

INTERNET APPENDIX

MCMC FOR QUANTILE BAYESIAN PREDICTIVE SYNTHESIS

Jiaxun Liu

October, 2026

IA.1 Overview and Initialization

This internet appendix provides the comprehensive algorithmic details of the Markov Chain Monte Carlo (MCMC) computations for the structurally stabilized Quantile Bayesian Predictive Synthesis (QBPS) model. The estimation relies on a customized block Gibbs sampler: the first block samples the dynamic synthesis states and parameters via Forward-Filtering Backward-Sampling (FFBS), the second samples the latent agent states, and the third samples the latent Asymmetric Laplace (AL) shock variables.

IA.1.1 QBPS Structure

To guarantee numerical stability and retain sequential variance discounting, the model utilizes a decoupled AL representation. The Quantile Dynamic Linear Model (DLM) is defined by the following hierarchical representation:

$$\begin{aligned} y_t | \boldsymbol{\theta}_t^\tau, e_t, v_t, \tau &\sim \mathcal{N}(\mathbf{F}_t' \boldsymbol{\theta}_t^\tau + a_\tau \sqrt{v_t} e_t, b_\tau v_t e_t) \\ \boldsymbol{\theta}_t^\tau | \boldsymbol{\theta}_{t-1}^\tau, \mathbf{W}_t, v_t &\sim \mathcal{N}(\boldsymbol{\theta}_{t-1}^\tau, v_t \mathbf{W}_t), \quad \mathbf{W}_t = \frac{1 - \delta_t}{\delta_t} \cdot \mathbf{C}_{t-1} \\ e_t &\sim \text{Exp}(1) \\ v_t | \mathcal{D}_{t-1} &\sim \text{Inverse-Gamma}(\beta n_{t-1}/2, \beta n_{t-1} s_{t-1}/2) \end{aligned} \tag{IA.1}$$

where $a_\tau = \frac{1-2\tau}{\tau(1-\tau)}$ and $b_\tau = \frac{2}{\tau(1-\tau)}$ are constants depending solely on the target quantile τ . Notice that the latent shock e_t follows a Standard Exponential distribution and is fully decoupled from the time-varying volatility v_t .

For notational simplicity, the superscript τ on the state vector $\boldsymbol{\theta}_t$ is omitted in the remainder of this document.

Standing at time T , the decision maker conditions on historical information $\mathcal{D}_{1:T} = \{y_{1:T}, \mathcal{H}_{1:T}\}$. The model is initialized with priors $\boldsymbol{\theta}_0 \sim \mathcal{N}(\mathbf{m}_0, \mathbf{C}_0 v_0/s_0)$ and $v_0^{-1} \sim \text{Gamma}(n_0/2, n_0 s_0/2)$, where $\mathbf{m}_0$ and $\mathbf{C}_0$ denote the initial state mean and unscaled covariance matrix, while n_0 and s_0 govern the initial degrees of freedom and structural variance estimate.

Bifurcated Prior Initialization: To structurally accommodate the non-parametric inertia of bounded, tree-based agents (e.g., Quantile Random Forests), the prior initialization of the state covariance matrix ($\mathbf{C}_0$) is bifurcated. The dynamic weights assigned to candidate agents are anchored with a tight precision variance (e.g., $C_{k,k} = 0.01$) to mathematically crush the initial Kalman gain, preventing the weights from wildly overfitting the extreme noise inherent in the AL tail variance. Conversely, the unconstrained baseline intercept is initialized with a significantly looser prior variance (e.g., $C_{1,1} = 0.5$). This calibrated flexibility mathematically designates the intercept as the primary structural shock absorber, granting it the immediate parametric freedom to plunge and absorb unprecedented level shifts that fall strictly outside the candidate agents' historical boundaries.

State Setup: The design matrix is initialized as $\mathbf{F}_t = (1, x_{t1}, \dots, x_{tJ})'$ using draws from the individual agent priors $x_{tj} \sim h_{tj}(x_{tj})$. The latent Standard Exponential shocks are initialized natively as $e_t \sim \text{Exp}(1)$. The observation-driven asymmetric discounting module is initialized at the square of the average training sample check-loss ($S_0 = \overline{\text{check} - \text{loss}_{in-sample}^2}$) with the baseline memory ceiling ($\delta_0 = \delta_{max}$).

