# Web Appendix to

# The risk management approach

# to macro-prudential policy*

*Sulkhan Chavleishvili,*[(a)] *Robert F. Engle,*[(b)] *Stephan Fahr,*[(c)]

*Manfred Kremer,*[(c)] *Frederik Lund-Thomsen,*[(c)] *Simone Manganelli,*[(c)]

*Bernd Schwaab*[(c)]

[(a)] Aarhus University, [(b)] NYU Stern School of Business, [(c)] European Central Bank

*E-mail contacts: sulkhan.chavleishvili@econ.au.dk; rengle@stern.nyu.edu; stephan.fahr@ecb.europa.eu; manfred.kremer@ecb.europa.eu; frederik_ole.lund-thomsen@ecb.europa.eu; simone.manganelli@ecb.europa.eu; bernd.schwaab@ecb.europa.eu (corresponding author). The views expressed in this paper are those of the authors and they do not necessarily reflect the views or policies of the European Central Bank.

# A Bayesian estimation

## A.1 Gibbs sampler

We obtain posterior inference relying predominantly on established methods for Bayesian quantile regressions (see e.g. Yu and Moyeed, 2001, Kozumi and Kobayashi, 2011, Khare and Hobert, 2012, and Korobilis, 2017). This still allows us, as in Chavleishvili and Manganelli (2023), to estimate the SQVAR equation by equation. We start by considering the endogenous variable $x_i$ at quantile $\gamma$,

$$x_{it} = w_{it}'\beta_i(\gamma) + \varepsilon_{it}^{\gamma}, \tag{A.1}$$

where $x_{it}$ is a scalar, $i = 1, \ldots, n$, $t = 1, \ldots, T$, $w_{it}$ is a vector of regressors (in our case, contemporaneous values and lags of the endogenous and exogenous variables, a constant, as well as dummy variables), and $\beta_i(\gamma)$ is a vector of quantile-specific coefficients. The error term $\varepsilon_{it}^{\gamma}$ is assumed to have an asymmetric Laplace distribution of the form

$$f\left(\varepsilon_{it}^{\gamma}\right) = \gamma\left(1-\gamma\right)\left[\sigma_i(\gamma)\right]^{-1} e^{-\rho_\gamma\left(\varepsilon_{it}^{\gamma}\right)}, \tag{A.2}$$

where $\sigma_i(\gamma)$ is a scale parameter, and $\rho_\gamma\left(\varepsilon\right) \equiv \varepsilon(\gamma - I(\varepsilon < 0))$ is the standard quantile regression check function; see also Koenker and Bassett Jr (1978) and Engle and Manganelli (2004). It is clear from (A.2) that minimizing the usual quantile regression objective function, as e.g. defined in Koenker and Bassett Jr (1978), is equivalent to maximizing a corresponding asymmetric Laplace log-likelihood; see e.g. Yu and Moyeed (2001).

An asymmetric Laplace random variable can be represented as a mixture of a standard normal and an exponential random variable (Yu and Moyeed, 2001, Kozumi and Kobayashi, 2011). As a result, (A.1) can be restated as

$$x_{it} = w_{it}'\beta_i(\gamma) + \theta(\gamma)\nu_{it}(\gamma) + \tilde{\tau}(\gamma)\sqrt{\sigma_i(\gamma)\nu_{it}(\gamma)}u_{it}^{\gamma}, \tag{A.3}$$

where $\tilde{\tau}^2(\gamma) = \frac{2}{\gamma(1-\gamma)}$, $\theta(\gamma) = \frac{1-2\gamma}{\gamma(1-\gamma)}$, $u_{it}^{\gamma} \sim N(0,1)$, $\nu_{it}(\gamma) \sim \mathcal{E}(\sigma_i(\gamma))$, and where $\mathcal{E}(\tilde{e})$ denotes an exponential distribution with mean $\tilde{e}$.

Since this section exclusively considers univariate regressions at a single quantile $\gamma$, we drop the quantile notation for convenience and leave the dependence implied. The mixture representation in (A.3) allows us to draw from the conditional posterior distributions using an appropriate set of prior distributions. We follow Kozumi and Kobayashi (2011) and choose the prior distributions

$$\beta_i \sim N\left(\underline{\mu}_{\beta,i}, \lambda_i \underline{\Sigma}_{\beta,i}\right), \qquad \sigma_i \sim IG\left(\underline{\alpha}_{\sigma,i}, \underline{\zeta}_{\sigma,i}\right), \tag{A.4}$$

where $IG(\alpha, \zeta)$ denotes an inverse-gamma distribution with shape and scale parameters $\alpha$ and $\zeta$. The addition of the scalar $\lambda_i \geq 0$ to the prior variance of $\beta_i$ lets us control the looseness/tightness of the respective prior. A lower value of $\lambda_i$ implies a tight, informative prior, while a high value of $\lambda_i$ implies a loose, uninformative prior. We give $\lambda_i$ its own prior density and let $\lambda_i \sim IG(\underline{\alpha}_{\lambda}, \underline{\zeta}_{\lambda})$.

With the above priors in place, and continuing to rely on Kozumi and Kobayashi (2011) and Khare and Hobert (2012), we end up with the following four-step Gibbs sampler for all $nq$ equations, dropping subscript $i$ for convenience. Steps 1 – 3 are identical to those in Khare and Hobert (2012). Step is 4 is new but straightforward; Section A.3 provides a derivation.

1. Draw $\sigma|x, w, \ldots \sim IG\left(\bar{\alpha}_{\sigma}, \bar{\zeta}_{\sigma}\right)$,
   where $\bar{\alpha}_{\sigma} = \underline{\alpha}_{\sigma} + \frac{3}{2}T$, and $\bar{\zeta}_{\sigma} = \zeta_{\sigma} + \sum_{t=1}^{T} \frac{(x_t - w_t'\beta - \theta\nu_t)^2}{2\tilde{\tau}^2\nu_t} + \sum_{t=1}^{T} \nu_t$.

2. Draw $\beta|x, w, \ldots \sim N\left(\bar{\mu}_{\beta}, \bar{\Sigma}_{\beta}\right)$,
   where $\bar{\Sigma}_{\beta}^{-1} = \sum^{T} \frac{w_t w_t'}{\tilde{\tau}^2\sigma\nu_t} + \lambda^{-1}\underline{\Sigma}_{\beta}^{-1}$, and $\bar{\mu}_{\beta} = \bar{\Sigma}_{\beta}\left[\sum^{T} \frac{w_t(x_t - \theta\nu_t)}{\tilde{\tau}^2\sigma\nu_t} + \lambda^{-1}\underline{\Sigma}_{\beta}^{-1}\underline{\mu}_{\beta}\right]$.

3. Draw $\nu_t^{-1}|x_t, w_t, \ldots \sim \mathcal{IGN}(\bar{\kappa}_1, \bar{\kappa}_2)$, where $\mathcal{IGN}$ denotes the inverse Gaussian distribution, with $\bar{\kappa}_1 = \frac{\sqrt{\theta^2 + 2\tilde{\tau}^2}}{|x_t - w_t'\beta|}$, and $\bar{\kappa}_2 = \frac{\theta^2 + 2\tilde{\tau}^2}{\sigma\tilde{\tau}^2}$.

4. Draw $\lambda|\beta, \ldots \sim IG\left(\bar{\alpha}_{\lambda}, \bar{\zeta}_{\lambda}\right)$,
   where $\bar{\alpha}_{\lambda} = \underline{\alpha}_{\lambda} + \frac{\tilde{k}}{2}$, $\bar{\zeta}_{\lambda} = \zeta_{\lambda} + \frac{1}{2}\left(\beta - \underline{\mu}_{\beta}\right)' \underline{\Sigma}_{\beta}^{-1}\left(\beta - \underline{\mu}_{\beta}\right)$, and $\tilde{k}$ is the dimension of $\beta$ (the number of regressors).

To conclude, we note that the Gibbs sampler of Kozumi and Kobayashi (2011) and Khare and Hobert (2012), steps 1–3 above, comes with theoretical guarantees: draws from the sampler converge to the intractable true posterior, and do so at a geometric rate. To prove this result, Khare and Hobert (2012) make no assumptions regarding the dimensions of the data (so that the result continues to hold even if $\tilde{k}$ were large relative to $T$).

## A.2 Specification of prior densities

As explained in the main text, we first estimate the SQVAR model's parameters for U.S. data. We do so using the same model specification (i.e., variables, deterministic terms, number of lags and zero restrictions). Next, we let these estimates inform the euro area parameters' priors. We are thus taking advantage of decades of data available for the U.S. between 1973Q1 and 2022Q4. Both the U.S. and euro area economies are advanced, market-based economies with institutional similarities along many dimensions.

### A.2.1 Priors for U.S. data

We need to specify the prior parameters $\underline{\mu}_{\beta,i}^{US}$ $\underline{\Sigma}_i^{US}$, $\underline{\alpha}_{\sigma,i}^{US}$, $\underline{\zeta}_{\sigma,i}^{US}$, $\underline{\alpha}_{\lambda,i}^{US}$, and $\underline{\zeta}_{\lambda,i}^{US}$. For U.S. data, we keep the prior parameters homogeneous across quantiles. We have therefore dropped a potential dependence on $\gamma$.

We employ a Minnesota prior for the vector of coefficients $\beta_i^{US}$ (see, e.g., Litterman, 1986, Luetkepohl, 2005, p. 225, Giannone et al., 2015). This means that the prior density is Gaussian, and pointing to a persistent autoregressive process of order one. Specifically, we set the coefficient in $\beta_i^{US}$ referring to variable $i$'s own first lag equal to either 0.9 or 1, depending on the variable. For the CISS, the financial cycle, and the real GDP growth rate, the own-lag coefficient is set to 0.9; for CPI inflation and the Federal Funds Rate the coefficient is set to one. All other elements of $\beta_i^{US}$ are given a prior mean of zero.

We further specify the Minnesota prior's covariance matrix $\underline{\Sigma}_{\beta,i}^{US}$ as diagonal, with $\underline{\sigma}_{\beta,ij,l}^{US}$ the element corresponding to the $l$th lag of the $j$th endogenous variable in equation $i$. The diagonal

elements are given by

$$\begin{aligned}
\underline{\sigma}^{US}_{\beta,ij,l} &= \phi_0^2 && \text{if } i=j, l=0 \\
\underline{\sigma}^{US}_{\beta,ij,l} &= \left(\frac{\phi_0}{l^{\phi_3}}\right)^2 && \text{if } i=j, l>0 \\
\underline{\sigma}^{US}_{\beta,ij,l} &= \left(\phi_0\phi_1\frac{\hat{\sigma}_i}{\hat{\sigma}_j}\right)^2 && \text{if } i\neq j, l=0 \\
\underline{\sigma}^{US}_{\beta,ij,l} &= \left(\frac{\phi_0\phi_1}{l^{\phi_3}}\frac{\hat{\sigma}_i}{\hat{\sigma}_j}\right)^2 && \text{if } i\neq j, l>0 \\
\underline{\sigma}^{US}_{\beta} &= (\phi_0\phi_2)^2 && \text{otherwise,}
\end{aligned}$$

where $\phi_0$ is a general tightness parameter, $\phi_1$ a tightness parameter on endogenous variables other than variable $i$, $\phi_2$ a tightness parameter on exogenous variables and deterministic terms, and $\phi_3$ a tightness parameter controlling the importance of lags of endogenous variables. We choose $\phi_0 = 0.2$, $\phi_1 = 0.5$, $\phi_2 = 10^5$, and $\phi_3 = 1$. These are common choices in the literature (see, e.g., Litterman, 1986). Parameter $\hat{\sigma}_i$ denotes the standard error of the residuals from a univariate quantile autoregression on endogenous variable $i$ at the median. Finally, we set $\underline{\alpha}^{US}_{\sigma,i} = \underline{\zeta}^{US}_{\sigma,i} = 0.01$, corresponding to a non-informative prior for the scale parameter $\sigma$. We also set $\underline{\alpha}^{US}_{\lambda,i} = 3$ and $\underline{\zeta}^{US}_{\lambda,i} = 6$, implying a prior mean and a prior standard deviation of 3 for $\lambda_i$.

