

Conditional density forecasting: a tempered importance sampling approach

Carlos Montes-Galdón
European Central Bank

Joan Paredes
European Central Bank

Elias Wolf
University of Bonn

Disclaimer: This presentation should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Forecasting literature offers different methods how to condition multivariate forecasts on future paths of a subset of model variables (e.g. Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) or more recently Chan et al. (2023))

—→ Limiting the information set to future paths of variables ignores the evolution of risks around these forecasts

Forecasting literature offers different methods how to condition multivariate forecasts on future paths of a subset of model variables (e.g. Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) or more recently Chan et al. (2023))

—→ Limiting the information set to future paths of variables ignores the evolution of risks around these forecasts

Conditional Density Forecasting: new methodology to condition model-based forecasting densities on off-model information (e.g. from options) about some of its marginal distributions.

Forecasting literature offers different methods how to condition multivariate forecasts on future paths of a subset of model variables (e.g. Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) or more recently Chan et al. (2023))

—→ Limiting the information set to future paths of variables ignores the evolution of risks around these forecasts

Conditional Density Forecasting: new methodology to condition model-based forecasting densities on off-model information (e.g. from options) about some of its marginal distributions.

Idea: Reweight the draws from the model-based marginals of the forecasting density $q_{\theta}(y)$ to match a target density $p_{\eta}(y)$ that satisfies external information.

Forecasting literature offers different methods how to condition multivariate forecasts on future paths of a subset of model variables (e.g. Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) or more recently Chan et al. (2023))

—→ Limiting the information set to future paths of variables ignores the evolution of risks around these forecasts

Conditional Density Forecasting: new methodology to condition model-based forecasting densities on off-model information (e.g. from options) about some of its marginal distributions.

Idea: Reweight the draws from the model-based marginals of the forecasting density $q_\theta(y)$ to match a target density $p_\eta(y)$ that satisfies external information.

How: We propose an algorithm that uses **tempered importance sampling** to adjust the model-based forecasting densities to the user specified target densities

Forecasting literature offers different methods how to condition multivariate forecasts on future paths of a subset of model variables (e.g. Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) or more recently Chan et al. (2023))

—→ Limiting the information set to future paths of variables ignores the evolution of risks around these forecasts

Conditional Density Forecasting: new methodology to condition model-based forecasting densities on off-model information (e.g. from options) about some of its marginal distributions.

Idea: Reweight the draws from the model-based marginals of the forecasting density $q_{\theta}(y)$ to match a target density $p_{\eta}(y)$ that satisfies external information.

How: We propose an algorithm that uses **tempered importance sampling** to adjust the model-based forecasting densities to the user specified target densities

Why: Understanding possibly asymmetric risks around forecasts of macroeconomic variables

Methodological contribution:

- Motivation is similar to the idea of **entropic tilting** developed in Robertson, Tallman, and Whiteman (2005) and Krüger, Clark, and Ravazzolo (2017)
- The performance of the entropic tilting methodology crucially hinges on the support of the original distribution Q_θ
- Application of Tempered Importance Sampling by Herbst and Schorfheide (2014)
→ Our approach is more flexible and robust.

Macro@Risk Literature:

- Analyse asymmetric risks around macroeconomic forecasts (Adrian, Boyarchenko, and Giannone (2019), Montes-Galdón and Ortega (2022) or Wolf (2022)).
- Multivariate skew-T distribution of Azzalini and Capitanio (2003) to capture asymmetric densities in option-implied moments.

Basic Idea: Importance Sampling

Proposal

The researcher wants to introduce external information on a subset of variables $y_i^e \in y_i$ that is given by the distribution P_η into a **model-based forecast**:

$$y_i \sim q_\theta(y)$$

Basic Idea: Importance Sampling

Proposal

The researcher wants to introduce external information on a subset of variables $y_i^e \in y_i$ that is given by the distribution P_η into a **model-based forecast**:

$$y_i \sim q_\theta(y)$$

Target

The researcher can re-weight the draws $\{y_i\}_{i=1}^N$ so that draws y_i^e from the final **forecasting density** satisfy the information in $p_\eta(y_i)$ using **importance sampling**:

$$\tilde{y}_i \sim \mathcal{MN}(\{y_i\}_{i=1}^N | \{W_i\}_{i=1}^N) \quad \text{with} \quad W_i = \frac{w_i}{\sum_{i=1}^N w_i} \quad \text{and} \quad w_i = \frac{p_\eta(y_i^e)}{q_\theta(y_i^e)}$$

The tuples $\{y_i, W_i\}_{i=1}^N$ provide a **particle approximation** of the final density.

Quality of the importance sampler depends on the Kullback-Leibler Divergence $D(P_\eta|Q_\theta)$.

- In our case, the proposal density is predetermined: \rightarrow It is a model-based (forecasting) density $q_\theta(y_i)$.
- Target density $p_\eta(y_i)$ might be "far apart" with a high Kullback-Leibler divergence.
 1. For high-dimensional $q_\eta(y_i)$ the probability mass is concentrated in a small region of the high dimensional space
 2. External information that will improve or alter the forecast density might often imply a very different mean and variance of $p_\eta(y_i)$ compared to $q_\theta(y_i)$.