IA.1.2 Observation-Driven Asymmetric Discounting Module

Applying a strict, static discount factor (e.g., $\delta_{max} = 0.995$) at extreme quantiles causes the FFBS algorithm to fail the Stability-Plasticity dilemma during structural breaks. To allow the model to rapidly adapt to crises without inducing MCMC thrashing, the framework utilizes an endogenous observation-driven asymmetric discounting module.

The module is mathematically anchored in the Generalized Autoregressive Score (GAS) framework, operating as a formally scaled observation-driven update governed by the gradient of the marginal Asymmetric Laplace (AL) log-likelihood. The mechanism first evaluates the out-of-sample forecast surprise via Z_t , the scaled check-loss, which serves as the sufficient statistic for local model failure. Using the prior state mean $\mathbf{m}_{t-1}$ and prior structural variance s_{t-1} , the standardized surprise is computed as:

$$Z_t = \frac{\rho_\tau(y_t - \mathbf{F}'_t \mathbf{m}_{t-1})}{\sqrt{s_{t-1}}} \quad (\text{IA.2})$$

where $\rho_\tau(\cdot)$ is the standard asymmetric check-loss function.

The localized surprises are then smoothed into a continuous latent state variable, S_t , which measures sustained macro-financial stress through a formal GAS autoregressive recursion:

$$S_t = \rho S_{t-1} + (1 - \rho) Z_t \quad (\text{IA.3})$$

To properly capture the macro-financial reality that crises detonate instantaneously but calm returns persistently, the asymmetric discounting module utilizes a path-dependent bounding link function. To prevent circular endogeneity within the likelihood evaluation, the variance regulator δ_t is determined strictly by the lagged stress state S_{t-1} .

First, the mechanism computes a provisional target level, δ_t^* , which maps the severity of the

structural stress directly to a required drop from the baseline stability ceiling δ_{max} :

$$\delta_t^* = \begin{cases} \max(\delta_{min}, \delta_{max} - \kappa_{drop}(S_{t-1} - \bar{S})) & \text{if } S_{t-1} \geq \bar{S} \\ \delta_{max} & \text{otherwise} \end{cases} \quad (\text{IA.4})$$

To execute the actual synthesis update, the realized discount factor δ_t takes the minimum between this required target and a graduated recovery from the previous period's state:

$$\delta_t = \min(\delta_t^*, \delta_{t-1} + \kappa_{recover}) \quad (\text{IA.5})$$

Hyperparameter Calibration: While this structural framework is generalized, the formal GAS theoretical bounds and empirical calibration across diverse macro-financial applications dictate the following optimal baseline parameterization utilized in this paper: an autoregressive persistence (EMA smoothing factor) of $\rho = 0.3$, a panic activation threshold of $\bar{S} = 1.75$, an asymmetric drop penalty of $\kappa_{drop} = 0.15$, an asymmetric recovery micro-step of $\kappa_{recover} = 0.01$, and a stable memory domain bounded between $\delta_{min} = 0.950$ and $\delta_{max} = 0.995$.

IA.2 Sampling Steps in Each MCMC Iterate

Preamble: Conditional Gaussianity and Marginal AL Recovery

Before detailing the Gibbs blocks, it is critical to clarify that the exact Asymmetric Laplace (AL) likelihood is never evaluated conditionally during the sequential MCMC steps. At any single step within the sampler, if the filter conditions on a specifically drawn, frozen value of the latent standard exponential shock e_t , the observation likelihood is strictly Gaussian:

$$p(y_t | \mathbf{F}_t, \boldsymbol{\theta}_t^\tau, v_t, e_t) = \mathcal{N}(\mathbf{F}_t' \boldsymbol{\theta}_t^\tau + a_\tau \sqrt{v_t} e_t, b_\tau v_t e_t) \quad (\text{IA.6})$$

Because this conditional space is perfectly Gaussian, the Forward-Filtering Backward-Sampling (FFBS) algorithm retains its exact linear-algebraic conjugacy, preventing computational bottlenecks. The target AL likelihood is inherently a *marginal* distribution. It does not exist at any isolated step of the algorithm; rather, it emerges only as the aggregate reality when this conditional Gaussian distribution is integrated across all possible values of e_t , weighted by its Exponential prior over the $[0, \infty)$ support. Because the marginalized likelihood is exactly AL, the marginal posterior distributions of the dynamic weights (θ_t^τ) are mathematically guaranteed to converge to the true AL posterior, ensuring the synthesis remains optimized for the targeted extreme quantile without ever subjecting the filter directly to the non-conjugate AL density.