### A.2.2 Priors for euro area data

We set the euro area prior parameters for $\beta^{EA}_i(\gamma)$ to match the corresponding posterior moments obtained from U.S. data. Specifically,

$$\begin{aligned}
\underline{\mu}^{EA}_{\beta,i}(\gamma) &= \frac{1}{N^S}\sum_{s=1}^{N^S}\hat{\beta}^{US}_{i,s}(\gamma) \\
\underline{\Sigma}^{EA}_{\beta,i}(\gamma) &= \frac{1}{N^S}\sum_{s=1}^{N^S}\left(\hat{\beta}^{US}_{i,s}(\gamma) - \underline{\mu}^{EA}_{\beta,i}(\gamma)\right)^2,
\end{aligned}$$

where $\hat{\beta}^{US}_{i,s}$ are the $N^S$ posterior draws of $\beta^{US}_i$. All other euro area hyperparameters are given values identical to their U.S. counterparts. Specifically, $\underline{\alpha}^{EA}_{\sigma,i} = \underline{\zeta}^{EA}_{\sigma,i} = 0.01$, $\underline{\alpha}^{EA}_{\lambda,i} = 3$, and $\underline{\zeta}^{EA}_{\lambda,i} = 6$.

## A.3 Posterior density for $\lambda$

This section derives the fourth step of the Gibbs sampler in Section A.1. This step is introduced to estimate a tightness parameter that determines the weight given to the U.S. parameter estimates in forming the priors for estimating the euro area model. Allowing the data to help decide on the tightness of the prior density across quantiles and variables appears particularly appropriate in our case, as we have no reliable prior knowledge of how informative the Minnesota prior is for the parameters of the U.S. model at any given quantile, nor how informative the U.S. posterior density is for the parameters of the euro area model. Bayesian updating of the tightness parameter through a hyper-prior is a common solution, for example, in the literature on shrinkage in Bayesian regression models (see e.g. Huber and Feldkircher, 2019 and Korobilis and Pettenuzzo, 2019). In addition, such updating is common in the analysis of sequential medical trials (see, e.g., Ibrahim et al., 2015 and Ibrahim and Chen, 2000).

To arrive at the conditional posterior distribution for $\lambda$, again dropping subscripts $\gamma$ and $i$ for clarity, we start by writing out the kernel of the joint posterior distribution, $\mathcal{P}(\cdot)$, with the terms relating to $\lambda$,

$$\mathcal{P}(\lambda|\cdot) \propto |\lambda\underline{\Sigma}_\beta|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\left(\beta-\underline{\mu}_\beta\right)'\lambda^{-1}\underline{\Sigma}_\beta^{-1}\left(\beta-\underline{\mu}_\beta\right)\right\} \times \lambda^{-\underline{\alpha}_\lambda-1}\exp\left\{-\underline{\zeta}_\lambda\lambda^{-1}\right\},$$

where the first term comes from the normally distributed prior for $\beta$ and the second term from the inverse-gamma prior distribution for $\lambda$. Rewriting the above expression, we obtain

$$\begin{aligned}\mathcal{P}(\lambda|\cdot) &\propto \lambda^{-\underline{\alpha}_\lambda-\frac{\tilde{k}}{2}-1}\exp\left\{-\lambda^{-1}\left(\frac{1}{2}\left(\beta-\underline{\mu}_\beta\right)'\underline{\Sigma}_\beta^{-1}\left(\beta-\underline{\mu}_\beta\right)+\underline{\zeta}_\lambda\right)\right\}\\ &\propto IG\left(\bar{\alpha}_\lambda,\bar{\zeta}_\lambda\right),\end{aligned}$$

where $\bar{\alpha}_\lambda = \underline{\alpha}_\lambda + \frac{\tilde{k}}{2}$ and $\bar{\zeta}_\lambda = \underline{\zeta}_\lambda + \frac{1}{2}\left(\beta-\underline{\mu}_\beta\right)'\underline{\Sigma}_\beta^{-1}\left(\beta-\underline{\mu}_\beta\right)$, and where $\tilde{k}$ denotes the number of regressors. As a result, we can immediately draw $\lambda$ from its posterior distribution conditional on a posterior draw of $\beta$.

# B Technical details

## B.1 Quantile impulse response functions

This section presents the simulation algorithm used to obtain quantile impulse response functions (QIRFs). For a given set of coefficients, we require two sets of conditional distributions to compare: First, we require the conditional forecast distribution of the data in a baseline scenario without an initial shock. Second, we require the conditional forecast distribution of the data in a counterfactual scenario in which a single shock arrives in the first forecasting period. The QIRFs are obtained as the difference between the latter and the former distribution.

### B.1.1 Notation

Throughout this section, we represent the SQVAR(p) as

$$x_t = \omega(\gamma) + A_0(\gamma)x_t + \sum_{j=1}^{p} A_j(\gamma)x_{t-j} + B(\gamma)d_t + \sum_{j=0}^{p} C_j(\gamma)z_{t-j} + \varepsilon_t^{\gamma}, \tag{B.1}$$

for a given $\gamma \equiv [\gamma_1, \ldots, \gamma_n]'$, and recall the identifying assumption $Q_\gamma\left(\varepsilon_t^{\gamma}|\Omega_t\right) = 0_{n\times 1}$. Similarly, each of the $r$ exogenous variables is modeled as a univariate quantile autoregressive (QAR) process,

$$z_{vt} = \mathcal{W}_v(\gamma_v^*) + \sum_{j=1}^{p} \mathcal{A}_{j,v}(\gamma_v^*)z_{v,t-j} + \mathcal{B}_v(\gamma_v^*)d_t + \xi_{vt}^{\gamma_v^*}, \tag{B.2}$$

for $\gamma^* \equiv [\gamma_1^*, \ldots, \gamma_v^*, \ldots, \gamma_r^*]'$. To establish notation for the remainder of this section, we let $H$ denote the forecast horizon, $o$ the forecast origin, $M$ the number of posterior draws required for posterior inference of the QIRFs, and $S$ the number of forward simulations for each posterior draw. We assume that $\mathcal{G} = \{0.05, 0.10, \ldots, 0.90, 0.95\}$ is a sufficiently large set of $q = 19$ quantiles distributed symmetrically around the median. We obtain $N^S = 2,500$ draws of each $\beta_i(\gamma)$, after discarding a burn-in sample of $N^B = 2,500$.

### B.1.2 Algorithm

The algorithm proceeds as follows.

1. **Obtain posterior draws.** Obtain and store $N^S$ posterior draws for all the SQVAR's parameters in (B.1) and (B.2) at all quantiles, see Appendix A.1. We store them in a four-dimensional array of dimensions $[n+r] \times [p(n+r)+r+1+k] \times q \times N^S$.

2. **Choose forecast origin.** Fix the initial conditions for the endogenous variables $x_{o-(p-1):o}$ and exogenous variables $z_{o-(p-1):o}$. We use the unconditional median over the estimation sample for all variables.

3. Set $m = 1$.

   3.1. **Draw parameters.** Let $\hat{\beta}^{(m)}$ be a random draw from the set of $N^S$ posterior draws, and $\hat{\beta}_i^{(m)}(\check{\gamma})$ the subset of parameters corresponding to the posterior draw of $\beta_i$ for some quantile $\check{\gamma}$.

   3.2. Let $\hat{\epsilon}_{o+1}$ be an $n \times 1$ vector of zeros.

   3.3. Set $s = 1$.

      3.3.1. Set $h = 1$.

         3.3.1.1. **Draw quantiles at random.** Obtain $n$ random draws from the uniform distribution $\mathrm{U}(0,1)$ and map them to the corresponding quantiles in $\mathcal{G}$ based on proximity. Stack the mapped quantiles in the $n \times 1$ vector $\varrho$. Define the $r \times 1$ vector $\varrho^*$ in a similar fashion.

         3.3.1.2. **Set up SQVAR system matrices.** Let $f_{A_j}\left(\hat{\beta}_i^{(m)}(\check{\gamma})\right)$ be a mapping from posterior draw $\hat{\beta}_i^{(m)}(\check{\gamma})$ to the $i$th row of $A_j$ in (B.1). Define similar mappings for the remaining matrix coefficients. Using these mappings, stack the variable-

specific posterior draws into matrices of quantile coefficients, such that

$$\tilde{A}_j^{o+h} = \begin{bmatrix} f_{A_j}\left(\hat{\beta}_1^{(m)}(\varrho_1)\right) \\ \vdots \\ f_{A_j}\left(\hat{\beta}_i^{(m)}(\varrho_i)\right) \\ \vdots \\ f_{A_j}\left(\hat{\beta}_n^{(m)}(\varrho_n)\right) \end{bmatrix} \qquad \tilde{\omega}^{o+h} = \begin{bmatrix} f_{\omega}\left(\hat{\beta}_1^{(m)}(\varrho_1)\right) \\ \vdots \\ f_{\omega}\left(\hat{\beta}_i^{(m)}(\varrho_i)\right) \\ \vdots \\ f_{\omega}\left(\hat{\beta}_n^{(m)}(\varrho_n)\right) \end{bmatrix}$$

$$\tilde{C}_j^{o+h} = \begin{bmatrix} f_{C_j}\left(\hat{\beta}_1^{(m)}(\varrho_1)\right) \\ \vdots \\ f_{C_j}\left(\hat{\beta}_i^{(m)}(\varrho_i)\right) \\ \vdots \\ f_{C_j}\left(\hat{\beta}_n^{(m)}(\varrho_n)\right) \end{bmatrix} \qquad \tilde{B}^{o+h} = \begin{bmatrix} f_{B}\left(\hat{\beta}_1^{(m)}(\varrho_1)\right) \\ \vdots \\ f_{B}\left(\hat{\beta}_i^{(m)}(\varrho_i)\right) \\ \vdots \\ f_{B}\left(\hat{\beta}_n^{(m)}(\varrho_n)\right) \end{bmatrix}$$

for $j = 0, \ldots, p$ and $h = 1, \ldots, H$. Define the corresponding matrices for the exogenous variables analogously using $\varrho^*$, resulting in matrices $\tilde{\mathcal{W}}_v^{o+h}$, $\tilde{\mathcal{A}}_{j,v}^{o+h}$ and $\tilde{\mathcal{B}}_v^{o+h}$.

3.3.1.3. **Iterate exogenous variables forward.** Compute the conditional forecast of each of the $r$ exogenous variables, $z_{o+h}^{(s)}$, using (B.2) and the relevant quantile coefficients determined in the previous step as

$$z_{v,o+h}^{(s)} = \tilde{\mathcal{W}}_v^{o+h} + \sum_{j=1}^{P} \tilde{\mathcal{A}}_{j,v}^{o+h} z_{v,o+h-j}^{(s)} + \tilde{\mathcal{B}}_v^{o+h} d_{t+h}.$$

3.3.1.4. **Iterate endogenous variables forward.** Compute the conditional forecast of $x_{o+h}^{(s)}$ using (B.1) and the relevant quantile coefficients determined in step 3.3.1.2 as

$$x_{o+h}^{(s)} = \left(I - \tilde{A}_0^{o+h}\right)^{-1} \begin{bmatrix} \tilde{\omega}^{o+h} + \sum_{j=1}^{P} \tilde{A}_j^{o+h} x_{o+h-j}^{(s)} \\ + \tilde{B}^{o+h} d_{t+h} + \sum_{j=0}^{P} \tilde{C}_j^{o+h} z_{o+h-j}^{(s)} + \hat{\epsilon}_{o+h} \end{bmatrix}$$

3.3.1.5. If $h < H$, set $h = h + 1$ and return to step 3.3.1.1.

3.3.2. If $s < S$, set $s = s + 1$ and return to step 3.3.1.

3.4. **Obtain predicted quantiles.** Let $\check{\mathbf{x}}_{o+h} = \{x_{o+h}^{(s)}\}_{s=1}^{S}$ be the set of $S$ simulated forecasts of $x_{o+h}$. Compute the $\gamma$th conditional quantile forecast of $x_{i,o+h}$ as

$$Q_\gamma\left(x_{i,o+h}|\Omega_o\right) = Q_\gamma\left(\check{\mathbf{x}}_{i,o+h}\right),$$

where $Q_\gamma\left(\cdot|\Omega_o\right)$ is the $\gamma$th quantile conditional on the information set $\Omega_o$, and letting $Q_\gamma\left(\cdot\right)$ denote the empirical quantile function. This concludes the baseline (no shock) forward simulations.