Quality of the importance sampler depends on the Kullback-Leibler Divergence $D(P_\eta|Q_\theta)$.

- In our case, the proposal density is predetermined: \rightarrow It is a model-based (forecasting) density $q_\theta(y_i)$.
- Target density $p_\eta(y_i)$ might be "far apart" with a high Kullback-Leibler divergence.
 1. For high-dimensional $q_\eta(y_i)$ the probability mass is concentrated in a small region of the high dimensional space
 2. External information that will improve or alter the forecast density might often imply a very different mean and variance of $p_\eta(y_i)$ compared to $q_\theta(y_i)$.

\rightarrow **To overcome these problems we resort to tempered importance sampling (Neal (2001), Herbst and Schorfheide (2014))**

Quality of the importance sampler depends on the Kullback-Leibler Divergence $D(P_\eta|Q_\theta)$.

- In our case, the proposal density is predetermined: \rightarrow It is a model-based (forecasting) density $q_\theta(y_i)$.
- Target density $p_\eta(y_i)$ might be "far apart" with a high Kullback-Leibler divergence.
 1. For high-dimensional $q_\eta(y_i)$ the probability mass is concentrated in a small region of the high dimensional space
 2. External information that will improve or alter the forecast density might often imply a very different mean and variance of $p_\eta(y_i)$ compared to $q_\theta(y_i)$.

\rightarrow **To overcome these problems we resort to tempered importance sampling (Neal (2001), Herbst and Schorfheide (2014))**

Additional Problems:

- External information might only be available for transformations of the model variables of interest (e.g. growth rates or log vs. levels) \rightarrow implies a change in the respective marginal distributions
- We consider Bayesian estimation setting without a closed form solution of the model-based predictive density $q_\theta(y_i)$.

Tempered Importance Sampling

Tempered Importance Sampling uses a sequence of N_ϕ bridge distributions

$$p_n(y_i|\mu_\eta, \Sigma_\eta/\phi_n) \quad \text{with} \quad 0 < \phi_1 < \dots < \phi_{N_\phi} = 1 \quad (1)$$

that converge towards the target distribution $p(y_i|\mu_\eta, \Sigma_\eta)$ for $\phi_n \rightarrow 1$.

1. Correction

$$W_{i,n} \propto \frac{p_n(y_{i,n-1}|\mu_\eta, \Sigma_\eta/\phi_n)}{p_{n-1}(y_{i,n-1}|\mu_\eta, \Sigma_\eta/\phi_{n-1})}$$

2. Selection

$$\tilde{y}_{i,n} \sim \mathcal{MN}(\{y_{i,n-1}\}_{i=1}^N | \{W_{i,n}\}_{i=1}^N)$$

3. Mutation

Propagate the resampled particles $\{\tilde{y}_i\}_{i=1}^N$ using M steps of an MH-Algorithm with a transition Kernel

$$y_{i,n} \sim K_n(y_n|\tilde{y}_{i,n})$$

that has the stationary distribution $p_n(y_i|\mu_\eta, \Sigma_\eta/\phi_n)$

Tempered Importance Sampling: Example

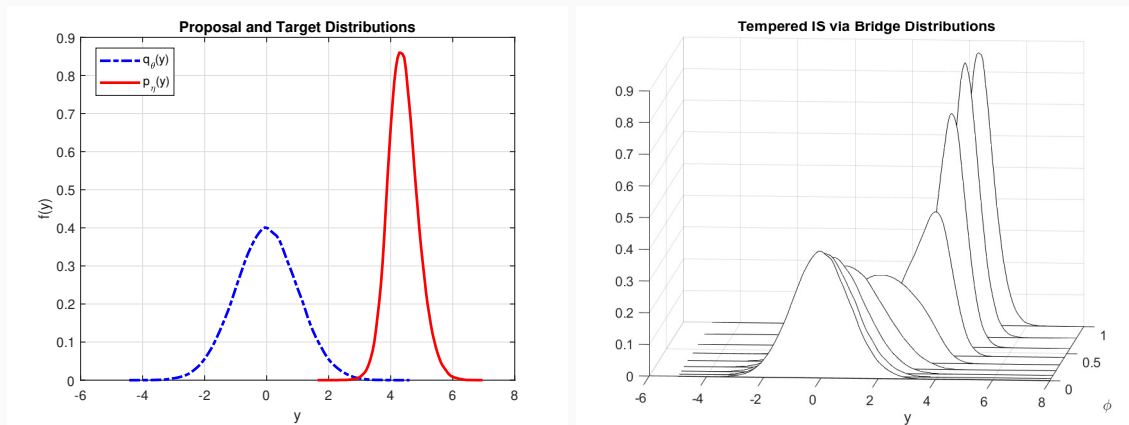


Figure 1: Proposal and Target densities $q_\theta(y)$ and $p_\eta(y)$ with probability masses concentrated in different regions of the real line and Kullback-Leibler divergence of 107,22. Bridge distributions are close by construction and converge to $p_\eta(y)$ as $\phi_n \rightarrow 1$.

