \int \int p(y^o_{i,t+h} | \boldsymbol{\theta}, \mathbf{y}_{1:t}, \mathbf{y}_{t+1:t+h-1}, \boldsymbol{\theta}_{t+1:t+h}) \right. \\
& \left. p(\mathbf{y}_{t+1:t+h-1}, \boldsymbol{\theta}_{t+1:t+h} | \boldsymbol{\theta}, \mathbf{y}_{1:t}) p(\boldsymbol{\theta} | \mathbf{y}_{1:t}) d\mathbf{y}_{t+1:t+h-1} d\boldsymbol{\theta}_{t+1:t+h} d\boldsymbol{\theta} \right] \\
\approx& \frac{1}{T - T_0 - h + 1} \sum_{t=T_0}^{T-h} \left[\log \frac{1}{R} \sum_{r=1}^{R} p(y^o_{i,t+h} | \boldsymbol{\theta}^{(r)}, \mathbf{y}_{1:t}, \mathbf{y}^{(r)}_{t+1:t+h-1}, \boldsymbol{\theta}^{(r)}_{t+1:t+h}) \right],
\end{aligned}$$

where $\boldsymbol{\theta}^{(r)}$ is a posterior sample of the (time-invariant) parameters of the TV skewness VAR model for $r = 1, \ldots, R$; $\mathbf{y}^{(r)}_{t+1:t+h-1}$ represents a simulated path of intermediate observations between $t+1$ and $t+h-1$; $\boldsymbol{\theta}^{(r)}_{t+1:t+h}$ represents a simulated path of the time-varying parameters between $t+1$ and $t+h$; and $p(y^o_{i,t+h} | \boldsymbol{\theta}^{(r)}, \mathbf{y}_{1:t}, \mathbf{y}^{(r)}_{t+1:t+h-1}, \boldsymbol{\theta}^{(r)}_{t+1:t+h})$ is the 1-step ahead posterior predictive density function (the score) evaluated at the realization of the variable, conditional on the posterior sample of parameters and the simulated path.

To compare both point and density forecasting performance, we apply the one-sided test of Diebold and Mariano (1995), where the standard errors are calculated using the Newey-West estimator for the standard errors, in line with previous work by Clark (2011) and Clark and Ravazzolo (2015). Forecasts are generated from January 2010 to March 2023 for all horizons, which allows for evaluating 159 forecasts uniformly over all horizons.

Tables 1 and 2 show forecast comparison results for point and density forecasts, respectively. The first row in each panel (for each variable) contains the MSFE and the LPS values for the univariate Gaussian model, and in the rest of the rows, relative values (relative MSFE and difference in LPS) are shown.

Looking at point forecasts, no quantitatively meaningful improvements can be found relative to the forecasts of the baseline univariate Gaussian model for industrial production growth or CPI inflation (even if the Diebold-Mariano test indicates that the univariate model with fat tails and dynamic skewness has a significantly lower MSFE at the three-month horizon for industrial production growth). In contrast, when forecasting the EPU index, allowing for fat tails and constant or dynamic skewness improves over Gaussian innovations in a univariate framework at the two longest horizons. Turning to the trivariate models, significant improvements over the benchmark model can only be found in one case, namely when forecasting the EPU index at the six-month horizon using the model with fat tails and dynamic skewness; it can, however, be noted that this MSFE is higher than that of the corresponding univariate model. The overall impression is that the trivariate models are not particularly successful compared to the univariate models when it comes to forecasting.

Regarding the density forecasts, we find no systematic improvements *vis-á-vis* the univariate Gaussian model. While there are some cases where the LPS is higher for the model with fat tails and dynamic skewness, these improvements are never significant. Similarly, using the trivariate VAR model never leads to significant improvements in density predictions.