3.5. **Choose the shock of interest.** Let $\mathcal{S}$ be an $n \times 1$ selection vector picking the endogenous variable $i$ to be shocked. (That is, the $i$th element of $\mathcal{S}$ is 1 and 0 otherwise.) Redefine $\hat{\epsilon}_{o+h} = \psi\mathcal{S}$, where $\psi$ is a scalar. A common choice of magnitude $\psi$ is the estimated standard deviation of quantile shocks obtained at the median. Repeat steps 3.3 - 3.4 to obtain the quantile projection conditional on the shock, $Q_\gamma\left(x_{i,o+h}|\Omega_o, \psi\right)$.

3.6. **Compute the $\gamma$th quantile impulse response function** as

$$\iota_{\gamma,i,o+h}^{(m)} = Q_\gamma\left(x_{i,o+h}|\Omega_o, \psi\right) - Q_\gamma\left(x_{i,o+h}|\Omega_o\right)$$

3.7. If $m < M$, set $m = m + 1$ and return to step 3.2.

Posterior inference of the QIRF of the conditional quantile $\gamma$ of variable $i$ in period $o+h$ is obtained from the sequence $\left\{\iota_{\gamma,i,o+h}^{(m)}\right\}_{m=1}^{M}$.

## B.2 Simulation algorithm for downside risk measures

We obtain measures of downside risks to economic growth between $o$ and $o+h$ by reusing the simulation algorithm in Section B.1, but now replace steps 3.4-3.6 with the following steps.

3.4. For each forecast horizon, compute the $o+h$ growth shortfall as

$$\text{GS}_{o,o+h}^{\tau,(m)} = \frac{1}{S}\sum_{s=1}^{S} \tilde{y}_{o+h}^{(s)} \cdot 1\{\tilde{y}_{o+h}^{(s)} < \tau\},$$

where $\tilde{y}_{o+h}^{(s)}$ denotes a simulated value for quarterly real GDP growth at time $o+h$.

3.5. Compute the final time-$o$ downside risk measure as averages across simulations

$$\text{AGS}_{o,o+1:o+h}^{\tau,(m)} = \frac{1}{H}\sum_{h=1}^{H} \text{GS}_{o,o+h}^{\tau,(m)}$$

Posterior inference is obtained from the sequence $\left\{\text{AGS}_{o,o+1:o+h}^{\tau,(m)}\right\}_{m=1}^{M}$. The above steps can similarly be used to compute average growth longrise.

## B.3 Counterfactual scenarios

Rather than moving through the tree of potential future values of $x_{t+h}$ completely at random, as done in Web Appendix B.2, we may sometimes wish to consider only a subset of paths, or even just one path, in isolation. Such paths can also be thought of as a 'counterfactual scenario,' or model-based thought experiment, that conditions on an arbitrary but fixed sequence of future shocks.

To this effect, let $\Gamma^*$ denote a design matrix of size $H \times (n + r)$. Each element of $\Gamma^*$ is either empty or in the interval $(0, 1)$. Any non-empty $h$th element of the $i$th column designates the quantile path to be traveled by variable $i$ in period $h$ in all $S$ simulations. An empty element means that the path is chosen at random. The matrix $\Gamma^*$ can then be readily applied in Step 3.3.1.1 in the simulation algorithm of Section B.1 to obtain density forecasts in a counterfactual scenario.

An illustration of this principle is given in Figure B.1. There, a single variable, $y$, is projected forward for $H = 3$ periods using two estimated quantiles, $\gamma \in \{\text{u}, \text{d}\}$. An unrestricted projection as in Figure B.1a yields a total of $2^3 = 8$ unique paths for $y$ to travel along. Consider now the counterfactual, in which we require $y$ to initially increase, then decrease, and then finally increase again. This maps to the path satisfying $\Gamma^* = \{\text{u}, \text{d}, \text{u}\}'$, highlighted by the red dashed lines in Figure B.1b.

Instead of a fully restricted tree, as illustrated above, one may also consider a partially restricted tree. In a partially restricted tree, the quantile paths are only fixed for some of the periods or variables. In the example above, one may, for instance, require that the chosen quantile from period 1 to 2 is always $\{\text{d}\}$, while leaving the remaining branches unrestricted. In this case, the total number of paths that can be traveled equals four $\left(Q^{\text{u,d,u}}, Q^{\text{u,d,d}}, Q^{\text{d,d,u}}, Q^{\text{d,d,d}}\right)$.

**Figure B.1: Illustration of counterfactual scenario analysis through quantile restrictions**

Notes: Filled blue lines indicate unrestricted paths, red dashed lines restricted paths, and transparent blue lines the paths excluded by the imposed restrictions.

**(a) Uncrestricted tree**

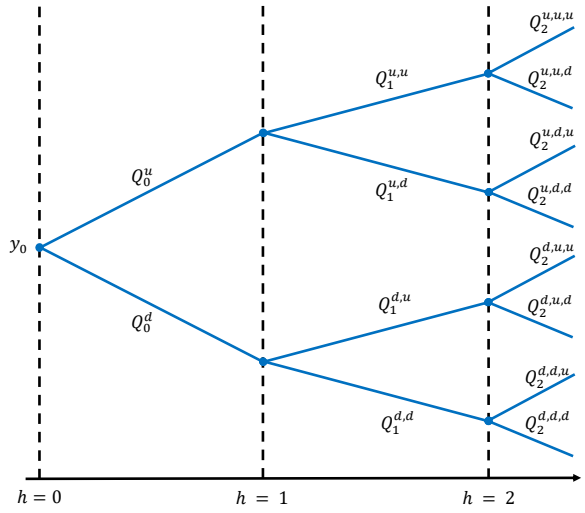


**(b) Fully restricted tree**

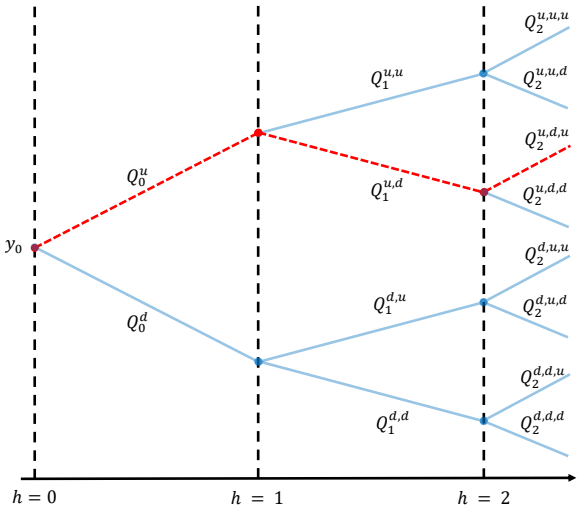

# C Data plots

**Figure C.1: Euro area and U.S. time series**

Notes: Annualised inflation and growth rates are calculated as 400 times the log-difference between quarterly averages.
Sources: European Central Bank, U.S. Bureau of Labor Statistics, U.S. Bureau of Economic Analysis, Board of Governors of the Federal Reserve System, LSEG, Haver Analytics, and authors' calculations.

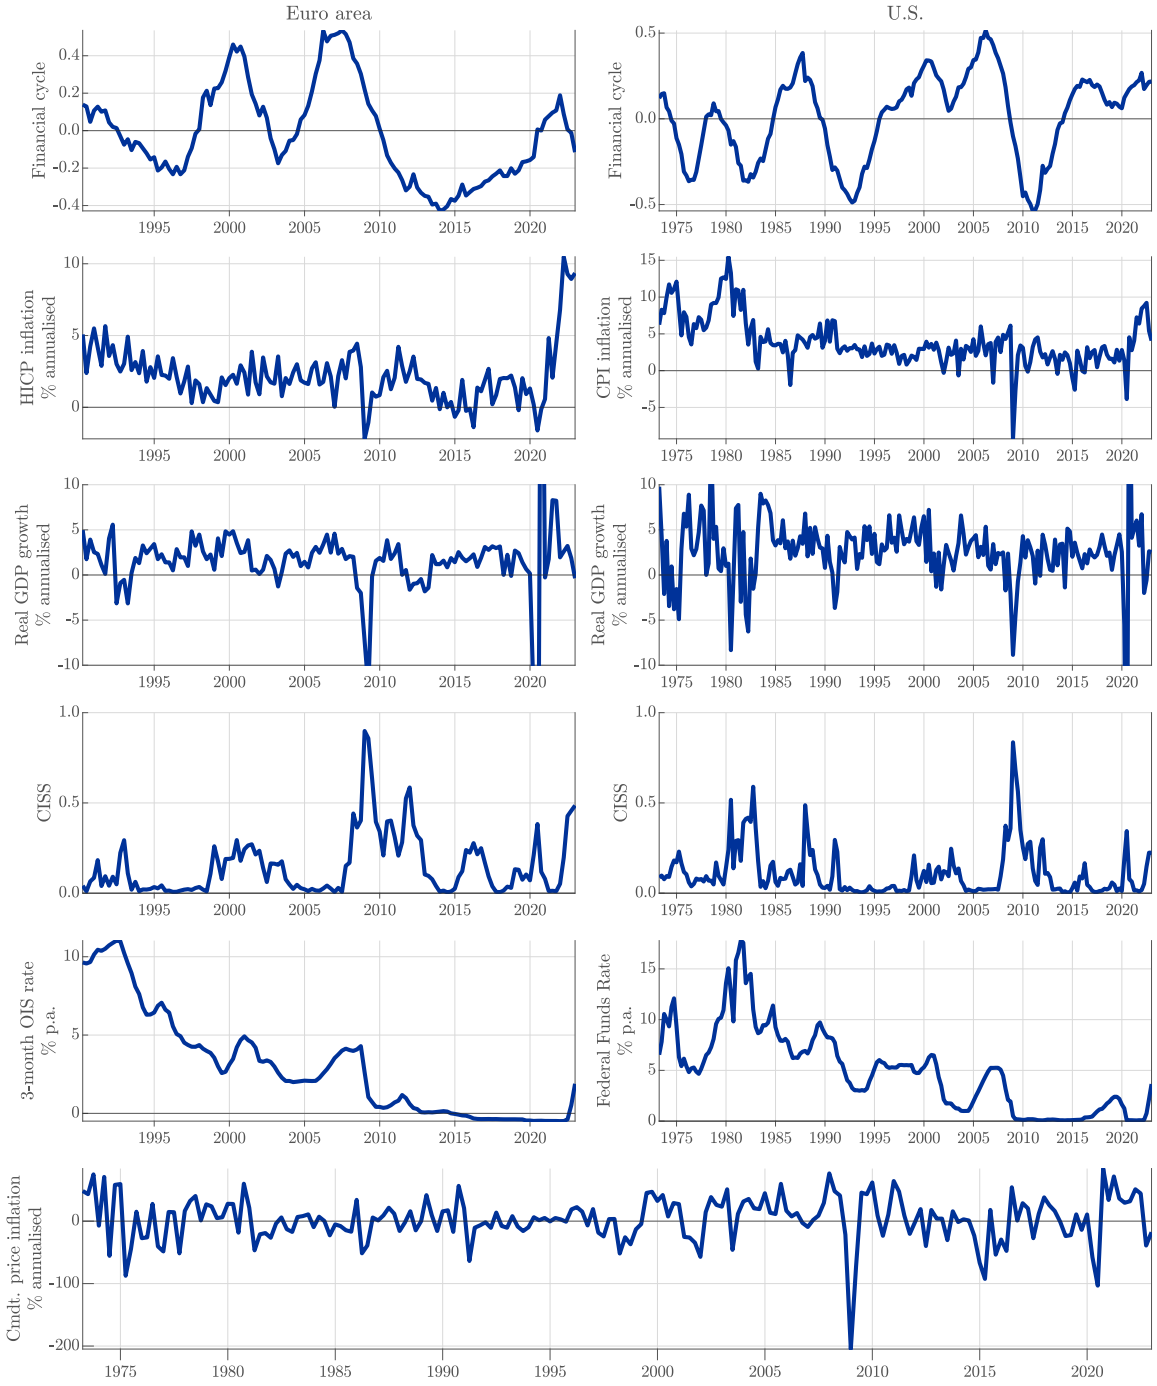

# D Euro area posterior estimates

**Figure D.1: Posterior inference for the euro area, $\omega(\gamma)$**

Posterior mean and 95% credible intervals obtained from 2,500 posterior draws. Least squares estimates (in red) for the conditional mean are provided for comparison.