- Draws from Q_θ are sequentially adapted to P_η using a number of bridge distributions that converge to the target density
- Number of bridge distributions N_{ϕ_N} is set adaptively to guarantee even weights. [Details](#)

Conditional Density Forecasting Algorithm:

For $i = 1, \dots, N$:

1. Draw $y_{i,1} \sim q\left(y_i | \mu_\theta^{(i)}, \Sigma_\theta^{(i)}\right)$
2. Select the subset $y_{i,1}^e \in y_{i,1}$ for which there exists external information on the transformation $h(y_{i,1}^e)$.

2.1 Obtain initial importance weights $W_{i,1} \propto \frac{p_1(h(y_{i,1}^e) | \mu_\eta, \Sigma_\eta / \phi_1)}{q(y_i | \mu_\theta^{(i)}, \Sigma_\theta^{(i)})}$

2.2 Resample $y_{i,1}^e \sim \mathcal{MN}(\{y_{i,1}^e\}_{i=1}^N | \{W_{i,1}\}_{i=1}^N)$

3. For $n = 2 : N_{\phi_N}$

3.1 **Correction:** Obtain weights $W_{i,n} \propto \frac{p_n(h(y_{i,n-1}^e) | \mu_\eta, \Sigma_\eta / \phi_n)}{p_{n-1}(h(y_{i,n-1}^e) | \mu_\eta, \Sigma_\eta / \phi_{n-1})}$

3.2 **Selection:** Resample $y_{i,n}^e \sim \mathcal{MN}(\{y_{i,n-1}^e\}_{i=1}^N | \{W_{i,n}\}_{i=1}^N)$

3.3 **Mutation:** For $j = 1 : H$

3.3.1 Draw $\hat{y}_{i,n}^e \sim q(y_i^e | y_{i,n}^e, \mu_\theta^{(i)}, c_n \Sigma_\theta^{(i)})$

3.3.2 Compute

$$\alpha = \frac{p_n(h(\hat{y}_{i,n}^e) | \mu_\eta, \Sigma_\eta / \phi_n)}{p_n(h(y_{i,n}^e) | \mu_\eta, \Sigma_\eta / \phi_n)} \times \left| \frac{\det(\mathcal{J}_{h^{-1}}(y_{i,n}^e))}{\det(\mathcal{J}_{h^{-1}}(\hat{y}_{i,n}^e))} \right| \quad (2)$$

where $\mathcal{J}_{h^{-1}}(y)$ denotes the Jacobian of the inverse transformation of $h(y)$.

3.3.3 Draw $u \sim U(0, 1)$.

Iff $u < \alpha$:

Set $y_{i,n}^e = \hat{y}_{i,n}^e$

3.3.4 Draw the other variables y_i^{-e} from conditional density

$$y_i^{-e} \sim q(y_i^{-e} | y_{i,N_\phi}^e, \mu_{\theta, -e|e}^{(i)}, \Sigma_{\theta, -e|e}^{(i)})$$

Relationship to Entropic Tilting

BASIC IDEA: Reweight draws from a model-based distribution $F(y)$ to adapt them to a target distribution $F'(y)$. $F'(y)$ is closest distribution that satisfies a number of constraints \bar{g} :

$$D(F|F') = \int f'(y) \log\left(\frac{f'(y)}{f(y)}\right) dy \quad \text{s.t.} \quad \int f'(y)g(y)dy = \bar{g} \quad \text{and} \quad \int f'(y)dy = 1$$

- Solution is given by Radon-Nikodym Derivative

$$f'(y) \propto \Lambda f(y) \quad \text{with} \quad \Lambda \propto \exp(\gamma' g(y)) \quad \text{and} \quad \gamma = \left[\int \exp(\gamma' g(y))(g(y) - \bar{g}) dF(y) \right]$$

→ In finite samples Λ does not exist if $D(F|F')$ becomes too large.

- Draws $\{y_i\}_{i=1}^N$ from F , are resampled using the normalized importance weights

$$W(y_i) = \frac{f'(y_i) \exp(\gamma' g(y_i))}{\sum_{i=1}^N f'(y_i) \exp(\gamma' g(y_i))}$$

- While entropic tilting works directly with the moment constraints, our approach takes the moments to build a target density
- Requires additional assumption on the density but allows to move the draws in the mutation step to overcome particle degeneracy.

Application: Transmission of oil-price risks to inflation

- We use our methodology to introduce information from option prices of oil in the density forecast of a small euro area BVAR model and in the NAWM (paper).
- Based on Breeden and Litzenberger (1978), it is possible to construct probability densities and implied moments of underlying assets from derivative prices at the date of expiry.
- Information content of option implied moments is also documented in Bauer and Chernov (2021) or Day and Lewis (1992).
- Option-implied probability densities are provided for variables such as exchange rates, interest rates or oil prices and are regularly published by the ECB, the BoE or the Federal Reserve.