Table 1: MSFEs

	1M	2M	3M	6M	12M
	(a) Industrial production growth				
Univariate AR(3)					
G.SV	4.465	11.131	14.202	16.373	23.982
Skew-T.SV	1.000	1.000	1.005	1.004	1.017
dynSkew-T.SV	0.994	0.997	0.997†	1.001	1.016
Trivariate VAR(3)					
G.SV	0.994	1.021	1.056	1.153	1.202
Skew-T.SV	0.994	1.030	1.069	1.152	1.191
dynSkew-T.SV	0.989	1.026	1.063	1.146	1.190
	(b) CPI inflation				
Univariate AR(3)					
G.SV	0.104	0.334	0.629	1.565	3.395
Skew-T.SV	1.003	1.003	1.006	1.005	1.010
dynSkew-T.SV	1.007	1.007	1.008	1.007	1.011
Trivariate VAR(3)					
G.SV	1.092	1.128	1.134	1.109	1.116
Skew-T.SV	1.094	1.123	1.127	1.099	1.101
dynSkew-T.SV	1.090	1.118	1.120	1.091	1.102
	(c) EPU/10				
Univariate AR(3)					
G.SV	16.699	33.934	43.480	59.102	68.458
Skew-T.SV	1.043	1.006	0.977	0.933*	0.883*
dynSkew-T.SV	1.046	1.010	0.981	0.939*	0.887*
Trivariate VAR(3)					
G.SV	1.004	0.999	0.989	0.986	0.986
Skew-T.SV	1.025	0.997	0.976	0.958	0.938
dynSkew-T.SV	1.024	0.992	0.970	0.949*	1.009

The first line in each panel reports the MSFE of the benchmark univariate model with Gaussian innovations based on forecasts generated from 2010M01 to 2023M03 (159 recursive estimations). The relative performance (all other rows) is computed as the ratio of the MSFE of alternative specifications over the benchmark, that is, entries less than 1 indicate that a given model has a lower MSFE. * denotes that the model in question significantly outperforms the benchmark model at the 5% level based on the one-sided Diebold and Mariano (1995) test where the standard errors of the test statistics are computed with the Newey–West estimator.

Table 2: LPSs

	1M	2M	3M	6M	12M
	(a) Industrial production growth				
Univariate AR(3)					
G.SV	-1.577	-2.149	-2.433	-2.652	-2.966
Skew-T.SV	0.009	0.109	-0.048	0.001	-0.006
dynSkew-T.SV	0.041	0.000	0.075	0.001	-0.007
Trivariate VAR(3)					
G.SV	-0.021	-0.119	-0.034	-0.209	-0.246
Skew-T.SV	-0.058	-0.138	-0.087	-0.181	-0.232
dynSkew-T.SV	-0.047	-0.116	-0.141	-0.214	-0.275
	(b) CPI inflation				
Univariate AR(3)					
G.SV	-0.294	-0.874	-1.197	-1.639	-1.999
Skew-T.SV	0.009	0.005	0.000	-0.010	-0.066
dynSkew-T.SV	0.005	-0.011	-0.030	-0.035	-0.219
Trivariate VAR(3)					
G.SV	-0.031	-0.048	-0.058	-0.037	-0.016
Skew-T.SV	-0.037	-0.058	-0.085	-0.098	-0.151
dynSkew-T.SV	-0.039	-0.076	-0.102	-0.121	-0.228
	(c) EPU/10				
Univariate AR(3)					
G.SV	-2.537	-2.816	-2.944	-3.081	-3.243
Skew-T.SV	-0.006	0.019	0.052	-0.014	0.073
dynSkew-T.SV	-0.011	0.036	0.043	-0.009	0.029
Trivariate VAR(3)					
G.SV	-0.001	0.006	0.022	0.024	0.045
Skew-T.SV	-0.002	0.018	0.042	0.028	0.078
dynSkew-T.SV	-0.002	0.010	0.004	-0.004	0.013

The first line in each panel reports the LPS of the benchmark univariate model with Gaussian innovations based on forecasts generated from 2010M01 to 2023M03 (159 recursive estimations). The relative performance is computed as the difference of the LPS of alternative specifications and the benchmark, that is, entries smaller than 0 indicate that a given model has a higher LPS. * denotes that the model in question significantly outperforms the benchmark model at the 5% level based on the one-sided Diebold and Mariano (1995) test where the standard errors of the test statistics are computed with the Newey–West estimator.

4 Conclusions

In this paper, we have extended the standard Bayesian VAR with stochastic volatility to allow for generalized hyperbolic skew Student's t distribution in the innovations, where the skewness parameter is also allowed to vary over time. This specification permits a flexible modelling of innovations in terms of fat tails and – potentially dynamic – asymmetry. These might be useful features when modelling the macroeconomy, given the recent literature documenting the non-Gaussianity of many macroeconomic variables. Our model nests specifications with Gaussian and fat-tailed innovations, as well as the specification where asymmetry is constant over time, allowing us to analyse the importance of these features in an empirical context.