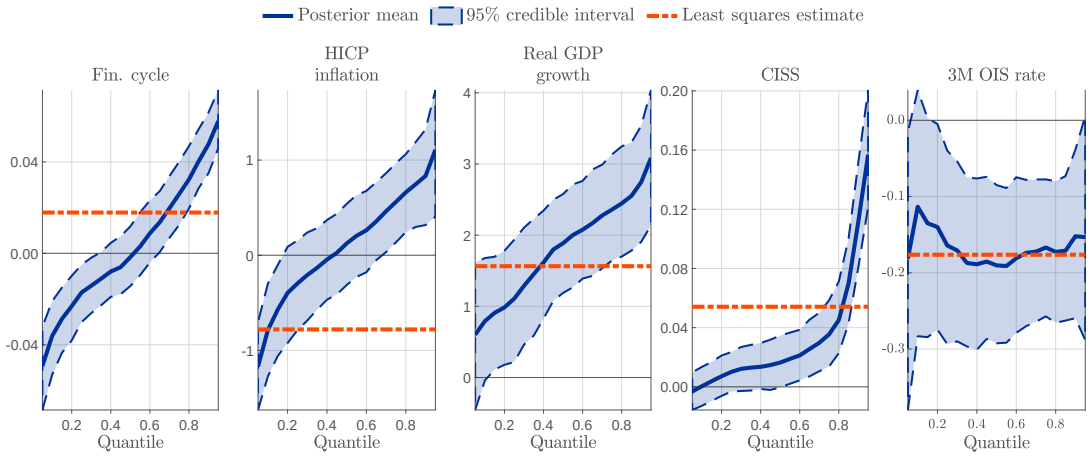


**Figure D.2: Posterior inference for the euro area, $A_0(\gamma)$**

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

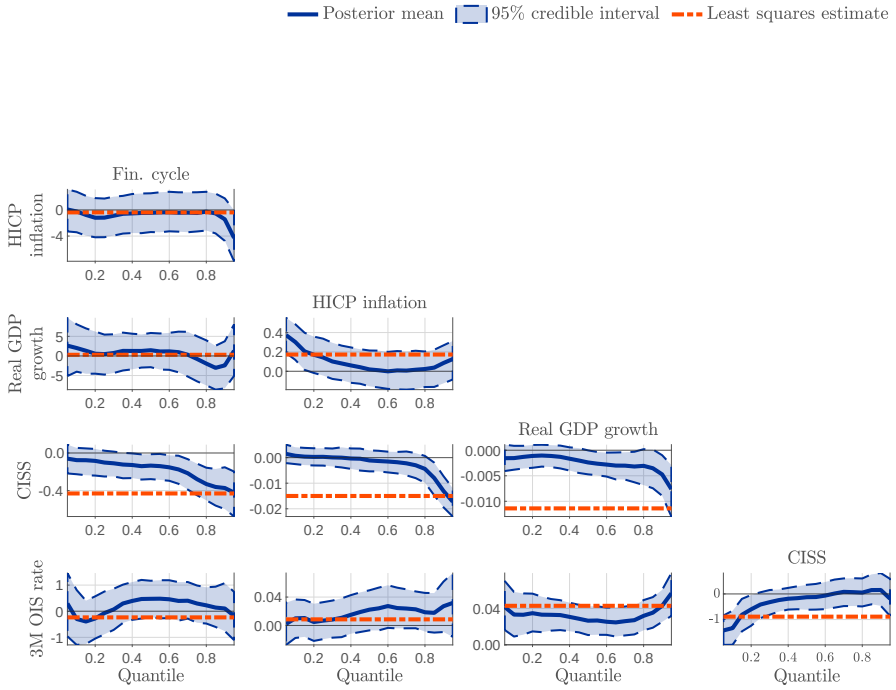

## Figure D.3: Posterior inference for the euro area, $A_1(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

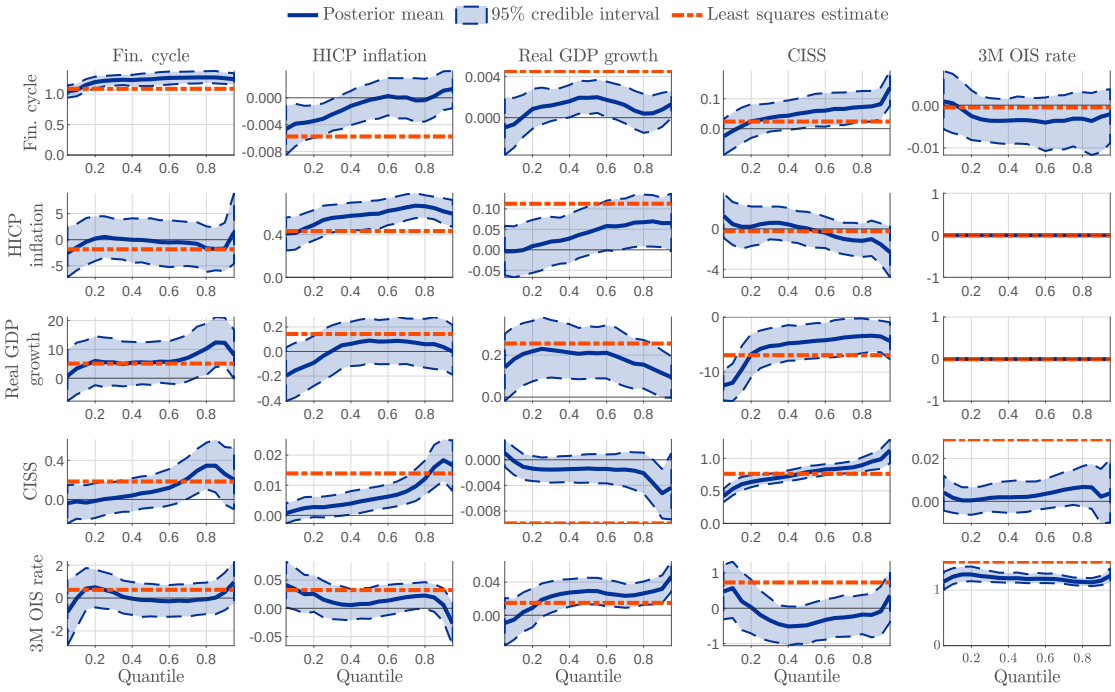

## Figure D.4: Posterior inference for the euro area, $A_2(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

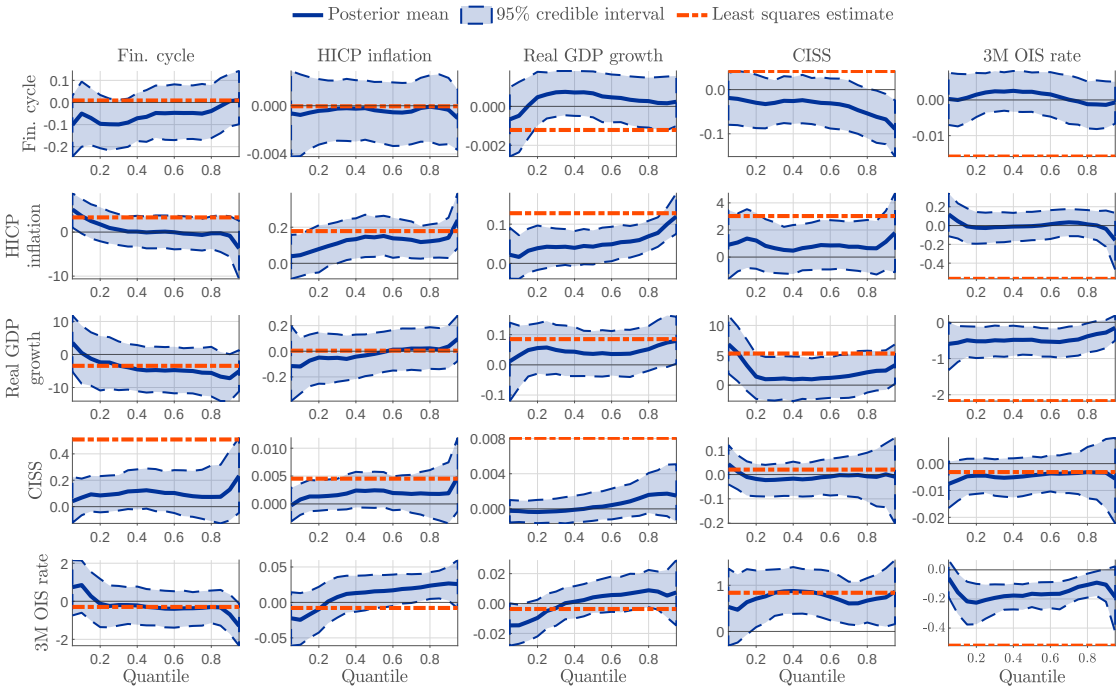

## Figure D.5: Posterior inference for the euro area, $A_3(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

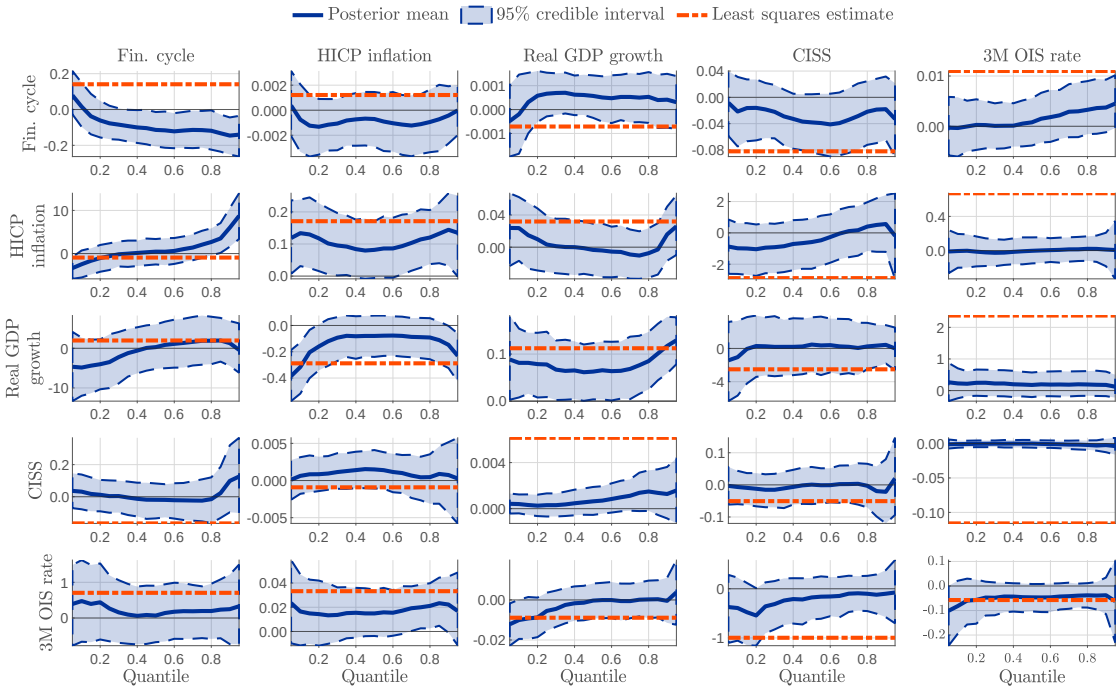

## Figure D.6: Posterior inference for the euro area, $A_4(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

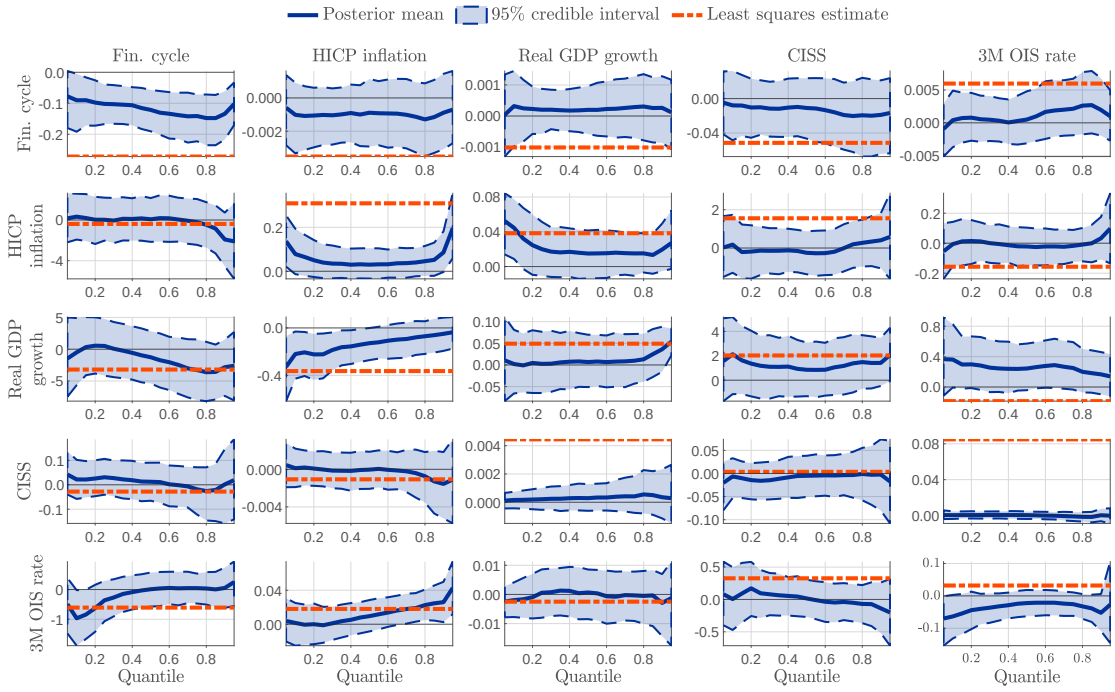

## Figure D.7: Posterior inference for the euro area exogenous variables, $C(\gamma)$

Posterior mean and 95% credible intervals are btained from 2,500 posterior draws. Least squares estimates are provided for comparison.

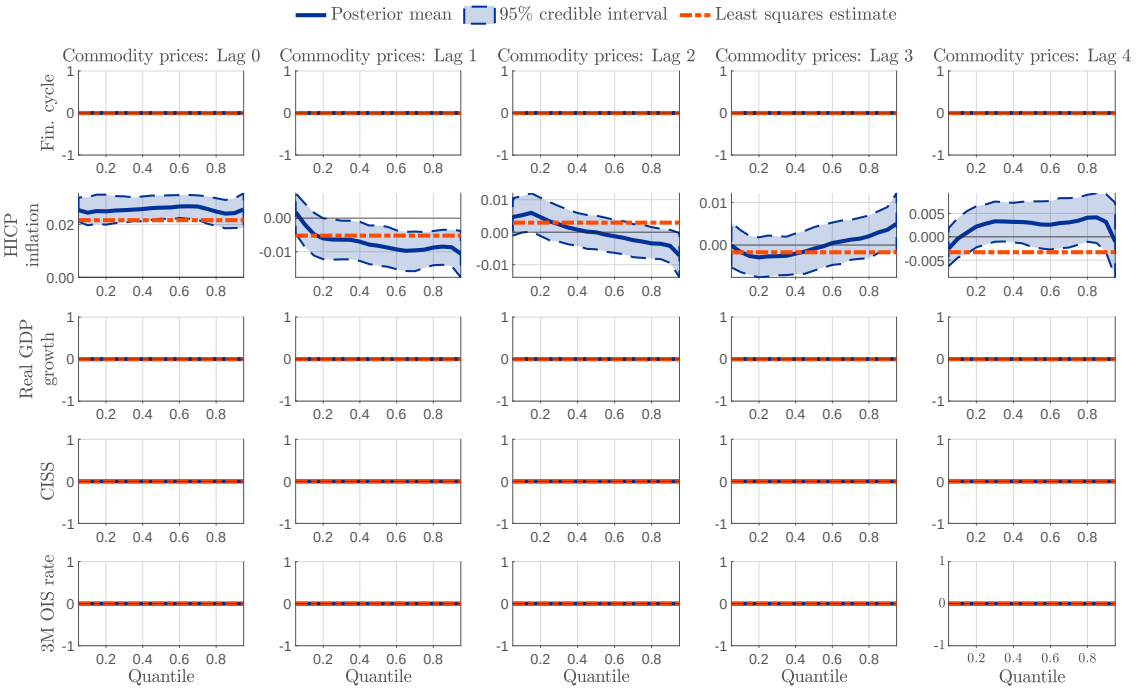