Application: Transmission of oil-price risks to inflation

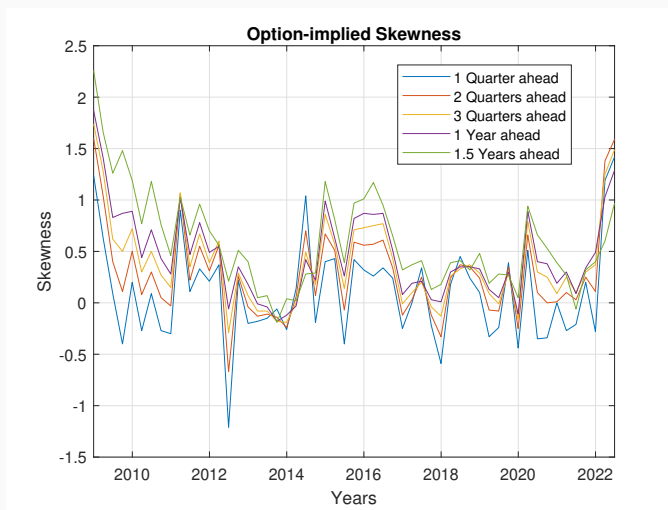


Figure 2: Skewness of the option implied probability density functions of oil prices using quarterly data from 2008 to 2022 for different forecast horizons obtained from the ECB's Statistical Data Warehouse. The probability density of the future oil price exhibits large fluctuations in the evolution of skewness for all horizons, and over the full sample.

Multivariate Skew-T

We model the option implied densities as a **multivariate Skew-T distribution** introduced by Azzalini and Capitanio (2003)

Multivariate Skew-T Distribution I

$y \in \mathbb{R}^p$, follows a multivariate Skew-T with density $\tau(y|\xi, \Omega, \lambda, \nu)$ where $\xi \in \mathbb{R}^p$ determines the location, Ω is a $p \times p$ Covariance matrix, $\lambda \in \mathbb{R}^p$ is the shape parameter and $\nu \in \mathbb{N}$ are the degrees of freedom.

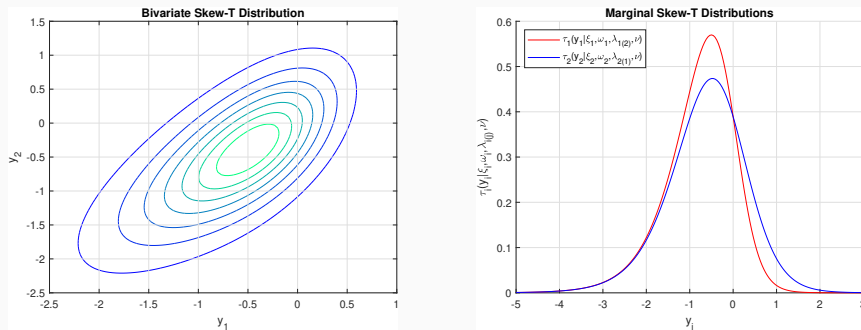


Figure 3: Simple two-dimensional example with $\xi_i = 0$ and $\omega_i = 1$ for $i = \{1, 2\}$. The correlation coefficient is $\rho = 0.8$ and the shape parameters are given by $\lambda_1 = -2$ and $\lambda_2 = 0$. This gives $\lambda_{1(2)} = -2$ and $\lambda_{2(1)} = -1.024$. Marginals

→ **Positive correlation between y_1 and y_2 results in negative skewness of both marginals.**

Skew-T: Theoretical Moments

Based on Azzalini and Capitanio (2003), the first three theoretical moments of the univariate skew-T distribution are given by the following set of equations:

Moment Matching

$$\gamma_i^{oil} = \kappa_{i(j)} \left[\frac{\nu (3 - \delta_{i(j)}^2)}{\nu - 3} - \frac{3\nu}{\nu - 2} + 2\kappa_{i(j)}^2 \right] \left[\frac{\nu}{\nu - 2} - \kappa_{i(j)}^2 \right]$$

$$\kappa_{i(j)} = \frac{\sqrt{\nu} \Gamma(\frac{1}{2}(\nu - 1))}{\sqrt{\pi} \Gamma(\frac{1}{2}\nu)} \delta_{i(j)}$$

$$\lambda_{i(j)} = \frac{\delta_{i(j)}}{\sqrt{1 - \delta_{i(j)}^2}}$$

$$\sigma_i^{oil} = \omega_i \sqrt{\left[\frac{\nu}{\nu - 2} - \kappa_{i(j)}^2 \right]}$$

$$\mu_i^{oil} = \xi_i + \omega_i \kappa_{i(j)}$$

In a second step we recover the vectors of the joint lambdas λ_i from the observed marginals $\lambda_{i(j)}$

[Details](#)

Introducing external information in a BVAR model

We obtain $q_\theta(y)$ from the forecasting density of a reduced form BVAR

$$y_t = \zeta + A_1 y_{t-1} + \dots + A_s y_{t-s} + u_t \quad \text{with} \quad u_t \sim \mathcal{N}(0, \Sigma)$$

- Variables: logs of the price of oil, real GDP, prices including and excluding energy, US/Dollar exchange rate, employment as well as the long and short term interest rates.
- Priors: Hierarchical approach of Giannone, Lenza, and Primiceri (2015) and Covid correction of Lenza and Primiceri (2020)

Introducing external information in a BVAR model

We obtain $q_\theta(y)$ from the forecasting density of a reduced form BVAR

$$y_t = \zeta + A_1 y_{t-1} + \dots + A_s y_{t-s} + u_t \quad \text{with} \quad u_t \sim \mathcal{N}(0, \Sigma)$$