IA.2.1 Step 1: Sampling BPS DLM parameters $\Phi_{1:T}$

Conditional on the latent agent states and AL shock variables, the parameters $\Phi_{1:T} = \{\theta_{1:T}, v_{1:T}\}$ are sampled using the FFBS algorithm augmented with observation-driven asymmetric discounting module on δ_t .

Forward filtering: The dynamic discount factor is initialized at the baseline stability ceiling, $\delta_1 = \delta_{max}$, and the GAS stress state S_1 is initialized dynamically using the first period's out-of-sample score. For each $t = 1 : T$, save the sequence of realized discount factors $\delta_{1:T}$, and perform the following sequential updates:

1. **Update to time t prior (Apply Current Delta):** Apply the strictly predetermined discount factor δ_t to inflate the state covariance:

$$\theta_t | v_t, \mathcal{D}_{t-1} \sim \mathcal{N}(\mathbf{m}_{t-1}, \mathbf{R}_t v_t / s_{t-1}) \quad (\text{IA.7})$$

where $\mathbf{R}_t = \mathbf{C}_{t-1} / \delta_t$, and $v_t^{-1} | \mathcal{D}_{t-1} \sim \text{Gamma}(\beta n_{t-1} / 2, \beta n_{t-1} s_{t-1} / 2)$.

2. **1-step predictive distribution:**

$$\begin{aligned}
V_{obs} &= b_\tau s_{t-1} e_t \\
f_t &= \mathbf{F}'_t \mathbf{m}_{t-1} + a_\tau \sqrt{s_{t-1}} e_t \\
q_t &= \mathbf{F}'_t \mathbf{R}_t \mathbf{F}_t + V_{obs}
\end{aligned} \tag{IA.8}$$

3. **Filtering update to time t posterior:** Based on the forecast error $e_{filter} = y_t - f_t$ and the Kalman gain $\mathbf{A}_t = \mathbf{R}_t \mathbf{F}_t / q_t$, update the moments utilizing the numerically robust Joseph Stabilized form:

$$\begin{aligned}
n_t &= \beta n_{t-1} + 1 \\
r_t &= (\beta n_{t-1} + e_{filter}^2 / q_t) / n_t \\
s_t &= r_t s_{t-1} \\
\mathbf{m}_t &= \mathbf{m}_{t-1} + \mathbf{A}_t e_{filter} \\
\mathbf{C}_t &= (\mathbf{I} - \mathbf{A}_t \mathbf{F}'_t) \mathbf{R}_t (\mathbf{I} - \mathbf{A}_t \mathbf{F}'_t)' + \mathbf{A}_t (V_{obs}) \mathbf{A}_t'
\end{aligned} \tag{IA.9}$$

4. **Observation-Driven Asymmetric Discounting (Dictating $t + 1$ Delta):**

Evaluate the out-of-sample surprise strictly based on prior information $\mathbf{m}_{t-1}$ to prevent endogeneity. Smooth the localized AL score into the stress state S_t :

$$Z_t = \frac{\rho_\tau(y_t - \mathbf{F}'_t \mathbf{m}_{t-1})}{\sqrt{s_{t-1}}} \tag{IA.10}$$

$$S_t = 0.3 S_{t-1} + 0.7 Z_t \tag{IA.11}$$

Define the target discount factor δ_{t+1}^* based on the activation threshold $\bar{S} = 1.75$ and penalty $\kappa_{drop} = 0.15$. Crucially, the drop is calculated relative to the stable baseline δ_{max} , not the current state:

$$\delta_{t+1}^* = \begin{cases} \max(\delta_{min}, \delta_{max} - 0.15(S_t - 1.75)) & \text{if } S_t > 1.75 \\ \delta_{max} & \text{otherwise} \end{cases} \tag{IA.12}$$