## Figure D.8: Posterior inference for the euro area dummy variables, $B(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

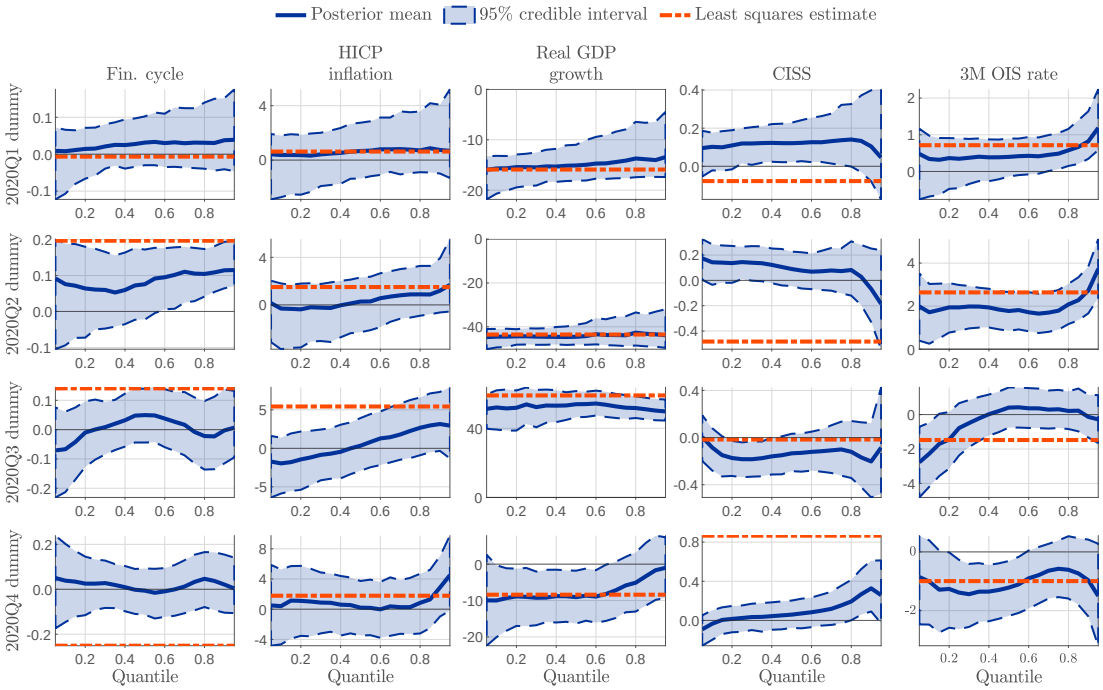


## Figure D.9: Posterior means for the euro area tightness parameter, $\lambda_i(\gamma)$

Posterior mean estimates of $\lambda_i(\gamma)$ are obtained from 2,500 posterior draws.

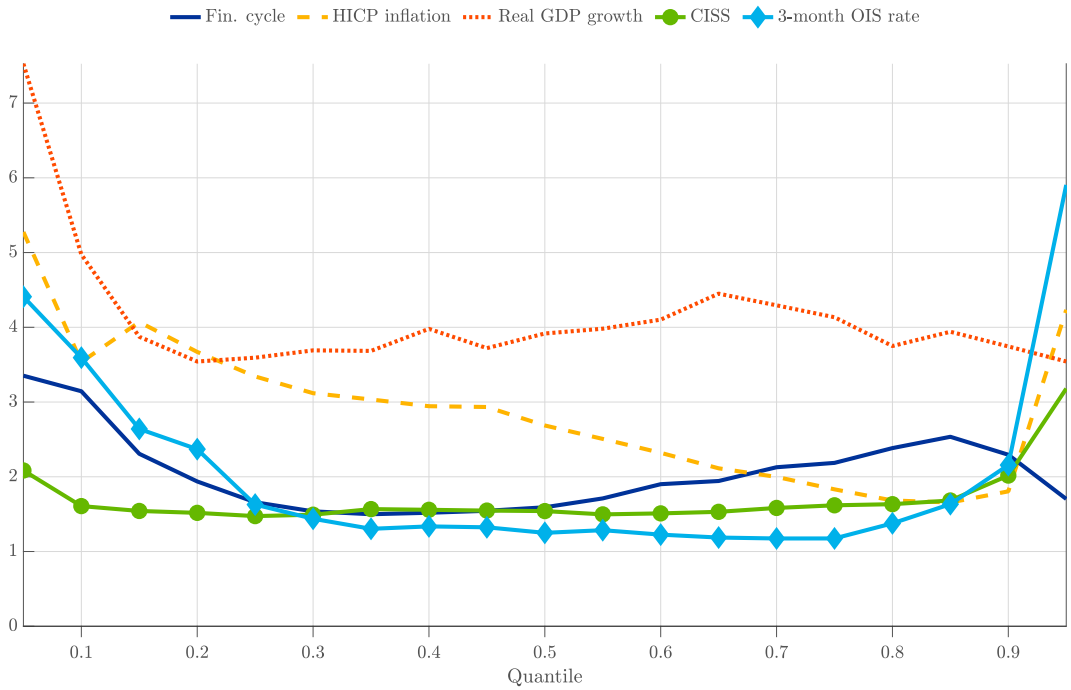

## Figure D.10: Quantile impulse response functions for the euro area

Impulse response functions implied by the posterior estimates reported in Figures D.1-D.8. Based on 400 draws from the posterior distribution and 20.000 forward simulations of the conditional distribution for each draw. Shocks are equal to the standard deviation of residuals from the median quantile regression of the respective variables. Variables are ordered financial cycle (respective first row), HICP inflation (second row), GDP growth (third row), CISS (fourth row), and the 3-month OIS rate (fifth row). Credible intervals are dashed and at a 95% level. Estimation sample is 1990Q1 to 2022Q4.