- Variables: logs of the price of oil, real GDP, prices including and excluding energy, US/Dollar exchange rate, employment as well as the long and short term interest rates.
- Priors: Hierarchical approach of Giannone, Lenza, and Primiceri (2015) and Covid correction of Lenza and Primiceri (2020)

The **proposal density** for the i^{th} draw $y_i = [y'_{T+1}, \dots, y'_{T+p}]$ is then given by the multivariate forecasting distribution (Proposal) of the model and takes the form

$$q_\theta(y_i) = \varphi(y | \mu^{(i)}, \Sigma^{(i)})$$

The **target density** is the multivariate skew-T fitted to the option implied moments.

$$p_\eta(y_i^{oil}) = \tau(y^{oil} | \hat{\xi}, \hat{\Omega}, \hat{\lambda}, \nu)$$

Results from the BVAR

We first use our method to investigate the effect of the strong increase in oil prices on inflation due to the begin of the war in Ukraine in the first quarter of 2022.

- We estimate the BVAR model using data up to 2022Q1 and introduce the information of the option-implied densities at the 4th of March.
- Table 1 displays the option-implied moments for all forecasting horizons p

p	Mean	SD	Skewness
1	110.2	38.88	1.8
2	103.16	40.64	1.56
3	98.92	40.59	1.28
4	95.3	41.1	1.14
5	92.13	41.88	1.09
6	89.75	41.99	1.01

Table 1: Option-implied moments Fitted

Given the debate about the pass-through of high energy prices to inflation we are particularly interested in the effect on both, inflation and and core inflation.

Inflation Risks

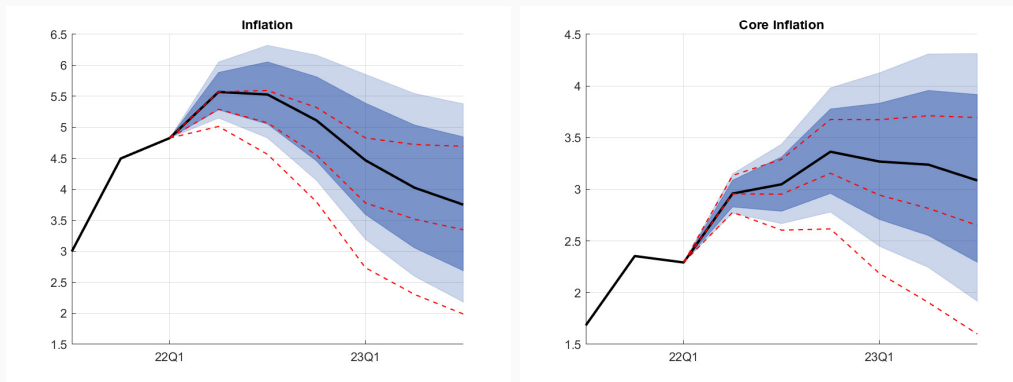


Figure 4: Forecasting densities for inflation and core inflation. The shaded areas show the 16, 25, 75 and 84 percent quantiles of the annualized inflation rate together with the median (solid black line). The dotted red lines show the 16 and 84 percent quantiles of the original distribution.

- Introducing the information of the options results in an upward shift of the full distribution
- Positive skewness in the distribution of the oil prices results in upside risks to inflation marginals
- Core inflation rates remain elevated over the forecasting horizon compared to the BVAR

We look at the probabilistic forecasting performance in a real time forecasting exercise to forecast GDP, inflation and core inflation.

- We estimate the same BVAR using data vintages starting in the last quarter of 2013 up until the third quarter of 2021.
- With the onset of the Covid pandemic we again use the method of Lenza and Primiceri (2020).
- We use our algorithm to impose the option-implied distribution at the end of the quarter to the forecasting density of the oil-price.
- Evaluation based on the continuous ranked probability score (CRPS)

$$CRPS(F, x) = \int_{-\infty}^{\infty} (F(y) - \mathbf{1}(y - x))^2 dy$$

Results: Forecasting Exercise

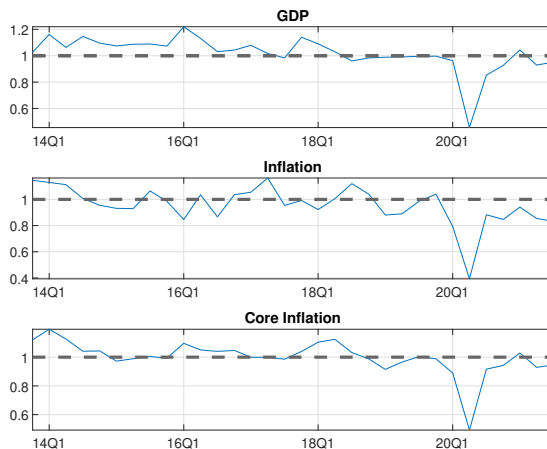


Figure 5: The plot shows the ratios of the average CRPS over all horizons for the symmetric density forecasts under Q_θ and the skew-T forecasts under P_η . Values lower than 1 indicate better probabilistic forecasts under P_η . Including additional information on the distribution from the options does not increase predictive accuracy in moderate periods but strongly increases the probabilistic forecasts accuracy in times of economic turmoil.