Finally, compute the next period's structural variance regulator δ_{t+1} using an asymmetric recovery step ($\kappa_{recover} = 0.01$):

$$\delta_{t+1} = \min(\delta_{t+1}^*, \delta_t + 0.01) \quad (\text{IA.13})$$

Backward sampling: Draw v_T^{-1} from $\text{Gamma}(n_T/2, n_T s_T/2)$ and $\boldsymbol{\theta}_T$ from $\mathcal{N}(\mathbf{m}_T, \mathbf{C}_T v_T/s_T)$. Recurse back over times $t = T - 1, \dots, 1$. *Crucially, the smoothed moments use the exact dynamically generated δ_{t+1} that was used to project from t to $t + 1$ during the forward pass:*

$$v_t^{-1} = \beta v_{t+1}^{-1} + \gamma_t, \quad \text{where } \gamma_t \sim \text{Gamma}((1 - \beta)n_t/2, n_t s_t/2) \quad (\text{IA.14})$$

$$\boldsymbol{\theta}_t | \boldsymbol{\theta}_{t+1}, v_t \sim \mathcal{N}(\mathbf{m}_t + \delta_{t+1}(\boldsymbol{\theta}_{t+1} - \mathbf{m}_t), \mathbf{C}_t(1 - \delta_{t+1})(v_t/s_t)) \quad (\text{IA.15})$$

IA.2.2 Step 2: Sampling the latent agent states $\mathbf{X}_{1:T}$

Isolate the candidate agent weights $\boldsymbol{\theta}_{t,1:J}$ and the intercept θ_{t0} from the sampled state vector $\boldsymbol{\theta}_t$. Construct $\mathbf{H}_t = \text{diag}(H_{t1}\phi_{t1}, \dots, H_{tJ}\phi_{tJ})$. The posterior distribution for the latent agent states $\mathbf{X}_t$ is derived as:

$$\mathbf{X}_t | \cdot \sim \mathcal{N}(\mathbf{h}_t + \mathbf{b}_t c_t, \mathbf{H}_t - \mathbf{b}_t \mathbf{b}_t' g_t) \quad (\text{IA.16})$$

where the conditioning components are defined as:

$$c_t = y_t - (\theta_{t0} + \mathbf{h}_t' \boldsymbol{\theta}_{t,1:J} + a_\tau \sqrt{v_t} e_t) \quad (\text{IA.17})$$

$$g_t = b_\tau v_t e_t + \boldsymbol{\theta}_{t,1:J}' \mathbf{H}_t \boldsymbol{\theta}_{t,1:J} \quad (\text{IA.18})$$

$$\mathbf{b}_t = \mathbf{H}_t \boldsymbol{\theta}_{t,1:J} / g_t \quad (\text{IA.19})$$

Conditional on the newly sampled x_{tj} , updated scales ϕ_{tj} are drawn via:

$$\phi_{tj} | \cdot \sim \text{Inverse-Gamma}\left(\frac{n_{tj} + 1}{2}, \frac{n_{tj} + (x_{tj} - h_{tj})^2 / H_{tj}}{2}\right) \quad (\text{IA.20})$$

IA.2.3 Step 3: Sampling the AL latent variables $e_{1:T}$

The latent shock variables are sampled using an exact block update from the Generalized Inverse Gaussian (GIG) distribution:

$$e_t | \cdot \sim \mathcal{GIG} \left(\lambda = 0.5, \chi = \frac{(y_t - \mathbf{F}_t' \boldsymbol{\theta}_t)^2}{b_\tau v_t}, \psi = \frac{a_\tau^2}{b_\tau} + 2 \right) \quad (\text{IA.21})$$

IA.3 On Scaled Parametrization

IA.3.1 Prior Specification and Conditional AL Equivalence

In static Bayesian quantile regression, such as the framework proposed by Kozumi and Kobayashi (2011), it is customary to place an Inverse-Gamma prior directly on the scale parameter σ_t . However, transitioning this scale-parameterization into a dynamic setting introduces severe computational bottlenecks. To achieve analytical sequential variance discounting within the FFBS algorithm, Normal-Gamma conjugacy is strictly required. Therefore, the proposed framework places an Inverse-Gamma prior directly on the observation variance v_t (implying a Gamma prior on the precision v_t^{-1}).