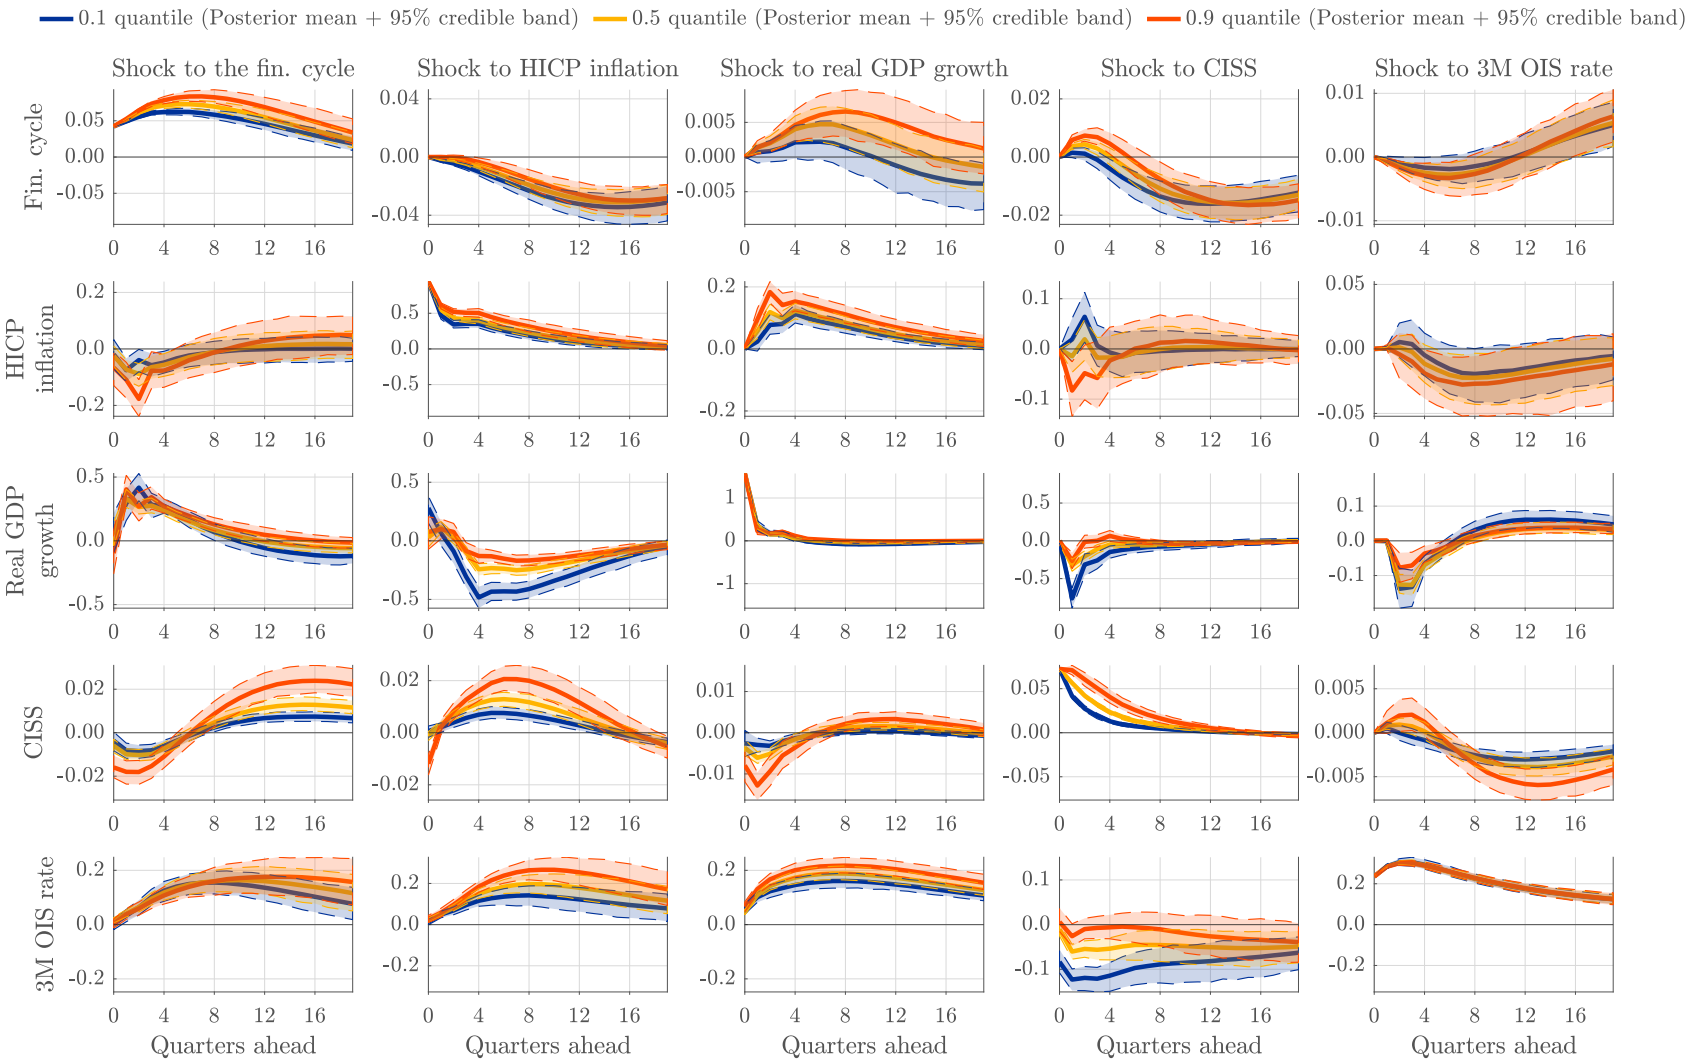

# E U.S. posterior estimates

**Figure E.1: Posterior inference for the U.S. , $\omega(\gamma)$**

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates (in red) for the conditional mean are provided for comparison.

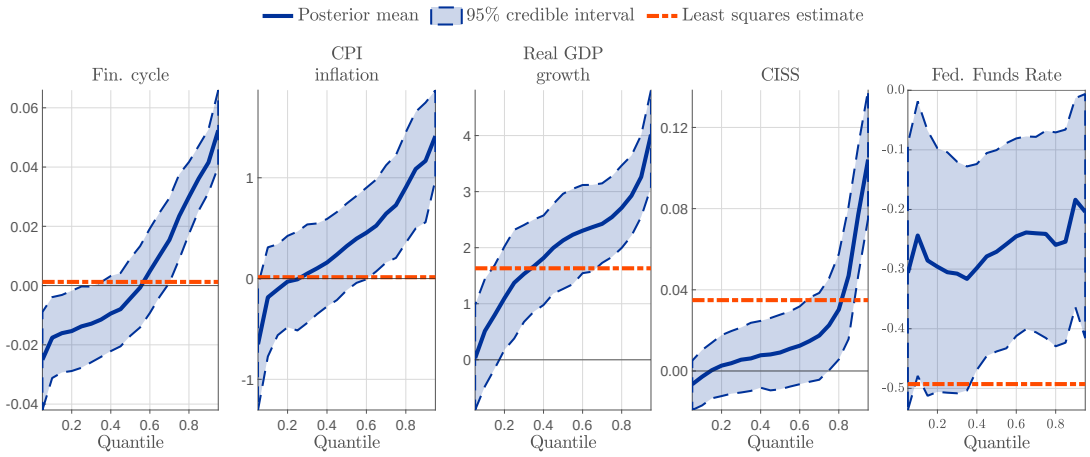


**Figure E.2: Posterior inference for the U.S., $A_0(\gamma)$**

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates for the conditional mean are provided for comparison.

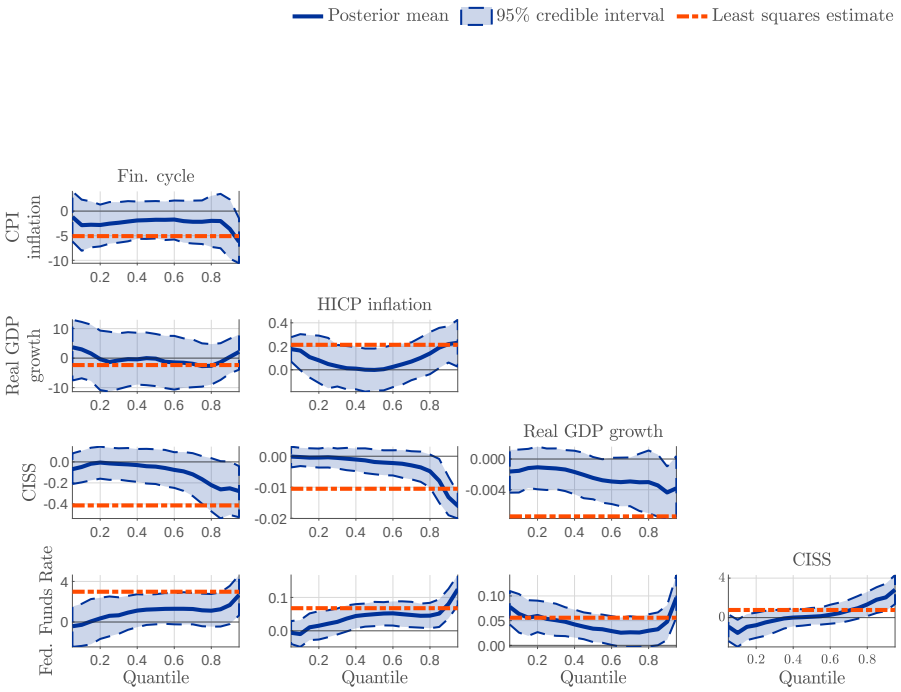

## Figure E.3: Posterior inference for the U.S., $A_1(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are for comparison.

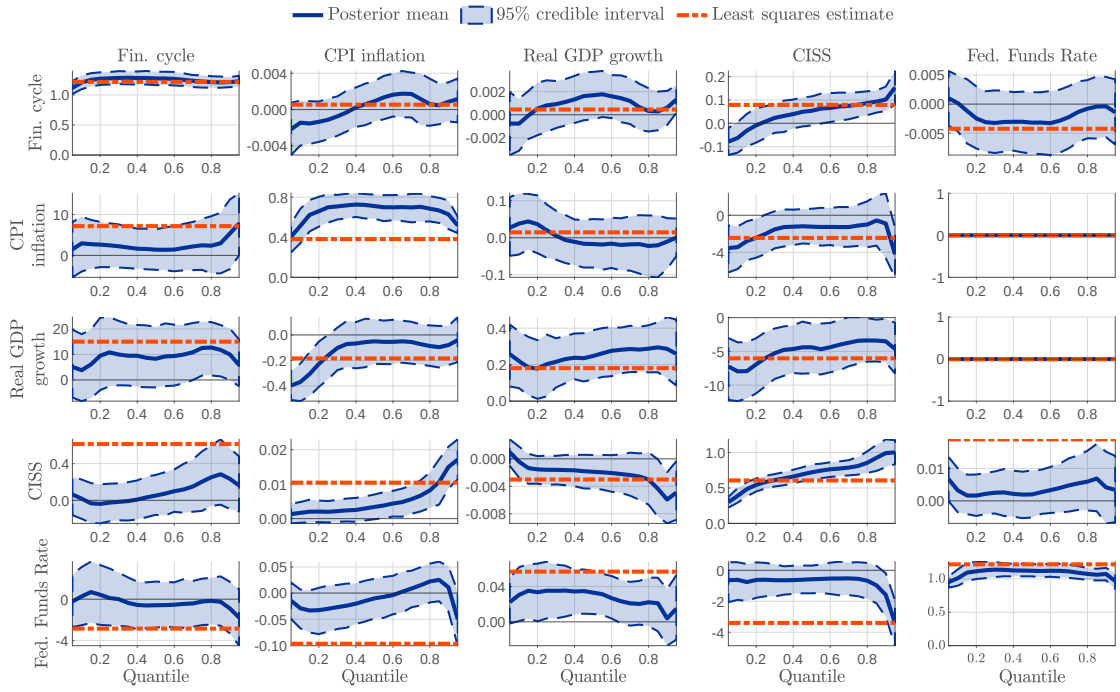

## Figure E.4: Posterior inference for the U.S., $A_2(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are for comparison.