GDP

Inflation

Core Inflation

We develop a methodology to adapt draws from a model-based distribution to a target distribution that is specified based on external information:

- The algorithm uses a modified version of the tempered importance sampling method of Herbst and Schorfheide (2014) to allow applications where the proposed draws are far away from the target density

We develop a methodology to adapt draws from a model-based distribution to a target distribution that is specified based on external information:

- The algorithm uses a modified version of the tempered importance sampling method of Herbst and Schorfheide (2014) to allow applications where the proposed draws are far away from the target density

Illustration of our algorithm by introducing external information about the distribution of future oil prices obtained from derivative prices into the forecasting densities of a BVAR and the NAWM

- Adapting the forecasting distributions of the BVAR to the option-implied densities results in upside risks to inflation and core inflation.
- Median forecasts of core inflation remain persistently higher over the forecasting horizon
- Real-time forecasting exercise indicates that introducing information about the marginal distribution of oil prices improves forecasts for GDP and inflation measures during the Covid pandemic compared to symmetric forecasting distributions.

Additional Material

We recover the parameters of the multivariate skew-T distribution in two steps:

1. We match the option implied mean μ_i^{oil} , standard deviation σ_i^{oil} and skewness γ_i^{oil} of the **marginal forecast densities** of oil prices to the moments of the skew-T distribution derived in Azzalini and Capitanio (2003) for each forecasting horizon $i = 1, \dots, P$
2. We use the results of Arellano-Valle and Genton (2010) to solve for the shape parameters of the **joint forecast density** ($\hat{\lambda}$) by solving the following system of equations

Joint Lambdas

$$\lambda_i = \lambda_{i(-j)} \sqrt{1 + \lambda'_{-i} \tilde{\Omega}_{ii|-i} \lambda_{-i}} - \bar{\Omega}_{ii}^{-1} \bar{\Omega}'_{-ii} \lambda_{-i} \quad \forall i = 1, \dots, P$$

Multivariate Skew-T: Density

Based on Proposition 3 in Arellano-Valle and Genton (2010) the multivariate skew T distribution is **closed under marginalization**.

Marginal Distributions

For a partition $y = (y_1, y_2)$, with dimensions p_1 and p_2 and parameters (ξ, Ω, λ) , the marginal densities of y_i with $i = 1, 2$ are given by

$$\tau(y|\xi_i, \Omega_{ii}, \lambda_{i(j)}, \nu)$$

with

$$\lambda_{i(j)} = \frac{\lambda_i + \bar{\Omega}_{ii}^{-1} \bar{\Omega}_{ij} \lambda_j}{\sqrt{1 + \lambda_j' \tilde{\Omega}_{ii|j} \lambda_j}} \quad \text{and} \quad \tilde{\Omega}_{ii|j} = \bar{\Omega}_{jj} - \bar{\Omega}_{ji} \bar{\Omega}_{ii}^{-1} \bar{\Omega}_{ij}$$

- shape parameter of the marginal distribution is a weighted sum of the vector of individual shape parameters
- weights depend on the correlation between y_i and y_j
- $\lambda_i = 0$ does not imply that the marginal distribution of y_i is symmetric.

Deriving the Forecasting Distribution

- BVAR-Model in Companion Form:

$$y_t = c_i + \Phi_i y_{t-1} + G_i \varepsilon_t$$

- Stacking the realizations over the full forecasting horizon in a vector y_i yields

$$\begin{bmatrix} y_{T+1} \\ y_{T+2} \\ \vdots \\ y_{T+h} \end{bmatrix} = \begin{bmatrix} \tilde{c}_{i,T+1} \\ \tilde{c}_{i,T+2} \\ \vdots \\ \tilde{c}_{i,T+h} \end{bmatrix} + \begin{bmatrix} G_i & 0 & 0 & 0 \\ \Phi_i G_i & G_i & 0 & 0 \\ \vdots & \vdots & \ddots & 0 \\ \Phi_i^{h-1} G_i & \Phi_i^{h-2} G_i & \cdots & G_i \end{bmatrix} \begin{bmatrix} \varepsilon_{T+1} \\ \varepsilon_{T+2} \\ \vdots \\ \varepsilon_{T+h} \end{bmatrix}$$

with $\tilde{c}_{i,T+h} = \sum_{j=1}^h \Phi_i^{j-1} c_i + \Phi_i^h y_T$.

- Redefining the terms results in the simple expression

$$y_i = \mu_i + G_i \varepsilon$$

- Given the distributional assumption about ε_t it follows that

$$y_i \sim \mathcal{N}(y | \mu^{(i)}, \Sigma^{(i)})$$

with $\Sigma^{(i)} = G_i G_i'$

- We use the the adaptive tempering schedule of Herbst and Schorfheide (2019) to obtain the optimal values for N_{ϕ_N} and ϕ_n .
- We determine the particle approximation using the Inefficiency Ratio $Ineff = \sum_{i=1}^M W_i^2$

Tempering Schedule

In each iteration, we solve

$$\phi_n = \operatorname{argmin} \frac{1}{M} \sum_i \left[\frac{w_{i,n}(\phi)}{\frac{1}{M} \sum_{i=1}^M w_{i,n}(\phi)} \right]^2 - r^*$$

- To obtain a precise approximation of the target density we set $r^* = 1.01$.