In our empirical application, we revisit how aggregate uncertainty (captured by the EPU index) affects the macroeconomy. This is done by estimating a trivariate VAR model including industrial production growth, CPI inflation and the EPU index. Apart from the usual strong evidence for time-varying volatility, we also document that the estimated skewness of the variables – primarily the EPU index – changes over time. However, we also find that modelling innovations flexibly typically does not improve the point or density forecasts. Finally, we confirm previous findings in the literature that shocks to uncertainty resemble aggregate demand shocks, dampening both real activity and inflation.

Acknowledgments

Hoang Nguyen, Stepan Mazur and Pär Österholm acknowledge financial support from the project "Improved Economic Policy and Forecasting with High-Frequency Data" (Dnr: E47/22) funded by the Torsten Söderbergs Foundation. Stepan Mazur also acknowledges financial support from the internal research grants at Örebro University.

References

D. Acemoglu and A. Scott. Asymmetric business cycles: Theory and time-series evidence. *Journal of Monetary Economics*, 40(3):501–533, 1997.

F. Q. Akram and H. Mumtaz. Time-varying dynamics of the Norwegian economy. *The Scandinavian Journal of Economics*, 121(1):407–434, 2019.

G. Ascari, G. Fagiolo, and A. Roventini. Fat-tail distributions and business-cycle models. *Macroeconomic Dynamics*, 19(2):465–476, 2015.

S. R. Baker, N. Bloom, and S. J. Davis. Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

M. Balcilar, R. Gupta, and M. Segnon. The role of economic policy uncertainty in predicting us recessions: A mixed-frequency Markov-switching vector autoregressive approach. *Economics*, 10 (1):20160027, 2016.

G. Caggiano, E. Castelnuovo, and J. M. Figueres. Economic policy uncertainty and unemployment in the United States: A nonlinear approach. *Economics Letters*, 151:31–34, 2017.

A. Carriero, T. E. Clark, and M. Marcellino. Realtime nowcasting with a Bayesian mixed frequency model with stochastic volatility. *Journal of the Royal Statistical Society. Series A,(Statistics in Society)*, 178(4):837, 2015.

A. Carriero, T. E. Clark, and M. Marcellino. Measuring uncertainty and its impact on the economy. *Review of Economics and Statistics*, 100(5):799–815, 2018.

A. Carriero, T. E. Clark, and M. Marcellino. Capturing macro-economic tail risks with Bayesian vector autoregressions. *Journal of Money, Credit and Banking*, 2020.

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

J. C. Chan. The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling. *Journal of Business and Economic Statistics*, 35(1):17–28, 2017.

J. C. Chan and E. Eisenstat. Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *Journal of Applied Econometrics*, 33(4):509–532, 2018.

J. C. Chan, T. E. Clark, and G. Koop. A new model of inflation, trend inflation, and long-run inflation expectations. *Journal of Money, Credit and Banking*, 50(1):5–53, 2018.

C.-W. J. Chiu, H. Mumtaz, and G. Pinter. Forecasting with VAR models: Fat tails and stochastic volatility. *International Journal of Forecasting*, 33(4):1124–1143, 2017.

T. E. Clark. Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility. *Journal of Business and Economic Statistics*, 29(3):327–341, 2011.

T. E. Clark and F. Ravazzolo. Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics*, 30(4):551–575, 2015.

T. Cogley and T. J. Sargent. Drifts and volatilities: Monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262–302, 2005.

J. Cross and A. Poon. Forecasting structural change and fat-tailed events in Australian macroeconomic variables. *Economic Modelling*, 58:34–51, 2016.

A. De Polis, L. Melosi, and I. Petrella. The taming of the skew: Asymmetric inflation risk and monetary policy. *Working Paper*, 2024. `https://andreadepolis.github.io/workingpapers/inflationrisk/DPMP2024.pdf`.

M. Del Negro and G. E. Primiceri. Time varying structural vector autoregressions and monetary policy: A corrigendum. *The Review of Economic Studies*, 82(4):1342–1345, 2015.