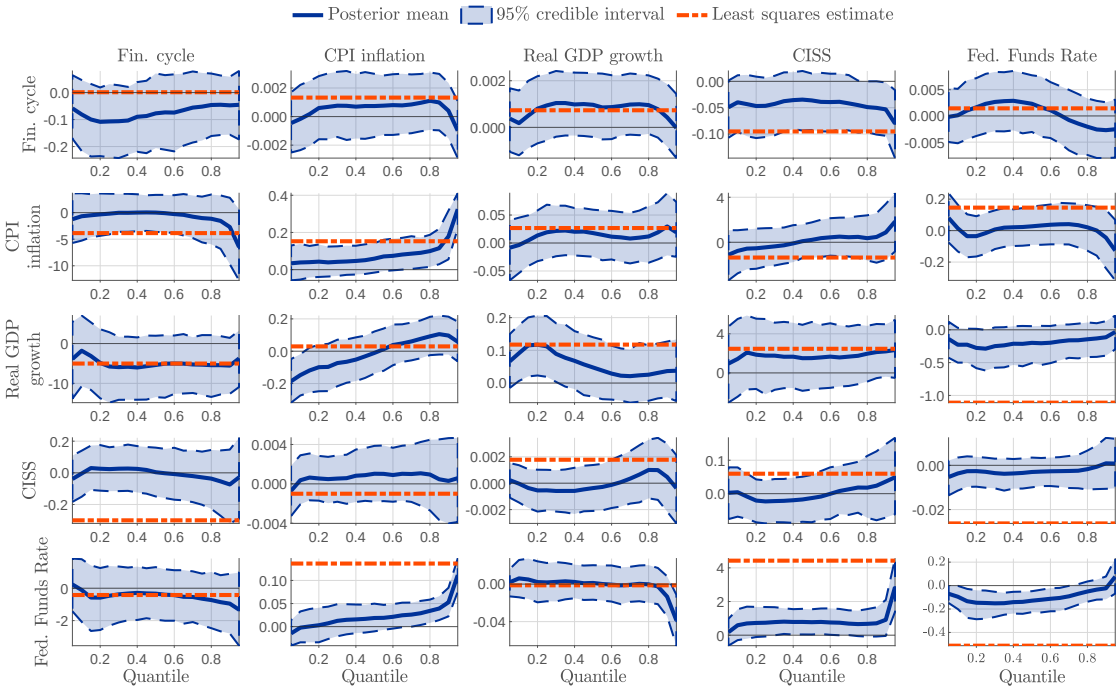

## Figure E.5: Posterior inference for the U.S., $A_3(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

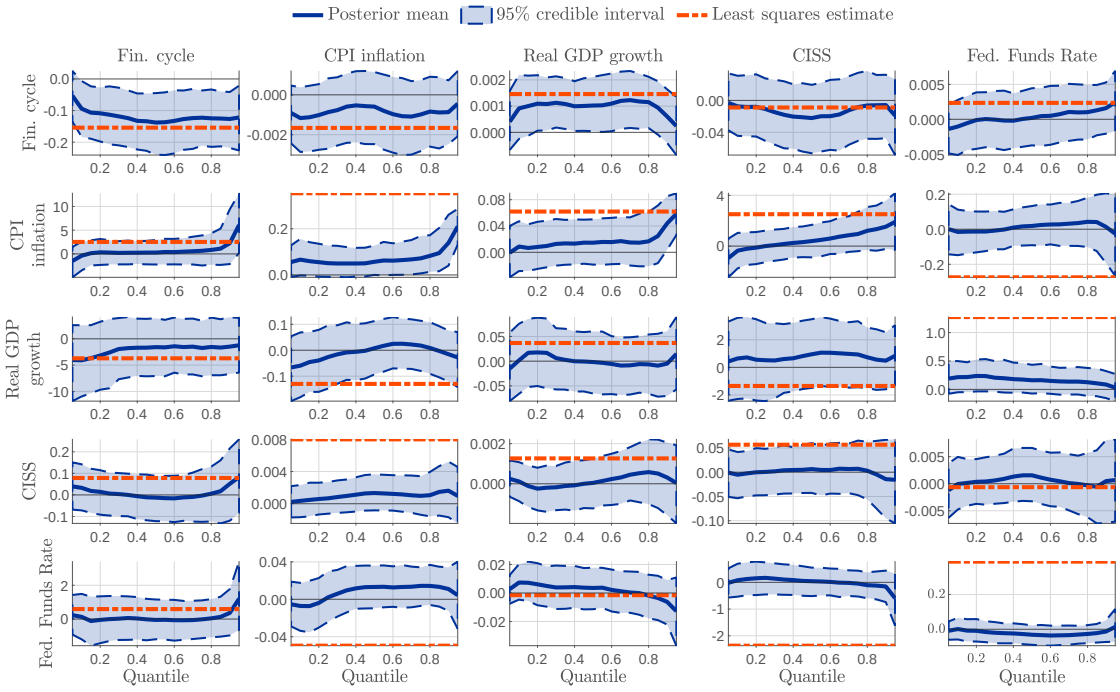

## Figure E.6: Posterior inference for the U.S., $A_4(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

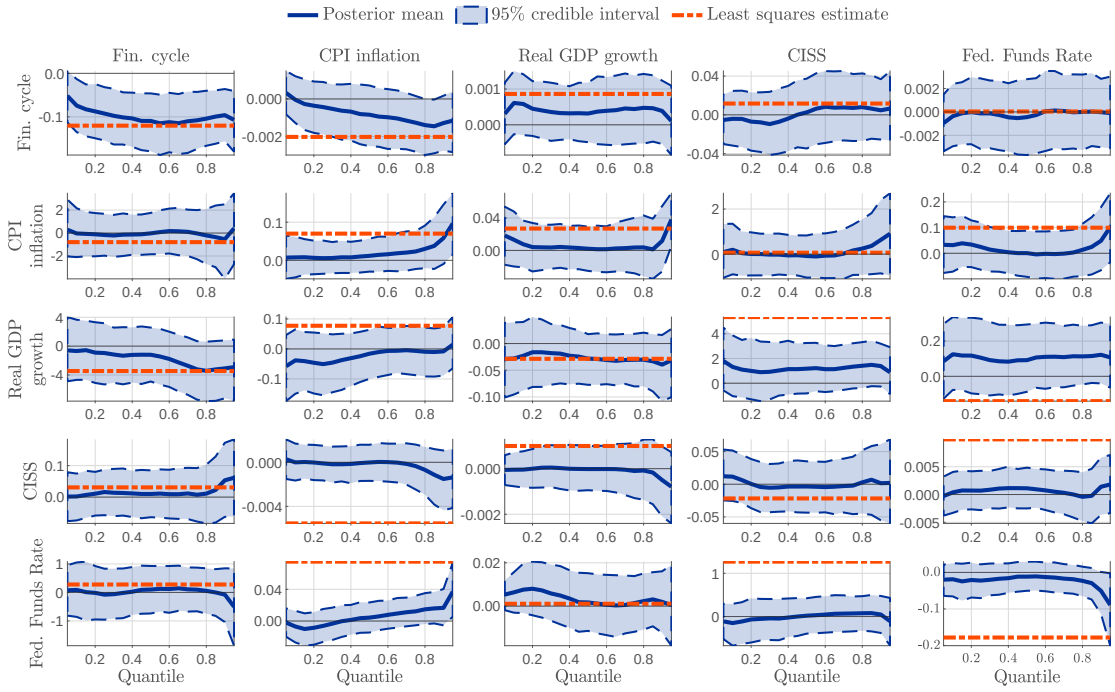

## Figure E.7: Posterior inference for the U.S. exogenous variables, $C(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

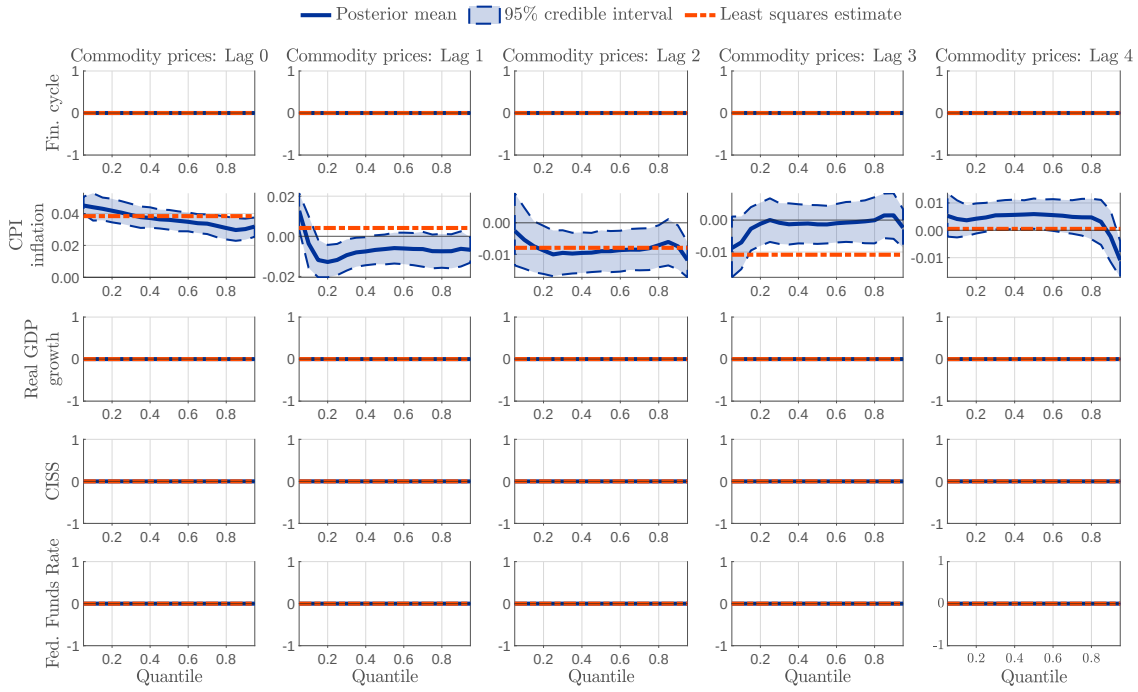

## Figure E.8: Posterior inference for the U.S. dummy variables, $B(\gamma)$

Posterior mean and 95% credible intervals are obtained from 2,500 posterior draws. Least squares estimates are provided for comparison.

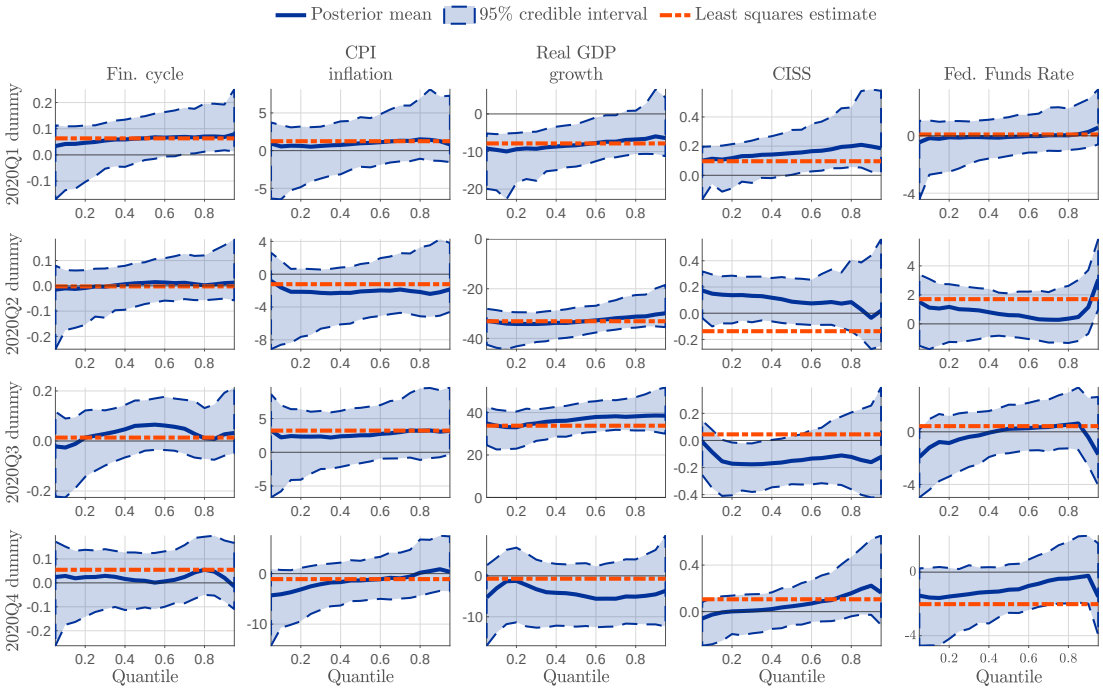


## Figure E.9: Posterior mean estimates for U.S. prior tightness parameter $\lambda_i(\gamma)$

Posterior mean estimates of $\lambda_i(\gamma)$ are obtained from 2,500 posterior draws.