Fitted Marginal Densities for oil price

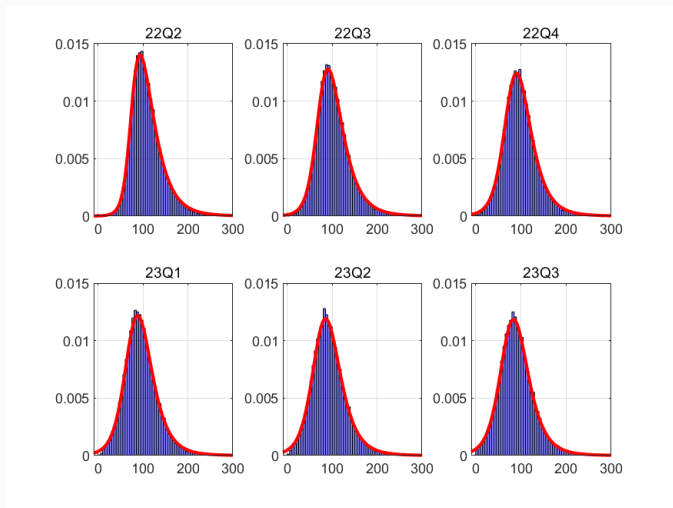


Figure 6: Fitted marginal skew T-densities that from the values in Table 1 together with the histogram of the final particles $\{y_{i,N_\phi}^{Oil}\}_{i=1}^N$. The algorithm required a number of $N_\phi = 40$ tempering steps. Draws are very well adapted to the target distribution and the positive skewness and increasing volatility is clearly visible. [Return](#)

Appendix IV: Marginal Densities

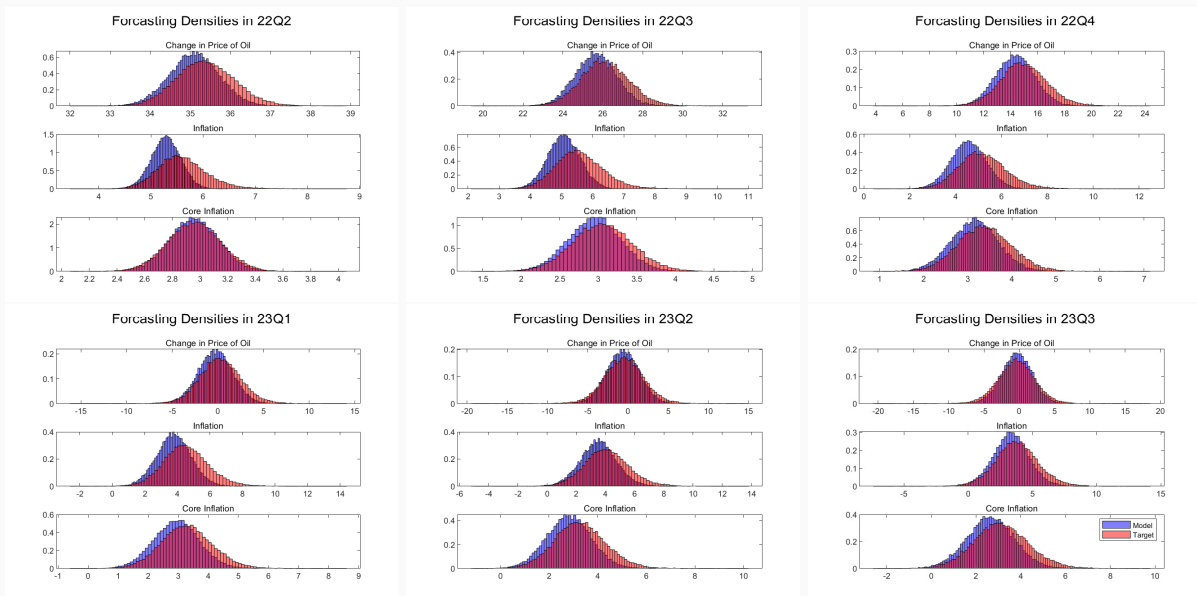


Figure 7: Marginal conditional forecasting densities of annual GDP growth, inflation and core inflation.

Results: Forecasting GDP

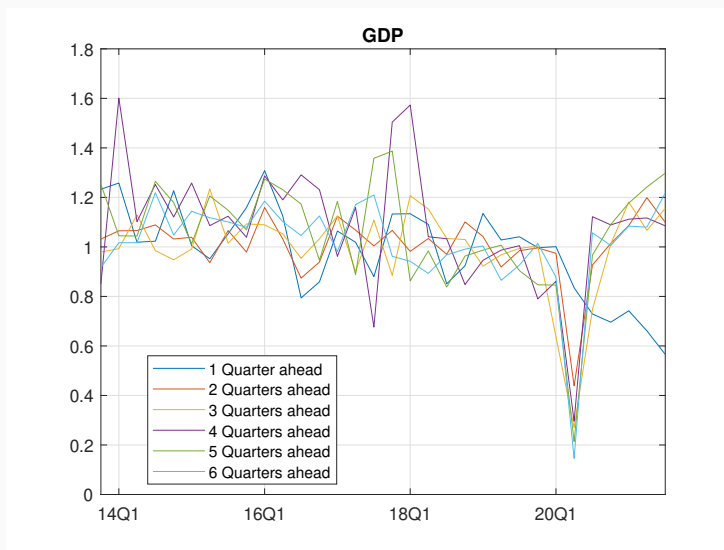


Figure 8: The plot shows the ratios of the CRPS for the symmetric density forecasts under Q_θ and the skew-T forecasts under P_η for GDP and for all forecasting horizons. Values lower than 1 indicate better probabilistic forecasts under P_η [Return](#)