D. Delle Monache, A. De Polis, and I. Petrella. Modeling and forecasting macroeconomic downside risk. *Journal of Business and Economic Statistics*, 42(3):1010–1025, 2024.

F. X. Diebold and R. S. Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3):134–144, 1995.

G. Fagiolo, M. Napoletano, and A. Roventini. Are output growth-rate distributions fat-tailed? Some evidence from OECD countries. *Journal of Applied Econometrics*, 23(5):639–669, 2008.

H. Gulen and M. Ion. Policy uncertainty and corporate investment. *The Review of Financial Studies*, 29(3):523–564, 2016.

W. Hörmann and J. Leydold. Generating generalized inverse Gaussian random variates. *Statistics and Computing*, 24(4):547–557, 2014.

M. Iseringhausen. The time-varying asymmetry of exchange rate returns: A stochastic volatility–stochastic skewness model. *Journal of Empirical Finance*, 58:275–292, 2020.

M. Iseringhausen. A time-varying skewness model for Growth-at-Risk. *International Journal of Forecasting*, 40(1):229–246, 2024.

P. M. Jones and E. Olson. The time-varying correlation between uncertainty, output, and inflation: Evidence from a DCC-GARCH model. *Economics Letters*, 118(1):33–37, 2013.

S. Karlsson and P. Österholm. A hybrid time-varying parameter Bayesian VAR analysis of Okun's law in the United States. *Economics Letters*, 197:109622, 2020a.

S. Karlsson and P. Österholm. A note on the stability of the Swedish Phillips curve. *Empirical Economics*, 59(6):2573–2612, 2020b.

S. Karlsson and P. Österholm. Is the US Phillips curve stable? Evidence from Bayesian vector autoregressions. *Scandinavian Journal of Economics*, 125(1):287–314, 2023.

S. Karlsson, S. Mazur, and H. Nguyen. Vector autoregression models with skewness and heavy tails. *Journal of Economic Dynamics and Control*, 146:104580, 2023.

L. Karnizova and J. C. Li. Economic policy uncertainty, financial markets and probability of us recessions. *Economics Letters*, 125(2):261–265, 2014.

G. Kastner and S. Frühwirth-Schnatter. Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics and Data Analysis*, 76:408–423, 2014.

S. Kim, N. Shephard, and S. Chib. Stochastic volatility: Likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393, 1998.

T. Kiss, S. Mazur, H. Nguyen, and P. Österholm. Modeling the relation between the US real economy and the corporate bond-yield spread in Bayesian VARs with non-Gaussian innovations. *Journal of Forecasting*, 42(2):347–368, 2023.

G. Koop and D. Korobilis. *Bayesian multivariate time series methods for empirical macroeconomics.* Now Publishers Inc, 2010.

G. Koop and D. Korobilis. Forecasting with high-dimensional panel VARs. *Oxford Bulletin of Economics and Statistics*, 81(5):937–959, 2019.

H. Le Bihan, D. Leiva-Leon, and M. Pacce. Underlying inflation and asymmetric risks. *ECB Working Paper Series 2848*, 2023.

S. Leduc and Z. Liu. Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, 82:20–35, 2016.

X. Liu. On tail fatness of macroeconomic dynamics. *Journal of Macroeconomics*, 62:103154, 2019.

M. W. McCracken and S. Ng. FRED-MD: A monthly database for macroeconomic research. *Journal of Business and Economic Statistics*, 34(4):574–589, 2016.

A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative risk management: Concepts, Techniques and Tools-revised edition.* Princeton university press, 2015.

E.-Y. Nam, K. Lee, and Y. Jeon. Macroeconomic uncertainty shocks and households' consumption choice. *Journal of Macroeconomics*, 68:103306, 2021.

S. Ni and D. Sun. Bayesian estimates for vector autoregressive models. *Journal of Business and Economic Statistics*, 23(1):105–117, 2005.

G. E. Primiceri. Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72(3):821–852, 2005.

G. O. Roberts and J. S. Rosenthal. Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.

G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, pages 461–464, 1978.

C. A. Sims and T. Zha. Were there regime switches in US monetary policy? *American Economic Review*, 96(1):54–81, 2006.

G. Tsiotas. On generalised asymmetric stochastic volatility models. *Computational Statistics & Data Analysis*, 56(1):151–172, 2012.