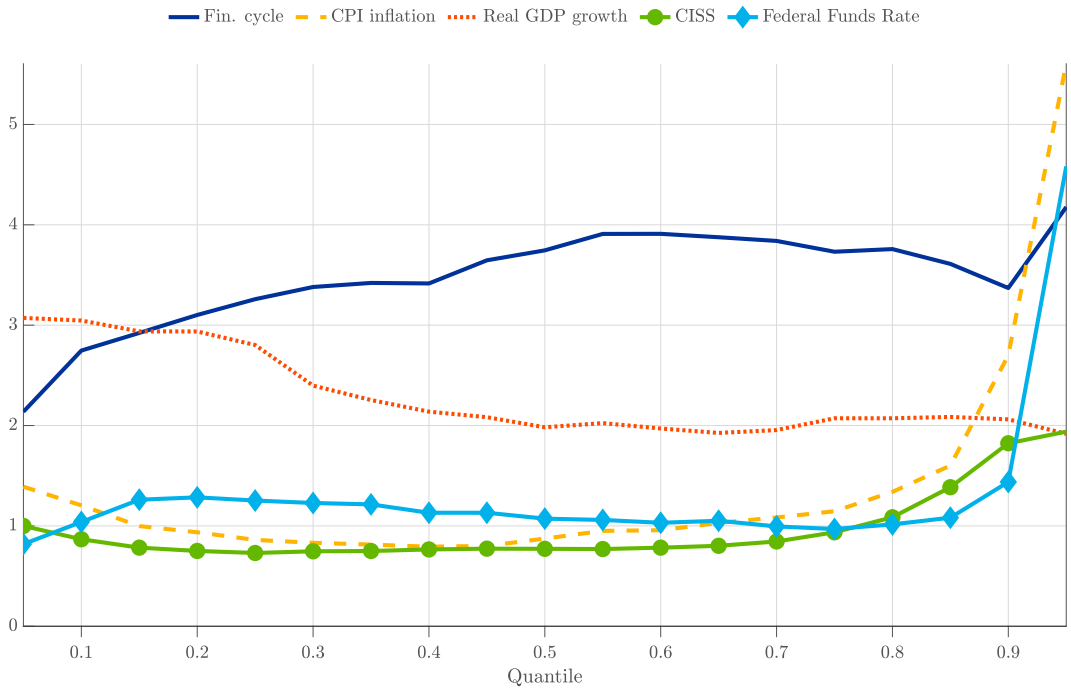

## Figure E.10: Quantile impulse response functions for the U.S.

Impulse response functions implied by the posterior estimates reported in Figures E.1-E.8. Based on 400 draws from the posterior distribution and 20.000 forward simulations of the conditional distribution for each draw. Shocks are equal to the standard deviation of residuals from the median quantile regression of the respective variables. Variables are ordered financial cycle (respective first row), CPI inflation (second row), GDP growth (third row), CISS (fourth row), and the Federal funds rate (fifth row). Credible intervals are dashed and at a 95% level. Estimation sample is 1973Q1 to 2022Q4.

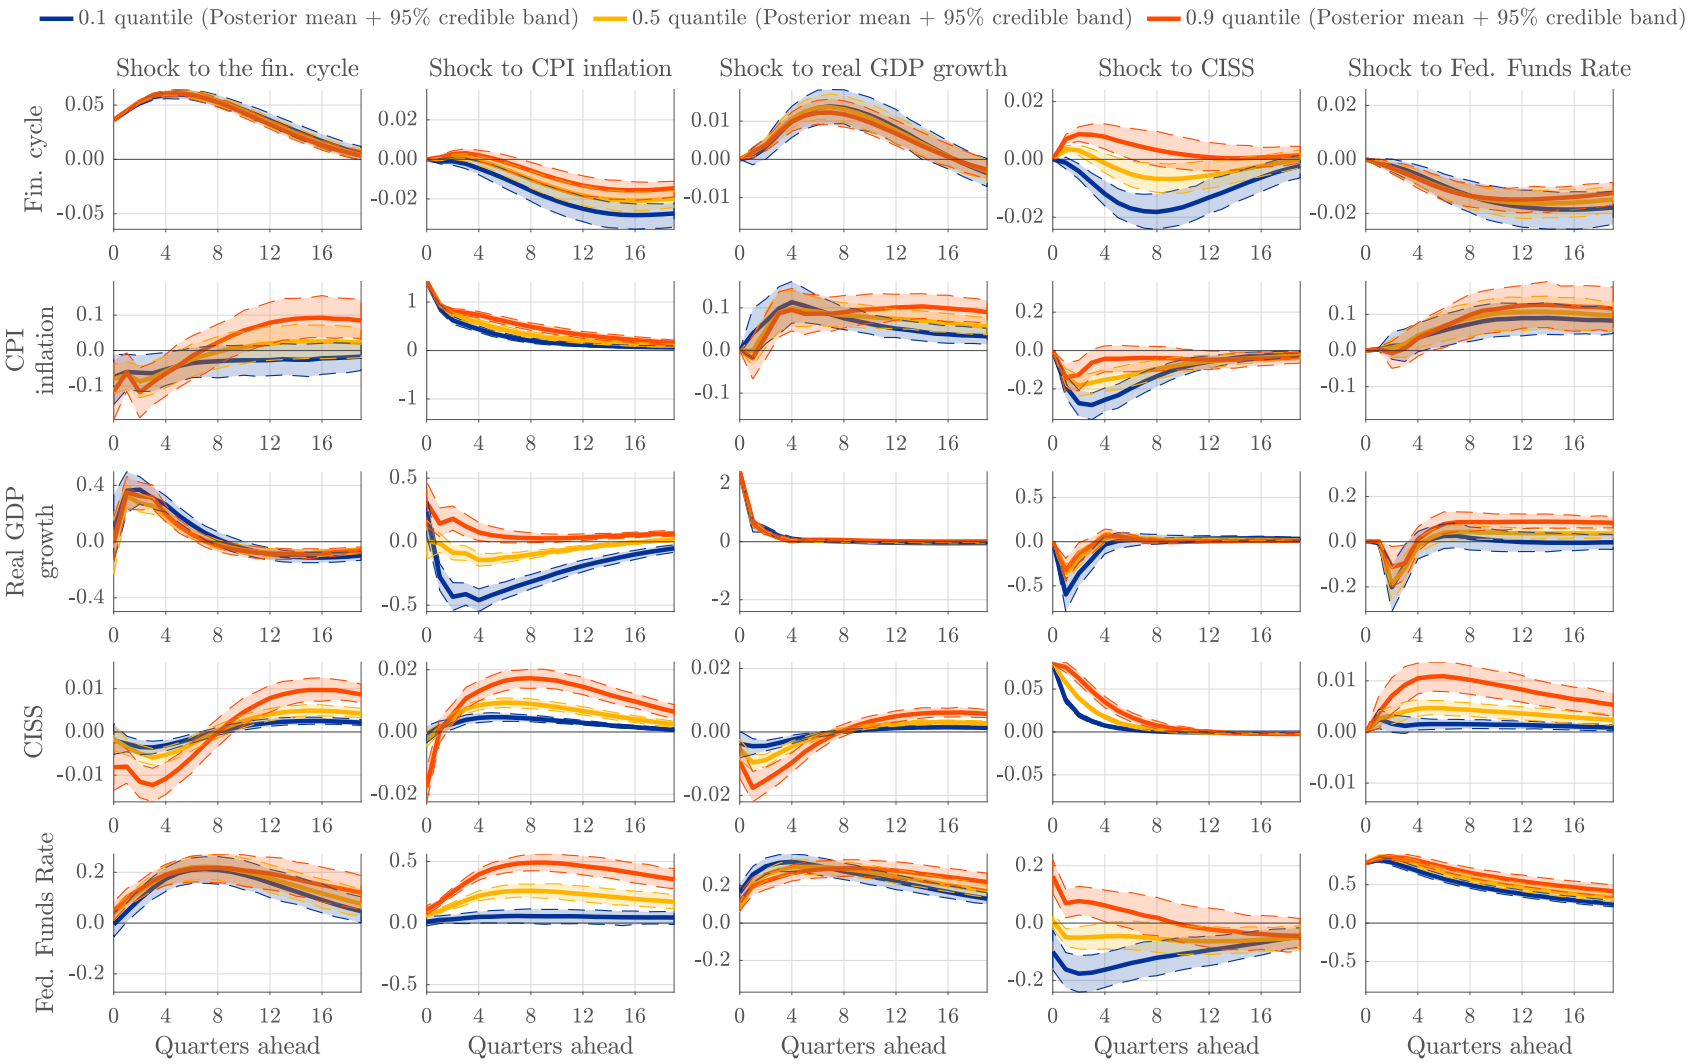

# References


Chavleishvili, S. and S. Manganelli (2023). Forecasting and stress testing with quantile vector autoregression. *Journal of Applied Econometrics, forthcoming*.

Engle, R. F. and S. Manganelli (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics 22*(4), 367–381.

Giannone, D., M. Lenza, and G. E. Primiceri (2015). Prior selection for vector autoregression. *The Review of Economics and Statistics 97*(2), 436–451.

Huber, F. and M. Feldkircher (2019). Adaptive shrinkage in bayesian vector autoregressive models. *Journal of Business & Economic Statistics 37*(1), 27–39.

Ibrahim, J. G. and M.-H. Chen (2000). Power prior distributions for regression models. *Statistical Science 15*(1), 46–60.

Ibrahim, J. G., M.-H. Chen, Y. Gwon, and F. Chen (2015). The power prior: Theory and applications. *Statistics in Medicine 34*(28), 3724–3749.

Khare, K. and J. P. Hobert (2012). Geometric ergodicity of the Gibbs sampler for Bayesian quantile regression. *Journal of Multivariate Analysis 112*, 108–116.

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

Korobilis, D. (2017). Quantile regression forecasts of inflation under model uncertainty. *International Journal of Forecasting 23*, 11 – 20.

Korobilis, D. and D. Pettenuzzo (2019). Adaptive hierarchical priors for high-dimenstional vector autoregressions. *Journal of Econometrics 212*(1), 241–271.

Kozumi, H. and G. Kobayashi (2011). Gibbs sampling methods for Bayesian quantile regression. *Journal of Statistical Computation and Simulation 81*(11), 1565–1578.

Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions - five years of experience. *Journal of Business & Economic Statistics 4*, 25–38.

Luetkepohl, H. (2005). *New introduction to multiple time series analysis*. Springer, Germany.

Yu, K. and R. A. Moyeed (2001). Bayesian quantile regression. *Statistics & Probability Letters 54*, 437–447.