Results: Forecasting Inflation

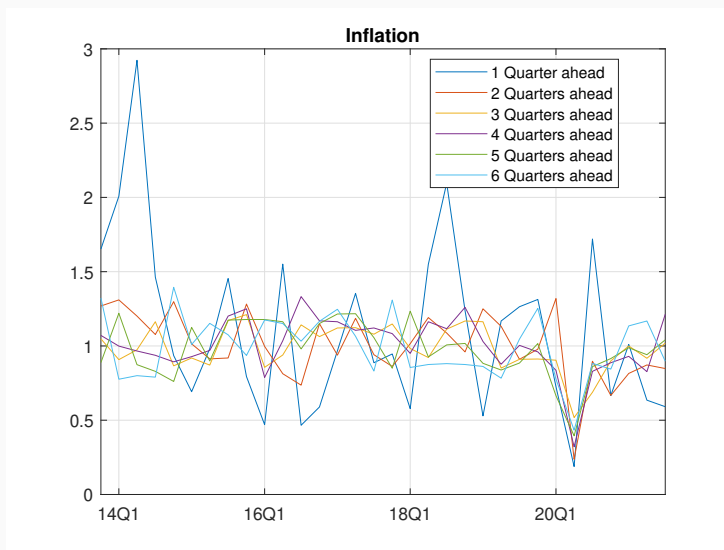


Figure 9: The plot shows the ratios of the CRPS for the symmetric density forecasts under Q_θ and the skew-T forecasts under P_η for inflation and for all forecasting horizons. Values lower than 1 indicate better probabilistic forecasts under P_η [Return](#)

Results: Forecasting Core Inflation



Figure 10: The plot shows the ratios of the CRPS for the symmetric density forecasts under Q_θ and the skew-T forecasts under P_η for core inflation and for all forecasting horizons. Values lower than 1 indicate better probabilistic forecasts under P_η [Return](#)

Our methodology requires knowledge of the conditional distribution of the proposal distribution. Second, we use a tempering method that requires a target density with a parameter to control the scale of the distribution.

- If the conditional distribution is not available, it is possible to approximate the proposal density $q(x)$ with a Gaussian mixture density that can be used to sample from in step 4.

Conditional Gaussian Mixture

$$q(x_1|x_2) = \sum_{k=1}^K \left[\frac{\pi_k \varphi(x_2|\mu_{k,2}, \Sigma_{k,22})}{\sum_{l=1}^L \pi_l \varphi(x_2|\mu_{l,2}, \Sigma_{l,22})} \right] \varphi(x_1|x_2, \mu_{k,1|2}, \Sigma_{k,1|2})$$

- We propose another way to define the bridge distributions as given in Neal (2001)

Alternative Bridge Distribution

$$p_n(y_i) = p_\eta(y_i)^{\phi_n} q_\theta(y_i)^{(1-\phi_n)}$$

Results from the NAWM

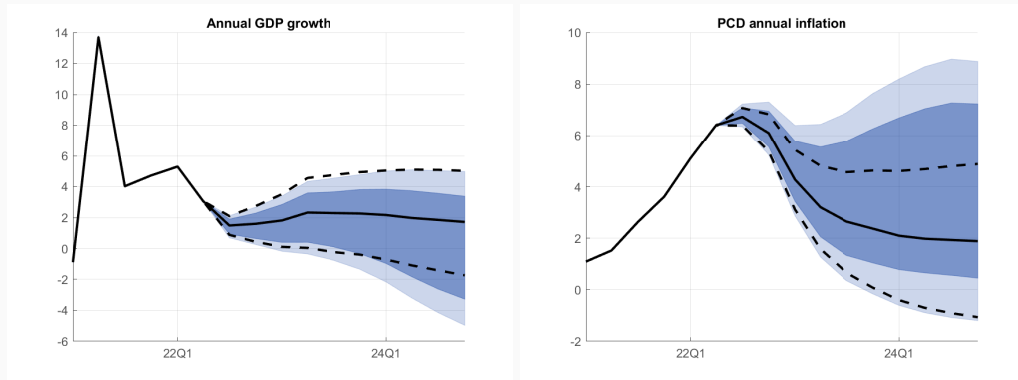


Figure 11: Forecasting densities for the year-on-year growth rate of GDP and PCD inflation. The shaded areas show the 16, 25, 75 and 84 percent quantiles of the tilted forecasting distributions. The dotted black lines show the 16 and 84 percent quantiles of the original distribution. The solid black lines show the BMPE values. The results are obtained using a modified version of the NAWM II model with an enhanced transmission channel of oil prices.

- Positive skewness in the distribution of the oil prices results in upside risks to inflation
- Positive skewness in the distribution of the oil prices results in downside risks to GDP growth