While this constitutes a deliberate departure from the unconditional prior structure of Kozumi and Kobayashi (2011)—effectively placing an Inverse-Nakagami prior on the underlying scale—it mathematically preserves the Asymmetric Laplace (AL) representation marginally. Conditional on the drawn variance v_t , integrating out the decoupled Standard Exponential shock $e_t \sim \text{Exp}(1)$ perfectly recovers the AL likelihood:

$$p(y_t | \mu_t, v_t, \tau) = \int_0^\infty \mathcal{N}(y_t | \mu_t + a_\tau \sqrt{v_t} e_t, b_\tau v_t e_t) \text{Exp}(e_t | 1) de_t \equiv \mathcal{AL}(y_t | \mu_t, \sqrt{v_t}, \tau) \quad (\text{IA.22})$$

Thus, the Gibbs sampler correctly targets the quantile surface while retaining the computational speed of closed-form DLM updates.

IA.3.2 Exact Prior Matching via Independence Metropolis-Hastings

Should a practitioner require exact matching to the static scale prior of Kozumi and Kobayashi (2011), the analytical variance discounting in the FFBS algorithm must be abandoned. If $\sigma_t \sim \text{Inverse-Gamma}(\alpha, \beta)$, the variance $v_t = \sigma_t^2$ is implied to follow a squared Inverse-Gamma distribution. Because this non-standard distribution breaks Normal-Gamma conjugacy, v_t must be isolated and sampled using an Independence Metropolis-Hastings (MH) step.

In this alternative block, a candidate variance is proposed directly from the exact prior by drawing $\sigma^* \sim \text{Inverse-Gamma}(\alpha, \beta)$ and setting $v^* = (\sigma^*)^2$. Because the proposal distribution $q(v)$ perfectly mirrors the prior distribution $p(v)$, they mathematically cancel in the MH acceptance ratio. The acceptance probability collapses entirely to the ratio of the conditional likelihoods:

$$r = \frac{\mathcal{N}(y_t | \mu_t + a_\tau \sqrt{v^*} e_t, b_\tau v^* e_t)}{\mathcal{N}(y_t | \mu_t + a_\tau \sqrt{v_{old}} e_t, b_\tau v_{old} e_t)} \quad (\text{IA.23})$$

While mathematically elegant, this exact-matching MH approach evaluates T likelihood ratios per MCMC iteration and is highly susceptible to poor mixing if the uninformative static prior severely disagrees with the dynamic posterior. Consequently, the conjugate Inverse-Gamma prior on v_t remains the recommended default for dynamic synthesis.

IA.3.3 Log-Scale Random Walk Metropolis-Hastings

While the Independence Sampler provides a mathematically elegant way to match the exact static prior, it is highly susceptible to poor MCMC mixing if the uninformative prior proposes values far from the high-density regions of the posterior. To ensure healthier local exploration of the parameter space while still permitting non-conjugate prior specifications, a standard Random Walk Metropolis-Hastings (RWMH) step can be employed.

Because the variance must strictly remain positive ($v_t > 0$), the random walk is formulated on the logarithmic scale. A candidate variance v^* is proposed via $\log(v^*) \sim \mathcal{N}(\log(v_{old}), s_{prop}^2)$,

where v_{old} is the variance from the previous MCMC iteration and s_{prop}^2 is a tunable proposal variance.

Unlike the Independence Sampler, the prior and proposal distributions do not cancel out. Furthermore, because the symmetric random walk occurs on the log scale, a Jacobian adjustment must be included when evaluating the acceptance ratio on the original scale. The acceptance probability $\min(1, r)$ is driven by the following ratio:

$$r = \frac{\mathcal{N}(y_t | \mu_t + a_\tau \sqrt{v^*} e_t, b_\tau v^* e_t) p(v^*) v^*}{\mathcal{N}(y_t | \mu_t + a_\tau \sqrt{v_{old}} e_t, b_\tau v_{old} e_t) p(v_{old}) v_{old}} \quad (\text{IA.24})$$

where $p(\cdot)$ evaluates the chosen prior density (e.g., the exact squared Inverse-Gamma), and the ratio v^*/v_{old} accounts for the Jacobian of the log transformation.

This RWMH approach guarantees much better acceptance rates than the Independence Sampler by exploring the localized posterior surface. However, it introduces the practical burden of tuning the proposal variance s_{prop}^2 to achieve optimal acceptance rates (typically between 20% and 40%), and it still requires evaluating likelihood ratios sequentially, making it computationally slower to the fully conjugate Inverse-Gamma FFBS framework.